

User's Reference
Volume II

AUTODESK®
3DS MAX®
9

Autodesk®

© 2007 Autodesk, Inc. All rights reserved.

This publication, or parts thereof, may not be reproduced in any form, by any method, for any purpose.
AUTODESK, INC., MAKES NO WARRANTY, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE REGARDING THESE MATERIALS, AND MAKES SUCH MATERIALS AVAILABLE SOLELY ON AN "AS-IS" BASIS. IN NO EVENT SHALL AUTODESK, INC., BE LIABLE TO ANYONE FOR SPECIAL, COLLATERAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH OR ARISING OUT OF PURCHASE OR USE OF THESE MATERIALS. THE SOLE AND EXCLUSIVE LIABILITY TO AUTODESK, INC., REGARDLESS OF THE FORM OF ACTION, SHALL NOT EXCEED THE PURCHASE PRICE OF THE MATERIALS DESCRIBED HEREIN.

Autodesk, Inc., reserves the right to revise and improve its products as it sees fit. This publication describes the state of this product at the time of its publication, and may not reflect the product at all times in the future.

The following are registered trademarks or trademarks of Autodesk, Inc., in the USA and other countries: 3DEC (design/logo), 3December, 3December.com, 3ds Max, ActiveShapes, Actrix, ADI, Alias, Alias (swirl design/logo), AliasStudio, Alias|Wavefront (design/logo), ATC, AUGI, AutoCAD, AutoCAD Learning Assistance, AutoCAD LT, AutoCAD Simulator, AutoCAD SQL Extension, AutoCAD SQL Interface, Autodesk, Autodesk Envision, Autodesk Insight, Autodesk Intent, Autodesk Inventor, Autodesk Map, Autodesk MapGuide, Autodesk Streamline, AutoLISP, AutoSnap, AutoSketch, AutoTrack, Backdraft, Built with ObjectARX (logo), Burn, Buzzsaw, CAiCE, Can You Imagine, Character Studio, Cinestream, Civil 3D, Cleaner, Cleaner Central, ClearScale, Colour Warper, Combustion, Communication Specification, Constructware, Content Explorer, Create>what's>Next> (design/logo), Dancing Baby (image), DesignCenter, Design Doctor, Designer's Toolkit, DesignKids, DesignProf, DesignServer, DesignStudio, Design|Studio (design/logo), Design Your World, Design Your World (design/logo), DWF, DWG, DWG (logo), DWG TrueConvert, DWG TrueView, DXF, EditDV, Education by Design, Extending the Design Team, FBX, Filmbox, FMDesktop, GDX Driver, Gmax, Heads-up Design, Heidi, HOOPS, HumanIK, i-drop, iMOUT, Incinerator, IntroDV, Kaydara, Kaydara (design/logo), LocationLogic, Lustre, Maya, Mechanical Desktop, MotionBuilder, ObjectARX, ObjectDBX, Open Reality, PolarSnap, PortfolioWall, Powered with Autodesk Technology, Productstream, ProjectPoint, Reactor, RealDWG, Real-time Roto, Render Queue, Revit, Showcase, SketchBook, StudioTools, Topobase, Toxik, Visual, Visual Bridge, Visual Construction, Visual Drainage, Visual Hydro, Visual Landscape, Visual Roads, Visual Survey, Visual Syllabus, Visual Toolbox, Visual Tugboat, Visual LISP, Voice Reality, Volo, and Wiretap.

The following are registered trademarks or trademarks of Autodesk Canada Co. in the USA and/or Canada and other countries: Backburner, Discreet, Fire, Flame, Flint, Frost, Inferno, Multi-Master Editing, River, Smoke, Sparks, Stone, Wire.

clothfx is a trademark of Size8 Software, Inc. mental ray is a registered trademark of mental images GmbH licensed for use by Autodesk, Inc. RE:Flex is a trademark of RE:Vision Effects, Inc. Intel is a registered trademark and the Intel Optimizer Logo is a trademark of Intel Corporation, used under license. Havok is a trademark or registered trademark of Havok.com, Inc. or its affiliates in the United States and/or in other countries. All other brand names, product names, or trademarks belong to their respective holders.

Third-Party Software Credits and Attributions

OpenEXR Bitmap I/O Plugin © 2003-2005 SplutterFish, LLC.
OpenEXR © 2003 Industrial Light and Magic a division of Lucas Digital Ltd. LLC.
Zlib © 1995-2003 Jean-loup Gaily and Mark Adler.
HDRI Import created 2002 by SplutterFish and Cuncyt Ozdas.
Portions Copyrighted © 2000-2005 Size8 Software, Inc.
Portions of this software are Copyright 1998-2004 Hybrid Graphics Limited.
This product includes Radiance software (http://radsite.lbl.gov/) developed by the Lawrence Berkeley National Laboratory (http://www.lbl.gov/).
The JPEG software is copyright © 1991-1998, Thomas G. Lane. All Rights Reserved.
Portions Copyrighted *mental images* GmbH 1989-2002.
Portions Copyright © IntegrityWare, Inc.; Npower Software LLC. All rights reserved.
Portions Copyright © 1991-1996 Arthur D. Applegate. All rights reserved.
Lenzfx and Max R2 Copyright © Digimation, Inc. All rights reserved.

GOVERNMENT USE

Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in FAR 12.212 (Commercial Computer Software-Restricted Rights) and DFAR 227.7202 (Rights in Technical Data and Computer Software), as applicable.
Published By: Autodesk, Inc.
111 McInnis Parkway
San Rafael, CA 94903, USA

toc

Contents

10 Precision and Drawing Aids 1
Precision and Drawing Aids 1
Tools for Precision 1
Helpers 2
Using Units 2
Using Grids 4
Using the Home Grid 4
Using Grid Objects 5
Viewing Grid Objects 6
AutoGrid 7
Aligning Objects 8
Aligning Normals 10
Setting Standard Snaps 11
Setting Snap Options 12
Measuring Distances 13
Measure Distance Tool 15
Select And Manipulate 15
Standard Helper Objects 16
Standard Helper Objects 16
Dummy Helper 16
Expose Transform Helper 17
Grid Helper 20
Point Helper 23
Tape Helper 24
Protractor Helper 26
Compass Helper 27
Manipulator Helper Objects 27
Manipulator Helper Objects 27
Cone Angle Manipulator 27
Plane Angle Manipulator 29
Slider Manipulator 31
Grid Commands 33
Grid Commands 33
Show Home Grid 34
Activate Home Grid 34
Activate Grid Object 34
Align Grid to View 35
Snap Commands 35
Snap Commands 35
2D Snap, 2.5D Snap, 3D Snap 35
Angle Snap Toggle 37
Percent Snap Toggle 38
Spinner Snap Toggle 38
Ortho Snapping Mode 38
Polar Snapping Mode 39
Grid and Snap Settings 41
Grid and Snap Settings 41
Snap Settings 41
Snap Override 45
Snap Options 46
Home Grid Settings 49
User Grids Settings 51
Drawing Aid Utilities 52
Drawing Aid Utilities 52
Measure Utility 52
Rescale World Units Utility 53

11 Space Warps and Particle Systems 55
Space Warps and Particle Systems 55
Space Warps 55
Space Warp Objects 55
Bind to Space Warp 58
Forces 59
Push Space Warp 59
Motor Space Warp 61
Vortex Space Warp 63
Drag Space Warp 66
PBomb Space Warp 68

Path Follow Space Warp 71
Gravity Space Warp 73
Wind Space Warp 75
Displace Space Warp 76
Deflectors 78
POmniFlect Space Warp 78
PDynaFlect Space Warp 81
SOmniFlect Space Warp 84
SDynaFlect Space Warp 85
UOmniFlect Space Warp 85
UDynaFlect Space Warp 86
SDeflector Space Warp 87
UDeflector Space Warp 89
Deflector Space Warp 90
Geometric/Deformable 91
FFD(Box) Space Warp 91
FFD(Cyl) Space Warp 95
Wave Space Warp 100
Ripple Space Warp 102
Conform Space Warp 103
Bomb Space Warp 105
Modifier-Based 107
Modifier-Based Space Warps 107
Particle Systems 108
Introduction to Particle Systems 108
Particle Flow 109
Particle Flow 109
Using Particle Flow 109
How Particle Flow Works 109
The Life of a Particle 110
Introductory Tutorial 112
Particle Flow FAQ 118
Particle Flow User Interface 125
Particle View 125
Particle View 125
Particle View Menu Bar 126
Particle View Event Display 131
Event Display Right-Click Menu 133
Particle Flow Source 135
Particle Flow Keyboard Shortcuts 140
Actions 141
Actions 141
Action Time Frames 141
Operators 142
Operators 142
Birth and Death 143
Birth Operator 143
Birth Script Operator 145
Delete Operator 146
Transforms 147
Position Icon Operator 147
Position Object Operator 148
Rotation Operator 153
Spin Operator 154
Scale Operator 156
Speed 159
Speed Operator 159
Speed By Icon Operator 162
Speed By Surface Operator 167
Keep Apart Operator 172
Shape 176
Shape Operator 176
Shape Facing Operator 176
Shape Instance Operator 178
Shape Mark Operator 183
Material and Mapping 186
Materials and Mapping in Particle View 186
Material Static Operator 187
Material Frequency Operator 189
Material Dynamic Operator 191
Mapping Operator 195
Additional Operators 197
Cache Operator 197
Display Operator 202
Force Operator 204
Notes Operator 206
Render Operator 206
Script Operator 208
Flows 208
Flows 208
Empty Flow 209
Standard Flow 209
Tests 210
Tests 210
Age Test 211
Collision Test 212
Collision Spawn Test 215
Find Target Test 218
Go To Rotation Test 224
Scale Test 227
Script Test 229
Send Out Test 230
Spawn Test 230
Speed Test 233
Split Amount Test 234
Split Selected Test 235
Split Source Test 236
Non-Event-Driven Particle Systems 237

Non-Event-Driven Particle Systems 237
Using Particle Systems 238
Particle System Usage 238
Creating a Particle Emitter 239
Using Materials with Particle Array 239
Achieving Particle Motion Blur 240
Using Mapped Materials with Particle Systems 240
Using Multi/Sub-Object Materials with Particle Systems 242
Using Spawned Particles 242
Using Interparticle Collision 243
Spray Particle System 244
Snow Particle System 246
Super Spray Particle System 249
Blizzard Particle System 251
PCloud Particle System 253
PArray Particle System 256
PArray Particle System 256
Basic Parameters Rollout (PArray) 258
Particle Generation Rollout 260
Particle Type Rollout 262
Rotation and Collision Rollout 268
Object Motion Inheritance Rollout 269
Bubble Motion Rollout 270
Particle Spawn Rollout 271
Load/Save Presets Rollout 274

12 Animation 275
Animation 275
Animation Concepts and Methods 275
Animation Concepts and Methods 275
Animation Concepts 276
Using Auto Key Mode 278
Using Set Key Mode 280
Spinner Right-Click Menu 282
Viewing and Copying Transform Keys 283
Controlling Time 285
Setting Time Segments 286
Moving Through Time 287
Choosing a Frame Rate and Playback Speed 288
Working with Controllers 289
Working with Controllers 289
Understanding Controllers 289
Changing Controller Properties 291
Assigning Controllers 292
Specifying Default Controllers 294
General-Purpose Controllers 295
Special-Purpose Controllers 296
Explicit Axis Keys 297
Float Controllers 297
Controlling Transforms 298
Controlling Position 298
Controlling Rotation 299
Controlling Colors 299
Morph Controllers 300
Motion Panel Commands 301
Trajectories 301
PRS Parameters Rollout 303
Key Info (Basic) Rollout/Dialog 304
Tangent Types 305
Key Info (Advanced) Rollout/Dialog 306
Animation Controllers 307
Animation Controllers 307
Audio Controller 309
Barycentric Morph Controller 309
Bezier Controllers 310
Block Controller 313
Boolean Controller 316
Color RGB Controller (Point3 XYZ Controller) 317
Euler XYZ Rotation Controller 318
Expression Controller 320
Expression Controller 320
Expression Controller Techniques 324
Layer Controller 325
Layer Controller 325
Animation Layers 326
Layer Properties Dialog 333
Create New Animation Layer Dialog 334
Limit Controller 335
Linear Controller 341
List Controller 342
Local Euler XYZ Rotation Controller 344
Look At Controller 344
Master Point Controller 346
Motion Capture Controller 347
Noise Controller 353
On/Off Controller 355
Position XYZ Controller 356
PRS Controller 357
Reaction Controllers 358
Reaction Controllers 358
Reaction Manager Dialog 361
Using Manipulators with Reaction Controllers 369
Scale XYZ Controller 371
Script Controller 372
Smooth Rotation Controller 374
Spring Controller 375
TCB Controllers 377

Transform Script Controller 379
Waveform Controller 381
XRef Controller 383
Controller Subdialogs 386
Audio Controller Dialog 386
Attach Controls Dialog (Block Controller) 388
Barycentric Morph Controller Key Info Dialog 388
Block Parameters Dialog (Block Controller) 389
Master Block Parameters Dialog (Block Controller) 390
Master Track Key Info Dialog (Master Point Controller) 391
Slave Parameters Dialog (Block Controller) 391
Track View Pick Dialog (Block Controller) 392
Animation Constraints 392
Animation Constraints 392
Attachment Constraint 393
Surface Constraint 396
Path Constraint 398
Position Constraint 401
Link Constraint 403
LookAt Constraint 406
Orientation Constraint 409
Wire Parameters 411
Wire Parameters 411
Parameter Wiring Dialog 412
Hierarchies and Kinematics 416
Hierarchies and Kinematics 416
Hierarchies 416
Hierarchies 416
Linking Strategy 418
Linking and Unlinking Objects 421
Linking and Unlinking Objects 421
Select and Link 422
Unlink Selection 422
Adjusting Pivots 423
Viewing and Selecting Hierarchies 424
Animating with Forward Kinematics 426
Animating with Forward Kinematics 426
Using Dummy Objects 429
Animating Links 430
Adjusting Object Transforms 432
Locking Object Transforms 433
Animating Attachment 433
Changing Link Inheritance 434
Link Inheritance (Selected) Utility 435
Animating with Inverse Kinematics (IK) 435
Inverse Kinematics (IK) 435
IK Terminology 437
IK Methods 439
Inverse Kinematics Methods 439
IK Solvers 440
IK Solvers 440
History-Independent (HI) IK Solver 446
History-Independent (HI) IK Solver 446
White Paper: Swivel Angle of the HI IK Solver 449
HI IK Solver Rollouts (Motion Panel) 453
IK Solver Rollout (HI Solver) 453
IK Solver Properties Rollout (HI Solver) 456
IK Display Options Rollout (HI Solver) 458
Sliding and Rotational Joints (HI Solver) 459
History-Dependent (HD) IK 461
History-Dependent (HD) IK Solver 461
Binding Objects (HD Solver) 461
Controlling IK Precision (HD Solver) 463
HD Solver Motion Panel Rollouts 463
HD Joint Settings 466
HD Joint Settings 466
Setting Joint Precedence (HD Solver) 467
Using Default Joint Precedence (HD Solver) 468
Choosing Child-to-Parent Precedence (HD Solver) 469
Choosing Parent-to-Child Precedence (HD Solver) 470
Setting Precedence Manually (HD Solver) 471
Terminating Chains (HD Solver) 471
IK Limb Solver 472
IK Limb Solver 472
Spline IK 473
Spline IK 473
Spline IK Solver Dialog 477
Spline IK Solver Rollouts 478
Interactive and Applied IK 480
Animating with Interactive IK 480
Animating with Applied IK 481
Joint Controls 483
Setting Joint Parameters 483
Activating Joint Axes 485
Limiting Joint Action 486
Hierarchy Panel Commands 487
Hierarchy Panel Commands 487
Pivot 487
Pivot 487
Adjust Pivot Rollout 488
Adjust Transform Rollout 489
IK 491
IK 491
HD IK Solver Rollouts 491

Object Parameters Rollout (HD Solver) 491
Position/Orientation/Bind to Follow Object (HD Solver) 492
Precedence (HD Solver) 494
Copying, Pasting, and Mirroring Joint Parameters (HD Solver) 495
Sliding and Rotational Joints Rollouts (HD Solver) 496
Interactive and Applied IK Rollouts 497
Inverse Kinematics Rollout 497
Auto Termination Rollout (Interactive IK) 499
Link Info 499
Link Info 499
Locks Rollout 500
Inherit Rollout 500
Track View 501
Track View 501
Working with Track View 503
Track View Workspace 504
Curve Editor 507
Dope Sheet 507
Time Ruler 510
Track View Shortcuts 510
Track View Hierarchy 512
Track View Hierarchy 512
Hierarchy Right-Click Menu 516
Properties (Track View Hierarchy) 519
Sound Options Dialog 520
Track View Menus 521
Track View Menu Bar 521
Modes Menu 521
Controller Menu 521
Controller Menu 521
Collapse Controller 522
Tracks Menu 524
Keys Menu 524
Keys Menu 524
Use Soft Select 525
Soft Selection Settings 525
Curves Menu 525
Time Menu 526
Options Menu 526
Options Menu 526
Interactive Update 526
Sync Cursor Time 527
Auto Expand 527
Auto Select 528
Auto Scroll 528
Modify Subtree 528
Modify Child Keys 529
Manual Navigation 529
Display Menu 530
Display Menu 530
Show All Tangents 531
Keyable Icons 531
Hide/Show Non-Selected Curves 532
View Menu 532
Utilities Menu 533
Track View Quad Menu 534
Track View Toolbars 535
Curve Editor Toolbars 535
Dope Sheet Toolbars 538
Controller Toolbar 540
Ranges: Dope Sheet Toolbar 541
Extras: Dope Sheet Toolbar 541
Filters 541
Filters Dialog (Track View) 542
Copy Controller 544
Paste Controller 545
Assign Controller 546
Delete Controller 549
Ignore Animation Range 549
Respect Animation Range 550
Make Controller Unique 550
Parameter Curve Out-of-Range Types 551
Add Note Track 552
Remove Note Track 553
Editing Keys 554
Editing Keys 554
Edit Keys 554
Snap Frames 554
Lock Selection 555
Align to Cursor 556
Add Visibility Track 556
Move Keys (Dope Sheet) 558
Slide Keys 558
Scale Keys - Time 559
Add Keys (Dope Sheet) 560
Properties (Track View Key Window) 560
Track View Utilities 561
Track View Utilities 561
Randomize Keys Utility 562
Create Out of Range Keys Utility 562
Select Keys By Time Utility 563
Euler Filter 564
Current Value Editor 565
Edit Time 566
Edit Time 566

Select Time 566
Delete Time 567
Cut Time 567
Copy Time 568
Paste Time 568
Reverse Time 569
Insert Time 570
Scale Time 570
Exclude Left End Point 570
Exclude Right End Point 571
Reduce Keys 572
Edit Ranges 573
Edit Ranges 573
Position Ranges 574
Recouple Ranges 574
Editing Tracks: Copying, Pasting, and Handling Instances and References 575
Copying and Pasting Items 575
Copying and Pasting Objects 576
Making Instance and Reference Controllers and Objects Unique 577
Curve Editor 578
Curve Editor 578
Move Keys (Curve Editor) 579
Scale Keys 580
Scale Values 581
Add Keys (Curve Editor) 581
Draw Curves 582
Show Tangents 582
Lock Tangents 583
Apply Ease Curve / Apply Multiplier Curve 584
Remove Ease/Multiplier Curve 585
On/Off (Curves) 585
Ease Curve Out-of-Range Types 585
Multiplier Curve Out-of-Range Types 586
Freeze Non-Selected Curves 587
Status Bar/View Controls 588
Status Bar/View Controls 588
Zoom Selected Object 588
Select By Name (Track View) 589
Track Sets List 590
Track Sets Editor Dialog 591
Key Time Display 594
Value Display 594
Show Selected Key Statistics 595
Pan (Track View) 595
Zoom Horizontal Extents 595
Zoom Value Extents 596
Zoom Track View Key Window 596
Zoom Region (Track View) 597
Managing Track View Windows 597
New Track View 597
Delete Track View 598
Delete Track View Dialog 598
Saved Track Views 599
Track View Customization 599
Track View Toolbar Right-Click Menu 603
Motion Mixer 604
Motion Mixer 604
Using the Motion Mixer 604
Using the Motion Mixer 604
Adding Tracks to the Mixer 607
Importing Clips to the Mixer 609
Working with Clips in the Mixer 611
Filtering Mixer Tracks 612
Adjusting Clip Timing 615
Working with Transitions 616
Adjusting Track Weight 619
Adding Time Warps 620
Adjusting Biped Balance in the Mixer 622
Exporting Animation to the Biped 624
Using the Reservoir 626
Motion Mixer User Interface 628
Motion Mixer Interface 628
Motion Mixer Menus 629
Motion Mixer Dialogs 634
Mixer Clip Source Options Dialog 634
Copy Biped Animation to Clip Dialog 635
Reservoir File Groups Dialog 635
Mixer Transition Editor Dialog (Biped Object) 636
Mixer Transition Editor Dialog (Non-Biped Object) 638
Pick Nodes Dialog (Motion Mixer) 641
Transition Optimization Dialog 641
Motion Mixer Toolbar 642
Trackgroup Filter Dialog (Biped Object) 645
Trackgroup Filter Dialog (Non-Biped Object) 645
Motion Mixer Editor 646
Reservoir 649
Mixer Preferences Dialog 651
Mixer Rollout 652
Animation Utilities 653
Animation Utilities 653
Follow/Bank Utility 653
Motion Capture Utility 655
MACUtilities Utility 665
Camera Tracker Utility 667
Camera Tracker Utility 667

Requirements for Camera Tracking 669
Camera Tracker: Movie Rollout 670
Camera Tracker: Movie Window 671
Camera Tracker: Motion Trackers Rollout 673
Camera Tracker: Movie Stepper Rollout 676
Camera Tracker: Error Thresholds Rollout 677
Camera Tracker: Batch Track Rollout 678
Camera Tracker: Position Data Rollout 679
Camera Tracker: Match Move Rollout 680
Camera Tracker: Move Smoothing Rollout 682
Camera Tracker: Object Pinning Rollout 683
Camera Tracker: Troubleshooting 685
Dynamics 686
Dynamics Utility 686
Edit Object Dialog 696
Edit Object List Dialog 700
Skin Utilities 700

13 reactor 703
reactor 703
Introducing Dynamics Simulation 707
Special Features in reactor 712
reactor Helpers 715
Rigid Bodies 716
Rigid Bodies 716
Rigid Body Basics 717
Rigid Body Properties 717
Compound Rigid Bodies 722
Rigid Body Collection 723
Constraints 724
Constraints 724
Constraint Concepts 725
Simple Constraints 727
Simple Constraints 727
Spring 727
Linear Dashpot 730
Angular Dashpot 732
Cooperative Constraints 735
Cooperative Constraints 735
Constraint Solver 736
Rag Doll Constraint 737
Defining Rag Doll Limits 743
Hinge Constraint 747
Point-Point Constraint 750
Prismatic Constraint 754
Car-Wheel Constraint 757
Point-Path Constraint 762
reactor Objects 764
Plane 764
Motor 765
Toy Car 766
Fracture 770
Fracture Tips 773
Storing and Accessing Collisions 774
Deformable Bodies 777
Deformable Bodies 777
Cloth 778
Cloth 778
Cloth Modifier 778
Cloth Collection 781
Soft Bodies 783
Soft Bodies 783
Soft Body Modifier 784
FFD Soft Bodies 786
Soft Body Collection 788
Rope 789
Rope 789
Rope Modifier 789
Rope Collection 792
Deforming Meshes 793
Deforming Meshes (Skin) 793
Deforming Mesh Collection 794
Deformable-Body Constraints 795
Constraining Deformable Bodies 795
Fixing Vertices in World Space 796
Keyframing Vertices 797
Attaching Vertices to a Rigid Body 798
Attaching Vertices to Deforming Meshes (Skin) 799
Soft Selection 800
Water 801
Water Simulation 801
Water Space Warp 801
Rendering Water 803
Wind 803
The reactor Utility 805
The reactor Utility 805
Preview & Animation Rollout 806
Havok 1 World/Havok 3 World Rollout 808
Collisions Rollout 810
Display Rollout 812
reactor Utilities 813
Properties Rollout 815
The Real-Time Preview 815
Scripts 817
Frequently Asked Questions 821
Troubleshooting 823

14 character studio 831
character studio 831
What is character studio? 831
What You Should Know to Use character studio 832
Understanding Biped 833
Understanding Physique 834
Understanding Track Editing 836
Understanding the Workbench 837
Understanding Motion Flow 837
Understanding Crowds 838
Understanding character studio Workflow 839
Biped 843
Biped 843
Creating a Biped 844
Understanding Biped Anatomy 846
Changing Initial Biped Anatomy 846
Naming the Biped 847
Posing the Biped 847
Scaling Links 851
Rubber-Banding Arms and Legs 852
Biped Display Options 853
Deleting a Biped 854
Linking Character Body Parts to the Biped 854
Saving and Loading FIG Files 855
Footstep Animation 856
Footstep Animation 856
Creating Footsteps 860
Creating Footsteps 860
Planning for Footsteps 860
Choosing a Gait 861
Creating Footsteps Automatically 862
Creating Footsteps Manually 863
Activating Footsteps 865
Deactivating Footsteps 865
Understanding Footstep and Body Keys 867
Editing Footstep Motion 868
Editing Footstep Motion 868
Editing Footstep Placement 869
Editing Footstep Timing 869
Editing Active Footsteps 871
Splicing Footsteps 872
Animating Legs and Feet 874
Animating the Upper Body 874
Adjusting Body Keys in Track View 875
Shifting the Biped's Balance 876
Adjusting Vertical Motion 878
Saving Footstep Animation 882
Footsteps and Freeform Animation 882
Footstep and Freeform Animation 882
Freeform Animation Between Footsteps 883
Converting Between Footstep and Freeform Animations 885
Freeform Animation 886
Freeform Animation 886
Creating Freeform Animation 888
Creating Freeform Animation 888
Selecting Biped Tracks 888
Animating by Moving Links 890
Animating by Rotating Links 891
Rotating Multiple Biped Links 895
Using Controllers 896
Using Props 898
Freeform and IK 900
Using IK Keyframe Parameters 900
Understanding Walk Cycle Constraints 902
Setting Keys for Feet and Hands 904
Animating Pivots 905
Animating a Quadruped 907
Animating IK Attachments 908
Editing Freeform Animation 909
Editing Freeform Animation 909
Editing Biped Keys 909
Copying and Pasting Postures and Poses 910
Mirroring Motion 913
Using Layers 913
Editing Trajectory Keys 914
Working with Biped Animation 916
Working with Euler Curves on Biped Animation 916
Working with XRef Bipeds 918
Loading, Saving, and Displaying Biped Motion 919
Working with Biped Motion Files 919
Loading and Saving BIP Animation 920
Importing and Exporting Animation Data 921
Motion Mapping: Retargeting Biped Motion 921
Merging and Cloning a Character 922
Combining BIP Motions 924
Loading and Saving STP Files 924
Using Motion-Capture Data 925
Correcting Posture 925
Copying and Pasting Tracks 926
Repositioning the Biped 927
Previewing Biped Motion 929
In Place Mode 930
Trajectory Display 931
Display Preferences 931
Biped User Interface 932
Biped User Interface 932

Motion Panel (Biped) 933
Center of Mass 933
Motion Panel Rollouts (Biped) 934
Assign Controller Rollout (character studio) 934
Biped Apps Rollout 935
Biped Rollout 936
Biped Rollout 936
Save As Dialog (Biped) 941
Open Dialog (Biped) 942
Display Preferences Dialog 944
Track Selection Rollout 945
Quaternion/Euler Rollout 948
Twist Poses Rollout 950
Bend Links Rollout 952
Key Info Rollout 954
Keyframing Tools Rollout 962
Set Multiple Keys Dialog 965
Copy/Paste Rollout 966
Layers Rollout 974
Dynamics & Adaptation Rollout 980
Figure Mode 982
Figure Mode 982
Structure Rollout 984
Footstep Mode Rollouts 988
Footstep Mode 988
Footstep Creation Rollout 988
Footstep Operations Rollout 990
Create Multiple Footsteps Dialog: Walk 992
Create Multiple Footsteps Dialog: Run 995
Create Multiple Footsteps Dialog: Jump 997
Convert to Freeform or Footsteps Dialogs 999
Footstep Mode Dialog 1000
Track View (Biped) 1002
Biped Color-coded Keys and Trajectories 1005
Biped Shortcuts 1006
Using the Workbench 1008
Working with the Workbench 1008
Navigating the Workbench 1010
Selecting Workbench Tracks 1011
Analyzing Curves 1011
Fixing Curves 1012
Workbench User Interface 1012
Animation Workbench 1012
Select Panel 1016
Analyze Panel 1017
Fix Panel 1020
Filters Panel 1023
Using Motion Flow 1026
Working with Motion Flow 1026
Placing Motions on the Motion Flow Graph 1027
Creating Transitions 1028
Creating a Motion Flow Script 1030
Saving, Loading, and Appending Motion Flow Graphs 1032
Customizing Transitions 1034
Creating Random Motion 1035
Unifying Motion 1038
Sharing Motion Flow 1039
Setting Up Paths for Motion Flow Files 1041
Motion Flow User Interface 1043
Motion Flow Mode 1043
Motion Flow Rollout 1045
Motion Flow Graph Dialog 1045
Motion Flow Scripts Group 1048
Transition Editor 1051
Create Random Motion Dialog 1055
Shared Motion Flow Dialog 1056
Transition Optimization Dialog 1058
Clip Properties Dialog 1059
Using Motion Capture 1059
Working with Motion-Capture Data 1059
Importing Motion-Capture Data 1061
Filtering Motion-Capture and Marker Data 1061
Sliding Footsteps 1064
Prop Bone 1065
Motion Capture User Interface 1065
Motion Capture Rollout 1065
Motion Capture Conversion Parameters Dialog 1070
Motion Capture Batch File Conversion Dialog 1075
Marker Display Dialog 1075
Physique 1076
Physique 1076
Using Physique 1076
Creating a Skin 1076
Using Physique with a Biped 1079
Using Physique with 3ds Max Bones 1080
Using Physique with 3ds Max Objects 1082
Applying and Initializing Physique 1083
Previewing Motion 1084
Envelopes and Vertex Assignments 1085
Envelopes 1085
Adjusting Default Envelope Shape 1086
Fine-Tuning Envelopes 1088
Customizing Vertex Assignments 1089
Envelope Display Options 1090
Working with Deformable Envelopes 1090
Working with Rigid Envelopes 1091

Working with Both Deformable and Rigid Envelopes....1091
Adjusting Link Parameters....1091
Partial Blending and Weight Assignments....1092
Bulges....1093
Bulges....1093
Creating Bulges....1094
Setting Bulge Angles....1095
Fine-Tuning Bulges....1096
Tendons....1096
Tendons....1096
Creating Tendons....1096
Saving and Loading Physique Data....1098
Reinitializing Physique Settings....1098
Working with an Initial Pose....1099
Improving Interactive Performance....1099
Scaling a Character....1100
Facial Animation....1100
Physique and Modifiers....1104
Combining Physique with Other Modifiers....1104
Physique and Free Form Deformations (FFDs)....1104
Physique User Interface....1106
Physique User Interface....1106
Physique Rollout....1106
Physique Level of Detail Rollout....1108
Floating Bones Rollout....1110
Physique Shortcuts....1111
Physique Dialogs....1111
Physique Initialization Dialog....1111
Reinitialize Physique....1113
Bulge Editor....1114
Physique Load Specification Dialog....1123
Physique Initialization Rollouts....1123
Vertex-Link Assignment Rollout....1123
Cross Sections Rollout....1125
Sub-Object Options Dialogs....1125
Blending Envelope Display Options Dialog....1125
Exclude Envelopes Dialog....1126
Bulge Angle Display Properties Dialog....1127
Tendon Display Options Dialog....1128
Type-In Weights Dialog....1128
Physique Modifier Sub-Object Levels....1129
Physique Sub-Objects....1129
Envelope Sub-Object....1130
Link Sub-Object....1135
Link Sub-Object....1135
Link Settings Rollout....1136
Joint Intersections Rollout....1140
Bulge Sub-Object....1141
Tendons Sub-Object....1147
Vertex Sub-Object (Physique)....1150
Crowd Animation....1154
Crowd Animation....1154
Creating a Crowd System....1155
Creating Crowd Helpers....1157
Adjusting Delegate Parameters....1159
Assigning Behaviors....1159
Directing Delegates....1162
Obstacle Avoidance....1164
Changing Delegate Orientation and Speed....1167
Solving the Simulation....1168
Linking Objects to Delegates....1169
Cognitive Controllers....1170
Using Motion Synthesis....1172
Using Motion Synthesis....1172
Biped Crowds....1172
Non-Biped Crowds....1179
Crowd Animation User Interface....1182
Crowd Animation User Interface....1182
Crowd Shortcuts....1182
Delegate Helper Object....1183
Geometry Parameters Rollout....1183
Motion Parameters Rollout....1183
Crowd Helper Object....1187
Crowd Helper Object....1187
Setup Rollout....1188
Setup Rollout....1188
Scatter Objects Dialog....1189
Random Placement Difficulty Dialog....1195
Object/Delegate Associations Dialog....1196
Edit Multiple Delegates Dialog....1197
Associate Bipeds With Delegates Dialog....1199
Behavior Assignments and Teams Dialog....1200
Select Behavior Type Dialog....1205
Select Delegates Dialog....1205
Cognitive Controllers....1206
Cognitive Controller Editor....1206
State Dialog....1207
State Transition Dialog....1208
Behavior Rollout....1211
Behavior Rollout....1211
Behaviors....1211
Avoid Behavior....1211
Orientation Behavior....1214
Path Follow Behavior....1216
Repel Behavior....1218
Scripted Behavior....1220
Seek Behavior....1220

Space Warp Behavior 1221
Speed Vary Behavior 1222
Surface Arrive Behavior 1223
Surface Follow Behavior 1226
Wall Repel Behavior 1227
Wall Seek Behavior 1229
Wander Behavior 1231
Solve Rollout 1232
Priority Rollout 1235
Set Start Frames Dialog 1237
Smoothing Rollout 1238
Collisions Rollout 1240
Geometry Rollout 1240
Global Clip Controllers Rollout 1241
Vector Field Space Warp 1241
Vector Field Space Warp 1241
Create Method Rollout 1242
Lattice Parameters Rollout 1242
Obstacle Parameters Rollout 1242
Motion Synthesis 1245
Motion Synthesis 1245
Synthesis Dialog 1246
Synthesis Dialog 1246
Motion Clips Panel 1246
State Panel 1248
Synthesis Panel 1250
MotionClip Parameters Dialog 1252
Track View Pick Dialog 1252
ClipState Dialog 1253
ClipState Dialog 1253
Speed Panel 1254
Acceleration Panel 1256
Pitch Panel 1257
Pitch Velocity Panel 1259
Heading Velocity Panel 1261
Script Panel 1262
File Formats and Index of Procedures 1263
character studio File Formats 1263
Procedures 1264

15 Lights and Cameras 1271
Lights and Cameras 1271
Lights 1272
Lights 1272
Name and Color Rollout (Lights) 1273
Using Lights 1274
Using Lights 1274
Working with Lights 1274
Properties of Light 1276
Lighting in 3ds Max 1279
Guidelines for Lighting 1280
Positioning Light Objects 1282
Animating Lights 1282
Light Include/Exclude Tool 1283
Light Lister 1285
Standard Lights 1288
Standard Lights 1288
Target Spotlight 1289
Free Spotlight 1290
Target Direct Light 1292
Free Direct Light 1293
Omni Light 1295
Skylight 1296
mr Area Omni Light 1298
mr Area Spotlight 1299
Photometric Lights 1301
Photometric Lights 1301
Photometric Lights: Preset Lights 1302
Target Point Light (Photometric) 1303
Free Point Light (Photometric) 1304
Target Linear Light (Photometric) 1305
Free Linear Light (Photometric) 1307
Target Area Light (Photometric) 1307
Free Area Light (Photometric) 1309
IES Sun Light (Photometric) 1309
IES Sky Light (Photometric) 1312
mental ray Sun & Sky 1313
mental ray Sun & Sky 1313
mr Sky 1318
mr Sun 1319
mr Physical Sky 1321
Isotropic Light Distribution (Photometric Lights) 1323
Diffuse Distribution (Photometric Lights) 1323
Spotlight Distribution (Photometric Lights) 1324
Web Distribution (Photometric Lights) 1325
Photometric Webs 1326
IES Standard File Format 1328
Example of Photometric Data File 1329
Common Lamp Values for Photometric Lights 1329
Common Lighting Rollouts 1331
General Lighting Parameters 1331
Exclude/Include Dialog 1335
Shadow Parameters 1337
Spotlight Parameters 1338
Advanced Effects Rollout 1341
mental ray Indirect Illumination Rollout (for Lights) 1343

mental ray Light Shader Rollout 1345
Additional Rollouts for Standard Lights 1345
Intensity/Color/Attenuation Parameters 1345
Directional Parameters 1348
Atmospheres and Effects for Lights 1349
Add Atmosphere or Effect Dialog....1351
Hair Light Attr(ibutes) Rollout 1351
Additional Rollouts for Photometric Lights 1352
Intensity/Color/Distribution Rollout 1352
Linear Light Parameters Rollout 1354
Area Light Parameters Rollout....1354
Area Light Sampling Rollout....1354
Web Parameters Rollout....1355
Rollouts for Specific Shadow Types 1356
Advanced Ray-Traced Parameters Rollout 1356
Area Shadows Rollout 1357
mental ray Shadow Map Rollout....1360
Optimizations Rollout....1361
Ray-Traced Shadow Parameters Rollout....1363
Shadow Map Parameters Rollout....1363
Cameras.... 1365
Cameras 1365
Free Camera 1370
Target Camera 1371
Using Cameras 1372
Characteristics of Cameras 1372
Common Camera Parameters 1373
Using Transforms to Aim a Camera....1379
Using Clipping Planes to Exclude Geometry....1379
Using the Horizon to Match Perspective 1380
Animating Cameras 1381
Multi-Pass Rendering Effects.... 1382
Multi-Pass Rendering Effects 1382
Depth of Field Parameter (mental ray Renderer) ..1383
Multi-Pass Depth of Field Parameters for Cameras 1383
Multi-Pass Motion Blur Parameters for Cameras ..1386
Camera Match Utility 1387
Camera Match Utility....1387
Camera Match Helper Objects.... 1391
Camera Match Helper 1391
CamPoint Helper....1391
Two-Point Perspective 1392
Camera Correction Modifier 1392

16 Material Editor, Materials, and Maps ... 1395
Material Editor, Materials, and Maps....1395
Designing Materials 1395
Designing Materials 1395
Sample Slots and Material Name 1396
Material Type....1397
Shading Type 1398
Material Components....1399
Lights and Shading....1399
Choosing Colors for Realism 1400
Using Maps to Enhance a Material 1403
Applying a Material to an Object 1405
Mapping Coordinates 1405
Saving a Material 1406
Material XML Exporter Utility 1407
UVW Remove Utility....1408
Material Editor.... 1409
Material Editor 1409
Material/Map Browser 1412
Copying and Pasting: Right-Click Menu for Materials, Maps, Bitmaps, and Colors 1418
Sample Slots 1420
Sample Slots....1420
Sample Slot Right-Click Menu....1422
Dragging and Dropping Maps and Materials....1423
Drag and Drop Sub-Object Material Assignment....1424
Creating a Custom Sample Object 1425
Material Editor Tools 1427
Material Editor Tools....1427
Material Editor Menu Bar 1428
Reflectance and Transmittance Display....1430
Sample Type 1432
Propagate Materials to Instances....1432
Backlight 1432
Sample Slot Background 1433
Sample UV Tiling 1433
Video Color Check 1434
Make Preview, Play Preview, Save Preview 1434
Material Editor Options 1436
Select By Material 1439
Get Material 1439
Put Material to Scene 1440
Assign Material to Selection 1441
Reset Map/Mtl to Default Settings 1442
Make Material Copy 1442
Make Unique (Material Editor) 1442
Put to Library 1443
Material ID Channel 1443
Material ID Channel Flyout 1444
Show Map in Viewport 1445
Show End Result 1446
Go to Parent 1446

Go Forward to Sibling 1447
Material/Map Navigator 1447
Pick Material From Object (Eyedropper) 1448
Name Field (Materials and Maps)....1448
Type Button (Materials and Maps) 1449
Animating Materials....1449
Animating Materials....1449
Creating and Playing Animated Material Previews....1450
Synchronizing an Animated Bitmap with the Scene....1450
Material Editor Subdialogs....1451
Copy (Instance) Map Dialog....1451
Copy or Swap Colors Dialog....1452
Create Material Preview Dialog....1452
Duplicate Name Dialog (Material Library)....1453
Merge Dialog (Material Library)....1453
Merge Material Library Dialog....1454
Put to Library Dialog....1455
Render Map Dialog....1455
Replace Map Dialog....1456
Replace Material Dialog....1456
Update Scene Materials Dialog....1457
Materials....1457
Types of Materials....1457
SuperSampling Rollout....1459
mental ray Connection Rollout....1461
DirectX Manager Rollout....1464
Standard Material....1465
Standard Material....1465
Shader Basic Parameters Rollout....1466
Understanding Shaders....1468
Basic Parameters Rollout (Standard Material)....1470
Extended Parameters Rollout (Standard Material)....1471
Maps Rollout (Standard Material)....1474
Dynamics Properties Rollout....1479
Basic Material Shaders....1480
Anisotropic Shader....1480
Blinn Shader....1480
Metal Shader....1481
Multi-Layer Shader....1481
Oren-Nayar-Blinn Shader....1482
Phong Shader....1482
Strauss Shader....1483
Translucent Shader....1484
Basic Parameters for Standard Materials....1485
Color Controls....1485
Self-Illumination Setting....1487
Opacity....1488
Diffuse Level....1489
Roughness....1490
Translucency Setting....1491
Specular Highlight Controls....1492
Anisotropic Highlights....1492
Blinn, Oren-Nayar-Blinn, and Phong Highlights....1493
Metal Highlights....1494
Multi-Layer Highlights....1495
Translucent Highlights....1496
Mapping Standard Material Components....1497
Ambient Color Mapping....1497
Diffuse Color Mapping....1498
Diffuse Level Mapping....1499
Diffuse Roughness Mapping....1500
Specular Color Mapping....1500
Specular Level Mapping....1501
Glossiness Mapping....1502
Self-Illumination Mapping....1502
Opacity Mapping....1503
Filter Color Mapping....1503
Anisotropy Mapping....1504
Orientation Mapping....1505
Metalness Mapping....1506
Bump Mapping....1506
Reflection Mapping....1508
Refraction Mapping....1509
Displacement Mapping....1511
Raytrace Material....1512
Raytrace Material....1512
Raytrace Basic Parameters Rollout....1514
Raytrace Extended Parameters Rollout....1519
Raytracer Controls Rollout....1521
Raytrace Maps Rollout....1523
Raytrace Dynamics Properties Rollout....1527
Raytracer Global Parameters Rollout....1528
Raytracing Acceleration Parameters Dialog....1531
Raytrace Exclude/Include Dialog....1531
Raytrace Antialiaser Dialog: Fast Adaptive Antialiaser....1533
Raytrace Antialiaser Dialog: Multiresolution Adaptive Antialiaser....1534
Architectural Material....1535
Architectural Material....1535
Templates Rollout....1536
Physical Qualities Rollout....1536
Special Effects Rollout....1539
Advanced Lighting Override Rollout....1540

Cutout Mapping 1542
mental ray Materials 1543
mental ray Materials....1543
mental ray Material 1544
mental ray Material 1544
Material Shaders Rollout (mental ray Material) 1544
Advanced Shaders Rollout (mental ray Material) 1548
Arch & Design Material (mental ray) 1549
Arch & Design Material (mental ray) 1549
Arch & Design Material (mental ray) – Overview....1562
Arch & Design Material (mental ray) – Tips & Tricks 1569
Car Paint Material/Shader (mental ray)....1576
DGS Material (mental ray)....1580
Glass Material (mental ray)....1582
Subsurface Scattering (SSS) Materials 1583
Matte/Shadow Material....1584
Compound Materials.... 1587
Compound Materials 1587
Blend Material 1588
Composite Material....1589
Double-Sided Material....1591
Morpher Material....1592
Multi/Sub-Object Material....1594
Shellac Material 1597
Top/Bottom Material....1599
Shell Material....1600
Advanced Lighting Override Material....1601
Lightscape Material....1604
Ink 'n Paint Material....1605
DirectX 9 Shader Material....1613
DirectX Viewport Shaders 1614
LightMap Shader Rollout 1614
Metal Bump Shader Rollout 1614
XRef Material....1616
Maps 1617
Types of Maps....1617
Real-World Mapping....1619
Output Rollout 1621
Missing Map Coordinates Dialog 1623
2D.... 1624
2D Maps 1624
Coordinates Rollout (2D) 1625
Noise Rollout (2D) 1630
Bitmap 2D Map 1631
Select Bitmap Image File Dialog 1635
Checker Map 1638
Combustion Map 1639
Gradient Map 1650
Gradient Ramp Map....1652
Flag Properties Dialog....1655
Swirl Map....1656
Tiles Map 1658
3D.... 1662
3D Maps 1662
Coordinates Rollout (3D) 1663
Cellular Map....1664
Dent Map....1667
Falloff Map 1670
Marble Map 1673
Noise Map....1674
Particle Age Map 1675
Particle MBlur Map 1676
Perlin Marble Map....1677
Planet Map....1678
Smoke Map....1679
Speckle Map 1680
Splat Map....1681
Stucco Map....1682
Waves Map....1683
Wood Map....1684
Compositor 1687
Compositor Maps....1687
Composite Map 1688
Mask Map....1689
Mix Map 1689
RGB Multiply Map 1691
Color Modifier.... 1692
Color Modifier Maps....1692
Output Map....1692
RGB Tint Map 1693
Vertex Color Map 1693
Reflection and Refraction 1695
Reflection and Refraction Maps 1695
Flat Mirror Map....1695
Raytrace Map 1698
Reflect/Refract Map 1699
Thin Wall Refraction Map....1703
Raytrace Map Rollouts 1704
Raytracer Parameters Rollout 1704
Raytrace: Attenuation Rollout 1706
Raytrace: Basic Material Extensions Rollout 1707
Raytrace: Refractive Material Extensions Rollout 1708
mental ray Shaders.... 1710
mental ray Shaders 1710

Custom Shaders for 3ds Max 1711
mental images Shader Libraries 1712
Shaders in the LumeTools Collection 1713
Connect Parameter to Shader Dialog (mental ray) 1713
3ds Max Custom Shaders 1714
3D Displacement Shader (mental ray) 1714
Bump Shader (mental ray) 1716
DGS Material Shader (mental ray) 1717
Dielectric Material Shader (mental ray) 1719
Environment Shader (mental ray) 1721
Height Map Displacement Shader (mental ray) 1722
Material to Shader (mental ray) 1723
Shader List (mental ray) 1723
UV Generator 1724
UV Generator Shader (mental ray) 1724
UV Generator Parameters Rollout 1725
Shaders Rollout (UV Generator) 1727
UV Coordinate Shader (mental ray) 1728
XYZ Generator 1729
XYZ Generator Shader (mental ray) 1729
XYZ Generator Parameters Rollout 1729
Shaders Rollout (XYZ Generator) 1730
XYZ Coordinate Shader (mental ray) 1730
Normal Bump Map 1731
Camera Map Per Pixel Map 1732
Material, Mapping, and Vertex Color Utilities .. 1734
Assign Vertex Colors Utility 1734
Channel Info Utility 1738
Clean MultiMaterial Utility 1742
Instance Duplicate Maps Utility 1744

Index 1749

10 Precision and Drawing Aids

3ds Max provides tools that give you control over the positioning and alignment of objects in 3D space. With these tools, you can do the following:

- Choose display units from the most common real-world measuring systems or define your own.
- Use the home grid as a construction plane, or use grid objects to position custom construction planes.
- Select different options to align objects with grids, points, and normals.
- Use 3D object snaps on a modeless dialog as you build and move geometry in your scene. Grid points and lines are among the many snap options.
- Use "helper objects" in your work. This category includes grid objects as well as objects used for positioning and measurement.

This section presents these brief topics designed to help you quickly start learning how to use the tools that make precision possible:

Tools for Precision (page 2–1)

Using Units (page 2–2)

Using Grids (page 2–4)

Using the Home Grid (page 2–4)

Using Grid Objects (page 2–5)

Aligning Objects (page 2–8)

Aligning Normals (page 2–10)

Setting Standard Snaps (page 2–11)

Setting Snap Options (page 2–12)

Measure Distance Tool (page 2–15)

Tools for Precision

A set of interrelated tools in the program gives you precise control of the scale, placement, and movement of objects in your scene. These are especially important tools for those who build precise models in real-world units of measurement.

Basic Tools

The tools for precision are grouped as follows:

Units—Define different measurement systems. Besides the generic unit, you have your choice of feet and inches in both decimals and fractions. Metric units range from millimeters to kilometers. You can also define other units.

Grids—Include the home grid and special grid objects. Both types of grid can act as construction

planes. The software constructs objects using the orientation and position of the active grid. While the home grid is fixed in world space, you can rotate grid objects and place them anywhere in a scene, and align them to other objects and surfaces. You can also give each grid object its own spacing, and display any grid as a dedicated viewport.

Object alignment—Matches an object with the position, orientation, or normal of another object, or to a point in space.

Object snaps—Ensure precise placement when creating and rearranging objects. Keyboard shortcuts let you change object snaps as you work. You can also set snaps to find grid lines and intersections. An angle snap sets the increment for rotation, and a percent snap sets the increment for scaling.

Helpers—Provide useful assistance, as the name implies. These are specialized tools in the same category as grid objects. A Tape object measures distances in current units, and a Protractor object measure angles. A Grid object defines a construction plane, and a Point object marks a particular spot in 3D space.

How the Tools Work Together

The tools themselves establish a general order of use and interaction, although you can always change settings as needed without following this sequence.

- Choose a measuring unit. The default is generic units, sufficient for many purposes.
- Set grid spacing (the size of the smallest square), based on the measuring unit. The home grid and grid objects can have their own spacing, separate from the grid spacing.
- Move and align grid objects to a useful orientation.
- Set or vary snap settings as needed in your work.
- Use other helper objects like Point and Tape as part of the precision process.

As you work, you can change your settings (including the measuring unit) without losing any precision.

Helpers 

Create panel > Helpers

Create menu > Helpers

Helper objects play a supporting role, like stage hands or construction assistants.

Several categories of helpers are available from the drop-down list on the Create panel:

Standard helpers (page 2–16)

Atmospheric Apparatus (page 3–304)

Camera Match Helper (page 2–1391)

Assembly Head Helper Object (page 1–111)

- *Character Assembly (page 1–102)*
- *Luminaire Helper Object (page 1–111)*

Manipulator Helper Objects (page 2–27)

Particle Flow (*Speed By Icon Operator (page 2–162)* and *Find Target Test (page 2–218)*)

VRML97 Helper Objects (page 3–597)

reactor (page 2–703)

Other helper objects might be available, depending on your configuration.

Using Units

Units are the key to connecting the three-dimensional world of 3ds Max with

the physical world. You define the units you want to use from the *Units Setup dialog (page 3–848)*.

Changing Display Units

When you change display units, 3ds Max displays measurements in the new unit for your convenience. All dimensions are displayed in the new unit. Essentially, you're using a new "measuring stick." No object is changed in this process. As in the physical world, objects in the scene maintain their absolute size, regardless of how you measure them.

Type-In Entry

When you enter any dimension, 3ds Max always assumes the number you enter is expressed in the current units. You can also enter a series of numbers: 3ds Max then converts their sum into the current unit. Here are some examples that assume the current units are in centimeters:

- When you enter a dimension of 1' (one US foot), it converts to 30.48cm.
- If you enter a series of numbers such as 14 286 175 (separated by spaces), the series is totaled to 475.0cm.
- If you enter 1' 1 (one US foot and 1 centimeter), this is converted and summed into 31.48cm.

When you use US Standard as the display unit scale, you can select either feet or inches as the default for type-in entry. If you select feet and enter 12, the result is 12' 0". However, if you enter 1' 2, the software identifies the second digit as inches, producing 1' 2" as the result.

In any unit system, you can enter fractional amounts. For example, assume you're working in US Standard with feet as the default:

- If you enter 18/3, the result is 6'0".
- If you enter 18/3", the result is 0'6".
- You can specify units in a different system, and they are converted on the fly. For example, if you enter 18/3cm, the result is 0'2.362".

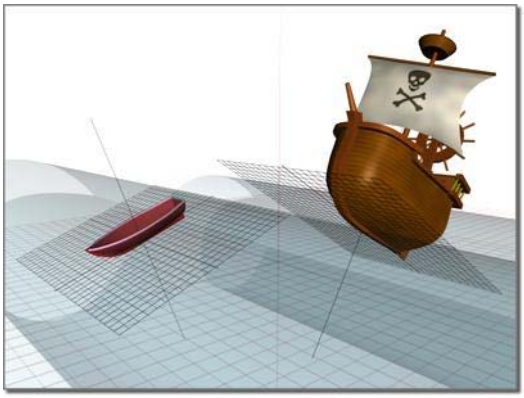
Understanding the System Unit

3ds Max keeps track of all measurements in its own internal system unit. No matter what kind of display units you use, measurements are stored in this absolute unit for storage and computation. The default system unit is defined as 1.000 inch. As long as the system unit is left at one inch, you can freely share models and change units on the fly with no effect on the underlying geometry. Except in rare circumstances, you never need to change this default scale. This means you can merge a model created with any standard unit into your scene at true scale.

You can change the system unit setting on the System Unit Setup dialog, available from the *Units Setup dialog (page 3–848)*. Changing the system unit is recommended only if your scene has very small (less than one inch) or very large dimensions. See *System Unit Setup Dialog (page 3–850)* for more information.

If you do need to change the system unit, change it *before* you create or import geometry. Do not change the system unit in an existing scene.

Using Grids



One grid establishes the pitch of the boat, another the pitch of the ship

Grids are two-dimensional arrays of lines similar to graph paper, except that you can adjust the spacing and other features of the grid to the needs of your work.

Grids have these primary uses:

- As an aid in visualizing space, scale, and distance.
- As *construction planes* where you create and align objects in your scene.
- As a reference system for using the Snap feature to align objects.

Home Grid and Grid Objects

3ds Max provides two kinds of grids: the home grid and grid objects. In addition, it includes the AutoGrid feature, an automated way of creating grid objects.

Home grid—The home grid is defined by three planes along the world X, Y, and Z axes. Each of these axes passes through the world coordinate system's origin point (0,0,0). The home grid is fixed; it cannot be moved or rotated.

- The home grid is visible by default when you start 3ds Max, but its display can be turned off.
- You can use any view of the home grid as a construction plane by drawing in the viewport in which the grid's view appears.

See *Viewing and Navigating 3D Space (page 1–21)* for a complete introduction to the home grid.

Grid objects—A grid object is a helper object you can create whenever you need a local reference grid or construction plane somewhere other than the home grid.

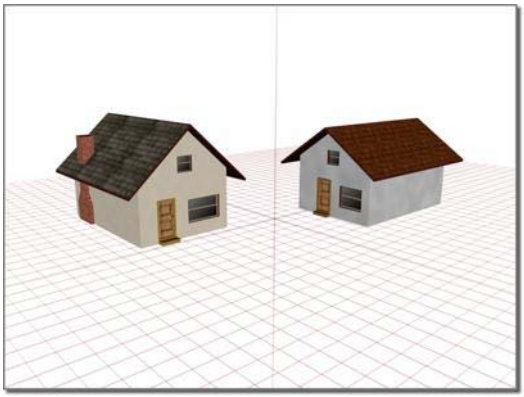
- You can have any number of grid objects in your scene, but only one can be active at a time.
- When active, a grid object replaces the home grid in all viewports.
- Each grid object has its own set of XY, YZ, and ZX planes. You can freely move and rotate grid objects, placing them at any angle in space, or attach them to objects and surfaces.
- You can change a viewport to display a plan (top) view of any active grid object.
- Grid objects can be named and saved like other objects, or used once and deleted.

AutoGrid—This feature lets you create new objects and grid objects off the surfaces of other objects on the fly. See *AutoGrid (page 2–7)*.

Using the Home Grid

The home grid provides ready-to-use construction planes, much like a leveled building site marked with stakes and strings. When you create an object in a viewport, the new object is placed on the home grid plane of that viewport.

To use the home grid effectively for construction, you often need to change the defaults to the job at hand, analogous to moving the stakes and strings to match your own site plan.



Using the home grid to position houses

Changing Home Grid Settings

The home grid is a single system; its three planes use the same settings for grid spacing and major line divisions. You change these settings from a single panel of the Grid and Snap Settings dialog.

To access the Home Grid panel, do one of the following:

- Choose Customize menu > Grid And Snap Settings, then click the Home Grid tab.
- Right-click any *snap button (page 2–35)*, then click the Home Grid tab.

Setting Grid Spacing

Grid spacing is the size, in current units, of the grid's smallest square. The basic idea is to choose a grid spacing that corresponds to your unit of measurement, then choose a larger spacing for multiple units.

For example, if you have units set to centimeters, you might make one grid space equal to 1.000 (one unit, or one centimeter in this case).

Setting Major Grid Divisions

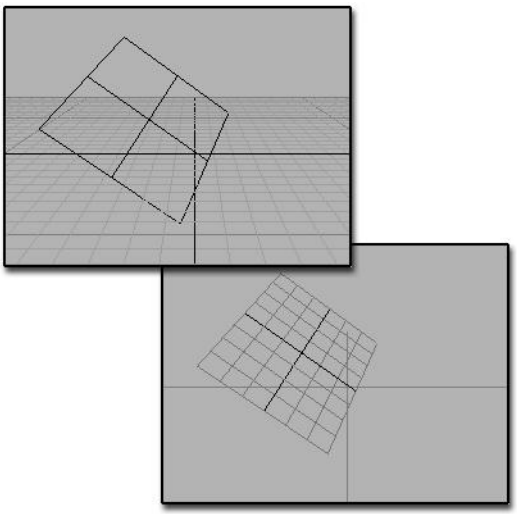
The home grid displays heavier or "major" lines to mark groups of grid squares. You can use these to represent larger units of measurement.

For example, if you use a grid spacing of one centimeter, you might use a value of 10 so the major grid divisions represent one decimeter.

Setting Color and Intensity

To improve grid visibility, you can change the intensity or color of the home grid. Choose Customize menu > Customize User Interface > *Colors panel (page 3–799)*.

Using Grid Objects



Above: Inactive grid object in a scene

Below: Activated grid object

Grid objects let you bypass the home grid and work on separately defined grids to create and position objects. You can use as many grid objects as you like, each serving as a custom construction plane with its own grid settings.

Grid objects are 2D parametric objects, with adjustments for overall size and grid spacing. You can adjust their orientation in world space, and match them to a particular surface or object.

To activate a grid object:

1. Create a grid object.
2. Choose Views > Grids > Activate Grid Object.

To deactivate a grid object, do one of the following:

- Select and activate another grid object.
- Choose Views > Grids > Activate Home Grid.
- Delete the grid object.

Creating and Modifying Grid Objects

You find Grid objects in the Create panel > Helpers category.

When creating an object, you can also create a temporary "AutoGrid" grid object that's tangent to an existing object surface at the point you click. To do this, turn on the *AutoGrid (page 2–7)* check box in the Object Type rollout. You can save this grid by holding the Alt key before you click.

Grid objects are named when you create them and are saved with the scene. You can delete them at any time.

Like other objects created in 3ds Max, standard grid objects are placed on the grid of the current viewport. By default, this is a plane of the home grid, but can also be another activated grid object.

For more information on grid object settings, see *Grid (page 2–20)*.

Viewing Grid Objects

An activated *grid helper object (page 2–20)* creates a true plane in 3D space. No matter how small an activated grid object appears on the screen, its plane is effectively infinite, just as if it were a plane of the home grid. However, you can view a given grid object in different ways.

Setting the Display Plane

You can adjust the visible plane of a selected or activated grid object on the Modify panel.

To set the display plane:

1. Open the Modify panel to display the Parameters rollout for the selected grid.
2. In the Display group, choose any of the three planes: XY, YZ, or ZX.

   The grid rotates in world space to display the corresponding plane, always based on the grid's local axis.

You can move or rotate a grid, either before or after setting its display plane.

Setting Grid Views

You can set the view in any viewport to an orthographic view as well as the display view.

To set grid view:

1. Activate the grid object.
2. Right-click a viewport label and choose Views > Grid to display a menu of possible views.
3. Choose Front, Back, Top, Bottom, Right, Left, or Display Planes. Display Planes corresponds to the current setting on the Parameters rollout (XY, YZ, or ZX).

   The viewport now shows this view of the grid.

You can set different viewports to different views of the grid. Objects are always created on a view's orthographic grid, even if you rotate the angle of view. In other words, you can set up a convenient view and still construct on the chosen plane.

Deactivating Grid Objects

You can deactivate a grid object in any of the following ways:

- Select and activate another grid object.

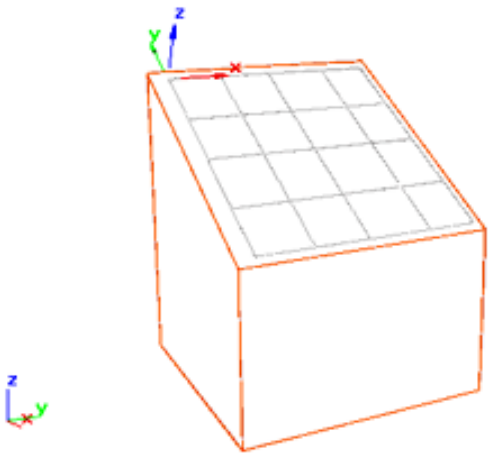
- Choose Views > Grids > Activate Home Grid.
- Delete the grid object.

When you deactivate a grid object, any viewports based on that grid switch to the corresponding orthographic view. For example, a Grid (Front) viewport becomes a Front viewport. A Grid (Display Planes) viewport always switches to Top view, regardless of the currently displayed plane.

AutoGrid

Create panel > Any object category > Object Type rollout

Extras toolbar > AutoGrid





AutoGrid used to position the second block on top of the first

AutoGrid lets you automatically create, merge, or import objects on the surface of other objects by generating and activating a temporary construction plane based on normals of the face that you click. This serves as a more efficient way of stacking objects as you create them, rather than building objects and then aligning them as a separate step.

You can use objects within an XRef scene as reference when AutoGrid is on.

Note: If the Smooth check box is turned on in the Parameters rollout of a parametric object, the construction plane is placed tangent to the face of the surface implied by any smoothing present on the surface, not the actual face of the surface.

Procedure

Example: To create a box aligned to another object using a temporary construction grid:

1. Create or load a scene containing an object to which you want to align a new box.

2.  On the Create panel > Standard Primitives > Object Type rollout, click Box.
3. Turn on AutoGrid.
4. Move your cursor over the object to which you want to orient the box you are about to create.

   The cursor includes an X,Y,Z tripod to help you orient the position of the new object. As you move over the object, the cursor aligns the Z axis to the surface normal.
5. When the orientation is as you want it, click and use the standard drag-release-move-click method to create the box. The box will be created on the designated surface.

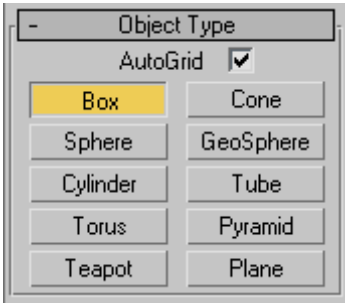
   When you click, a temporary, automatic grid is created, and the newly created object is aligned to that grid.

After creating the object, 3ds Max places it on the temporary construction grid. When creating an object, if the cursor isn't over another object, 3ds Max places the object on the current active grid.

Tip: When you create an object that requires multiple clicks, AutoGrid applies only to the first click. So, for instance, if you want to create a *Line spline (page 1–270)* that snaps to the faces of a sphere, turn on *Snaps (page 2–35)* (press the S key), and then open the *Grid And Snap Settings dialog (page 2–41)* and turn on Face.

Note: To make the grid permanent, hold down the Alt key before you click. The grid becomes active and 3ds Max turns AutoGrid off.

Interface



Object Type rollout

AutoGrid—AutoGrid is available only after you select an object button (such as Box). When you turn AutoGrid on, the cursor includes an axis tripod to help you orient the grid. Before clicking and as you position the cursor over a visible mesh object, the cursor jumps to the nearest point on that surface. The tripod's X and Y axes form a plane tangent to the object surface (forming an implied construction grid), and the Z axis is perpendicular to the plane.

Aligning Objects

You can align a selection consisting of one or more objects, called the *source*, with a *target* object. There are many uses for this feature. For precision, an important use is grid alignment. You can create a new grid object and align it to an existing object, either manually after creation or automatically during creation using the *AutoGrid (page 2–7)* feature. Conversely, you can move an object onto a grid anywhere in your scene.



Aligning objects by the center, bottom, or top of their bounding boxes

Source and Target Objects

Alignment involves two entities: one is the source object or selection set, where the process begins; the other is the target object, where the selection process ends.

Source object—Object or objects you want to move into alignment with another object. You select one or more source objects to begin the alignment process.

Target object—Object used to define the alignment. You select the target object during the alignment process. It cannot be selected beforehand.

After selecting the source, choose Tools menu > Align or click the Align button on the Main toolbar, and then select the target object.

Next, the Align dialog appears. For discussion of the Align options, see *Align (page 1–462)*.

Setting a Coordinate System

The effects of alignment depend on the current reference coordinate system, such as View, World, or Local. You should decide which system you want to use before beginning alignment.

Reference Coordinate System—Determines the axes used for positional alignment and the size of the bounding box for maximum and minimum positions.

To align objects using active grid axes:

- Choose Grid as the reference coordinate system from the list on the toolbar.

To align two objects using their own axes:

- Choose Local as the reference coordinate system. Alignment is then strictly between the two objects. Object bounding boxes determine maximum and minimum positions.

As a reminder to you during the alignment process, the current reference coordinate system appears in parentheses following the Align Position label in the Align Selection dialog.

Basics of Aligning Objects

Alignment controls are on a single dialog. As you make a setting, the object being aligned moves immediately to the new position. This lets you experiment with alignment until you get what you want. Settings can be made in any order.

You can also work step-by-step, by applying position choices, for example, before deciding on final orientation. You can cancel at any time, returning the scene to its original state. You can also undo any alignment and start over.

Aligning Multiple Objects



Aligning multiple objects to the arrow object (each photo is adjusted in a different correction)

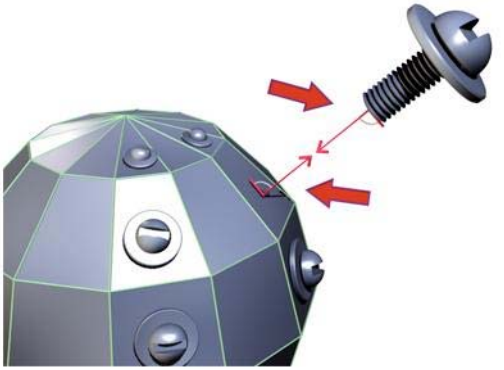
When you select multiple objects for alignment, the same settings apply to all of them. However, the effect on each source object can be different. In practice, you're aligning separate objects at the same time with the same settings.

To align a collection of objects as a single unit, select the objects and *group (page 1–96)* them. The alignment now takes place relative to the pivot and bounding box of the entire group.

Sub-Object Alignment

You can use Align with any selection that can be transformed. The tripod axis becomes the source for alignment. You do this by accessing the sub-object level of the object before clicking Align.

Aligning Normals



Normal on the front of the bolt is aligned with the normal of a face on the assembly.

3ds Max lets you align the normals between any two objects. In the case of mesh objects, the alignment is between individual faces, because each face has its own normal. See *Normal Align (page 1–465)*.

Basics of Aligning Normals

Before you begin, select a view that lets you see both objects you want to align. If necessary, you can navigate the view after selecting the first normal. To see face normals clearly, work in a wireframe viewport.

For details on the procedure and dialog settings, see *Normal Align (page 1–465)*.

Other Align Options

The Align flyout on the main toolbar has other buttons that provide specialized alignments.

 **Align Camera—**Orients a camera viewport to a selected face normal, with the normal in the center of the viewport and on the axis of the camera. See *Align Camera (page 1–468)*.

Align to View—Orients a local axis of an object or sub-object selection with the current viewport. Options on the dialog are interactive, as they are on the Normal Align dialog. See *Align to View (page 1–468)*.

Place Highlight—Orients a face normal to a light. See *Place Highlight (page 1–467)* and *Lighting in 3ds Max (page 2–1279)*.

Setting Standard Snaps



Snapping to: endpoints, midpoints, edges

Standard snaps give you control in creating, moving, rotating, and scaling objects. You access the snap features in the program from buttons on the main toolbar.

You make most snap settings on Grid And Snap Settings, a modeless dialog with four panels. You can move this dialog to any convenient location on your screen and turn options on and off as you work. Panels are reset to defaults for each new session.

You can also make commonly used snap settings on the *Snaps toolbar (page 2–42)*.

Snap settings are stored in the *3dsmax.ini (page 1–18)* file rather than in the MAX file. This means that the state of the snap settings persists from session to session without your having to modify the *maxstart.max* file.

Grid And Snap Settings Dialog and Snaps Toolbar

The most commonly used grid and object snaps appear on the *Grid and Snap Settings dialog (page 2–41)*. This is the general sequence for using these snaps:

1. Turn on snapping by clicking the Snaps Toggle button on the main toolbar, and then right-click this button to display the Grid And Snap Settings dialog.

   Two sets of snap types are available: Standard (the default) and NURBS. Both are discussed in *Snap Settings (page 2–41)*.

   You can also make commonly used snap settings on the *Snaps toolbar (page 2–42)*.

2. By default, only the Grid Points snap type is active. Turn on other snap types to activate them. When you create or move objects, these snaps are now in effect anywhere in 3D space, and unaffected by the current transform coordinate system.

   As you move the cursor, each snap type is marked by a distinctive icon shown on the Snap panel. The current icon indicates the type and position of the next snap.

   Grid points and grid lines are secondary to the other snap types. For example, if both Grid Point and Vertex are active, the software snaps to a vertex in preference to an equally close grid point.

3. Turn snaps on and off as needed; the easiest way is by pressing the S key.

For more information on snaps, see *2D Snap, 2.5D Snap, 3D Snap (page 2–35)* and *Snap Settings (page 2–41)*.

Snap Override

Snap Override bypasses currently selected snaps. You use a keyboard-mouse combination or a keyboard shortcut to define a new snap for the next click. You can override on the fly, one snap at a time.

For example, while creating a spline between grid points, you might need to snap to a vertex or midpoint of an object. This is the general procedure:

1. As you use preset snaps to create or move an object, hold down **Shift** and right-click. This pops up a three-quadrant right-click menu, which you can use to override the current snap setting(s).

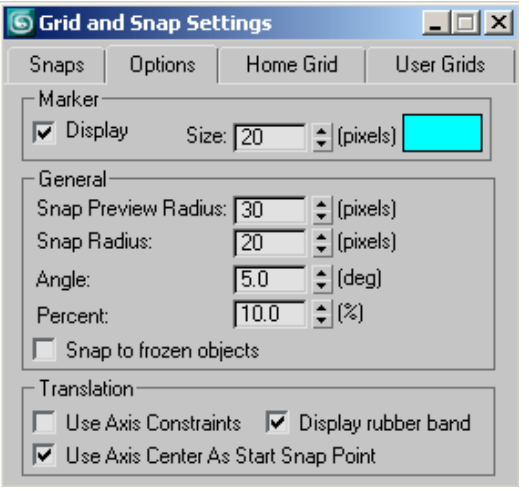
   Note: You can begin the creation or transformation process before accessing the override menu. This typically means that you'll be holding down the left button as you **Shift**+right-click.

2. From the Snap Override quadrant, choose the snap type you want to use. The cursor switches to this type.
3. Make the snap. After the snap, currently set snaps are back in effect.

In addition, the Snap Override quadrant lets you reuse the last override you used (listed by name), and gives you an option for None. None turns off snapping entirely on the next click. The Snap quad menu also contains an Options quadrant, with toggles for snapping within current transform constraints (default=off) and snapping to frozen objects (default=off).

Setting Snap Options

You can access a number of snap features from the Options tab of the Grid and Snap Settings dialog. Right-click any of the snap buttons on the main toolbar to display the Grid And Snap Settings dialog or choose Customize menu > Grid And Snap Settings, then click the Options tab.



See also

Snap Options (page 2–46)

Display and General Settings

Marker settings—Determine the color and size of the snap cursor. To prevent the snap cursor from appearing, turn off Display.

Snap Radius settings—Determines how close the cursor needs to approach a snap point before the snap preview or actual snap takes place. These are global settings, affecting all snap interactions, and are measured in terms of the pixels in a "search region" around the active point of the cursor.

Snap to Frozen Objects—Normally, if an object is frozen you can't snap to it; this option lets you snap to frozen objects.

Settings for Angle and Percent Snap

The following Options settings are for two snap buttons that operate independently of standard snaps.

Angle (deg)—A global setting, in degrees, that determines the angle of rotation for a number of features in the program, including the standard Rotate transform. As you rotate an object (or group of objects), the object moves around a given axis in the increment you set. Angle snap also affects the following:

- Pan/Orbit camera controls
- FOV and Roll camera settings
- Hotspot and Falloff light angles

For more information, see *Angle Snap (page 2–37)*.

Percent—Sets a percentage increment during a scaling operation.

For more information, see *Percent Snap (page 2–38)*.

Rotating and Scaling with Snaps

The effect of rotating and scaling with snaps depends on whether *Auto Key (page 3–717)* is on or off:

- With Auto Key *on*, snap toggles are disabled, while Angle and Percent snaps remain active. Rotation and scaling occur around the pivot point of the object.
- With Auto Key *off*, rotations and scales occur around the snap point. For example, using Vertex snap, you can rotate a box about any of its corners.

Translation Options

By default, the *Use Axis Constraints (page 1–437)* option is off. The current setting on the Axis Constraints toolbar (XY, for example) has no effect. Turning on this option lets you use snaps in conjunction with axis constraints.

The Translation group also lets you toggle display of a rubber band line between the start and end points during a snap operation.

Settings for Spinner Snap

You set the spinner snap on the *General panel (page 3–815)* of the Preferences dialog. Right-click the Spinner Snap button on the main toolbar to display this panel.


Spinner snap—Sets a numerical increment for spinner fields. If you're using generic units of 1 inch, a setting of 12 would let you resize objects by one foot with every click, or add 12 segments to a sphere.

The same setting applies to all spinner fields. Since spinner snap is a toggle, you can easily turn it on when needed and use the default at other times. Spinner snap has no effect on dragging a spinner, only on single clicks.

For more information, see *Spinner Snap (page 2–38)*.

Measuring Distances

3ds Max provides several options for measuring various aspects of your scene. The tape, protractor and compass helpers measure distance, angles and direction respectively. The Measure utility has a floater to display various measurements of any selected object.

You can quickly measure the distance between two points with the *Measure Distance tool (page 2–15)*.

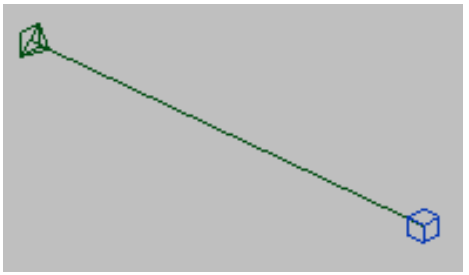
Helper Objects for Measurement

Tape Helper

You can create a *Tape helper (page 2–24)* by dragging between any two points in 3D space, much like using a physical tape measure. You then read the length on the Parameters rollout.

If you turn on Specify Length, the length field lets you enter a value in current units. This is like

locking a tape measure to a set length. You can position the tape object and snap to its ends.



The pyramid icon is the tape helper object; the cube is the helper object's target.

- To display length and angle settings, select only the tape helper object (the pyramid icon).
- To move the entire tape including its target (the cube icon), select the connecting line.

Two sets of World Space Angles give you current readouts to the three world axes (X, Y, Z) and the three world planes (XY, YZ, ZX).

Tip: When measuring with precision, it is helpful to use snaps to force the ends of the tape object to exact locations. If there are no specific points to snap to, you can use dummy or point helpers to set points.

To measure with a tape helper using snaps:

1. Locate points in the scene to which you can snap to accurately make a measurement. If no points like this exist, create *point objects (page 2–23)* at measurement extremes.
2. Choose 3D from the Snap Toggle flyout.
3. Right-click the Snap Toggle button to display the Grid and Snap Settings dialog. On the Snaps tab, set the snap type to the type you will use for measurement. For example, if you are going to use point objects for measurement, check the Pivot option so you can snap to the point objects' pivot points. Close the dialog.
4.  On the Create panel, click the Helpers button, then click Tape. Move the cursor over the first measurement point until the snap cursor appears, then click and drag to the other measurement point.
5. Read the length of the tape on the Parameters rollout.

Protractor

The *Protractor helper (page 2–26)* measures the angle between pivot points of two objects and the protractor object.

As with the tape helper, point objects and snap tools can be used to aid in precise measurement of angles.

To measure an angle with a protractor:

1.  On the Create panel, click the Helpers button, then click Protractor. Click in a viewport to place the protractor object.
2. On the Parameters rollout, click Pick Object 1 and select an object in your scene. The name of the selected object appears above the pick button. A line connects the protractor to the center of the object.
3. Click Pick Object 2 and pick a second object. The angle between the protractor and the two objects appears on the rollout.
4. To move either object while watching the Angle readout, turn on the Pin Stack button in the Modifier Stack rollout. This locks the stack to the Protractor's Parameters rollout.

Compass

The *Compass helper (page 2–27)* establishes an orientation for your scene. You create this object with a single click and drag to define the radius, usually in a Perspective or Top view. On its parameters rollout, you can adjust the apparent

radius of the compass rose. Like all helper objects, this object is for reference only and does not render.

Measure Utility

The *Measure Utility (page 2–52)* provides measurements of a selected object or shape.

To use the Measure utility:

1. On the Utilities panel, click Measure to display the Measure rollout.
2. Select any shape or object and read out the measurements. If you select multiple objects, the sum of the measurements is displayed.
3. Click New Floater to display a modeless Measure dialog with all the same information.
4. Use the Measure dialog to display the length of a spline, like a line or circle, as you create it.

Measure Distance Tool

The Measure Distance tool quickly calculates the distance between two points.

To measure the distance between two points:

1. Choose Tools menu > Measure Distance
2. Click in the viewport where you want to start measuring.
3. Click in the viewport where you want to measure to.

   A distance is returned in the Mini Listener. Additional information, detailing the distance along the X, Y, and Z coordinates is displayed in the status bar.

Measure Distance Tool

Tools menu > Measure Distance

The Measure Distance tool lets you quickly calculate the distance between two points, using just two clicks. The calculated distance is displayed in *display units (page 3–848)* in the status bar.

Procedure

To measure the distance between two points:

1. Choose Tools menu > Measure Distance.
2. Click in the viewport where you want to start measuring from.
3. Click again in the viewport where you want to measure to.

   A distance is returned in the status bar. Additional information, detailing the distance along the X, Y, and Z coordinates is displayed as well.

Select And Manipulate 

Main Toolbar > Select And Manipulate button

The Select And Manipulate tool lets you edit the parameters of certain objects, modifiers, and controllers by dragging "manipulators" in viewports.

Note: Unlike Select And Move and the other transforms, this button's state is non-exclusive. As long as Select mode or one of the transform modes is active, and Select And Manipulate is turned on, you can manipulate objects. However, you must turn off Select And Manipulate before you can select a manipulator helper.

You can add these custom manipulators to your scene:

Cone Angle Manipulator (page 2–27)

Plane Angle Manipulator (page 2–29)

Slider Manipulator (page 2–31)

These features have manipulators built in, which you can use to change parameters on these objects:

IK Solver Properties Rollout (HI Solver) (page 2–456)

Reaction Controllers (page 2–358)

Target Spotlight (page 2–1289)

Target Direct Light (page 2–1292)

mr Area Spotlight (page 2–1299)

Free Spotlight (page 2–1290)

Spotlight Parameters (page 2–1338)

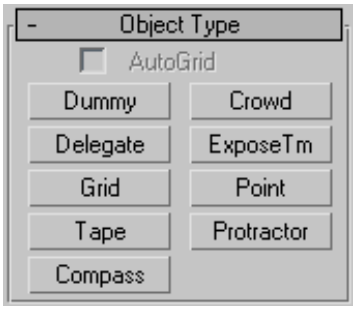
UVW Map Modifier (page 1–922)

Primitives (page 1–169) with a Radius parameter have a built-in manipulator for the radius.

Standard Helper Objects

Create panel > Helpers > Standard > Object Type rollout

Create menu > Helpers



The standard set of helper objects includes the following:

Dummy Helper (page 2–16)

Crowd Helper (page 2–1187)

Delegate Helper

Expose Transform Helper (page 2–17)

Grid Helper (page 2–20)

Point Helper (page 2–23)

Tape Helper (page 2–24)

Protractor Helper (page 2–26)

Compass Helper (page 2–27)

See also

Atmospheric Apparatus (page 3–304)

Camera Match Helper (page 2–1391)

Assembly Heads Helper (page 1–111)

Manipulator Helper Objects (page 2–27)

Particle Flow helpers: *SpeedByIcon (page 2–162)* and *Find Target (page 2–218)*

VRML97 Helper Objects (page 3–597)

reactor Helpers (page 2–715)

Dummy Helper

Create panel > Helpers > Standard > Object Type rollout > Dummy

Create menu > Helpers > Dummy

A Dummy helper object is a wireframe cube with a pivot point at its geometric center. It has a name but no parameters, you can't modify it, and it doesn't render. Its only real feature is its pivot, used as a center for transforms. The wireframe acts as a reference for transform effects.

The dummy object is used primarily in hierarchical linkages. For example, you can use a dummy object as a center of rotation by linking a number

of different objects to it. When you rotate the dummy, all of its linked children rotate with it. A dummy is often used this way to animate linked motion.

Another common usage of a dummy object is in the animation of target cameras. You can create a dummy and position a target camera within the dummy object. Then you can link both the camera and its target to the dummy, and animate the dummy with a path constraint. The target camera will follow the dummy along the path.

Dummy objects are always created as cubes. You can change the proportions of dummies by using non-uniform scaling, but avoid this on dummies that are within a hierarchical linkage; this can introduce unexpected results.

Procedure

To create a dummy object:

- Click Dummy and drag a cube to any convenient size.

Interface

Dummy objects don't have any parameter rollouts or settings.

Expose Transform Helper

Create menu > Helpers > Standard > ExposeTm

The Expose Transform helper exposes values of non-keyed objects for use in expressions and scripts. This lets riggers and animators access a select set of transforms for an object, and between an object and its parent, such as a bone in an IK chain. For example, using the Expose Transform helper, you could write a script or use parameter wiring that tests the angle of an character's arm bone that's controlled by IK (thus, has no keys), and once it exceeds a specific value, rotate the corresponding clavicle bone to deform the shoulder area.

In the viewports, the Expose Transform helper object looks just like a *Point helper (page 2–23)*.

Procedure

To use the Transform helper:

1. Click ExposeTm and choose the appearance characteristics from the Display rollout.
2. Adjust the size of the helper object by using the Size spinner on the Display rollout. Default=20.0
3. Click anywhere in your scene to add the helper object.

   The helper object's position has no bearing on its functionality.
4. On the Modify panel > Parameters rollout, click the Expose Node button (default label: None) and select the object whose transform values you want to expose.

   The object's name appears on the button, and its transform values appear in the fields on the Expose Values rollout.
5. Optionally, turn off Parent and choose a Local Reference Node object.

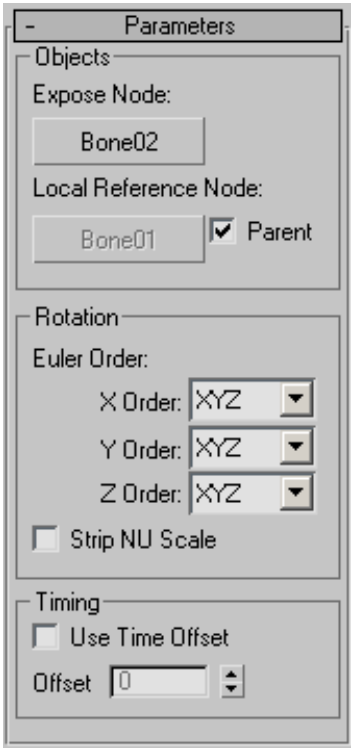
   This option is used by the Local Euler Angles, Distance To Reference, and Angle parameters. If the object has no parent and you don't specify a Local Reference Node object, these fields use the world center (0,0,0) as the reference node.
6. To use a transform value in a script or expression, first click the M button next to the parameter value.

   This copies the transform name, using MAXScript notation, to the copy buffer.
7. Paste the contents of the buffer to your script or expression.

A sample result of this is "$ExposeTransform01.localPositionX". 3ds Max interprets this as the local position, on the X axis, of the Expose Node object.

Interface

Parameters rollout



These settings let you specify the exposed node, a reference object other than the parent, and rotation and timing parameters.

Expose Node—The object from which the values are generated. Click the button, and then select the object. Thereafter, the object name appears on the button.

Local Reference Node—The object whose relationship with the Expose Node object is used to generate local data for rotation, distance, and angle. Click the button, and then select the object. Available only when Parent is off.

By default, this is the parent object, but you can turn off Parent and then specify a different object.

Parent—When on, the Local Reference Node is set automatically to the parent of the Expose Node. When off, you can pick an object to refer to for local data. Default=on.

This is used by the Local Euler Angles, Distance To Reference, and Angle parameters. If the object has no parent and you don't specify a Local Reference Node object, these fields use the world center (0,0,0) as the reference node.

Rotation Order—These three settings determine the order in which the Expose Transform helper will look for an Euler rotation.

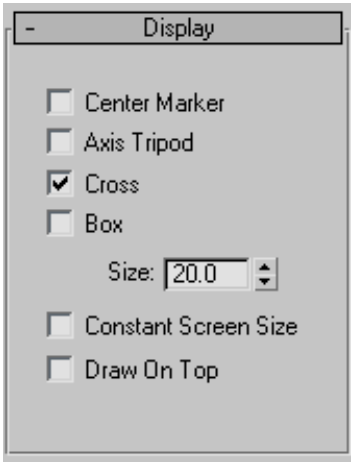
This parameter corrects for anomalies that can be introduced when a local rotation value is generated in relation to other rotation values. For example, when Z Order is set to XYZ, the Z value is generated in reference to X and then Y. Or when X Order is set to ZXY, the Z value is generated directly in reference to the parent, regardless of X and Y rotations.

Strip NU Scale—When on, removes any non-uniform scaling that could influence the rotation data. It does not strip non-uniform scaling from the object; only from the values generated.

Use Time Offset—When on, lets you specify a frame other than the current one from which to gather data. The Offset value is added to the current frame to derive the frame from which data is gathered.

For example, if the current frame is 20 and you wish to gather data from frame 15, turn on Use Time Offset and set Offset to -5.

Display rollout



Lets you specify how the helper object appears in the viewports. By default, the display type is set to Cross only, and Size is set to 20.0. You can activate multiple display settings if you like.

Center Marker—Displays a small X marker at the center of the helper object.

Axis Tripod—Displays a tripod axis indicating the position and orientation of the helper object.

Cross—Displays an axis-aligned cross.

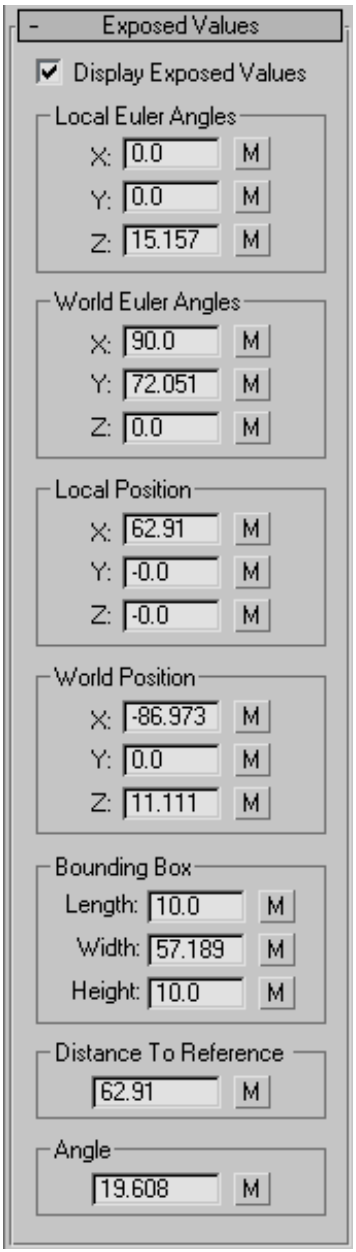
Box—Displays a small axis-aligned box at the center of the helper object.

Size—Sets the size for the helper object. Use this setting to minimize the helper object, or increase its size to aid in locating it. Default=20.0.

Constant Screen Size—Keeps the size of the helper object constant, regardless of how much you zoom in or out.

Draw On Top—Displays the helper object on top (or in front) of all other objects in the scene.

Exposed Values rollout



This rollout provides readouts of values for all transform values of the Expose Node object. Each value has a button labeled “M” next to it; click this

button to copy the MAXScript associated with the value to the clipboard, which you can then paste into a script.

Display Exposed Values—When on, transform values for the Expose Node object appear on this rollout, and update as the values change. When off, values don't appear, but the M buttons still function normally.

Local Euler Angles—Rotation values with respect to the parent or Local Reference Node object.

World Euler Angles—Rotation values in the world coordinate system.

Local Position—Position values with respect to the Expose Node object's local coordinate system.

World Position—Position values in the world coordinate system.

Bounding Box—The maximum dimensions of the object.

Distance to Reference—The distance between the Expose Node object and its parent or the designated Local Reference Node object. If neither exists, this shows the distance to the world center (0,0,0).

Angle—The angle between the Expose Node object and its parent or the designated Local Reference Node object.

Grid Helper

Create panel > Helpers > Standard > Object Type rollout > Grid

Create menu > Helpers > Grid

The grid, also called User Grid or Custom Grid, is a 2D parametric object with adjustments for overall size and grid spacing. You can move and orient user grids anywhere in world space.

You can create any number of grid objects in your scene. You name them when you create them and save them with the scene. You can delete them at any time.

Like other objects you create in the software, grid objects are placed on the grid of the current viewport. By default, this is a plane of the home grid, but it can also be another activated grid object.

Note: You can use *AutoGrid (page 2–7)* to create a temporary or permanent user grid off the surface of any object.

Using the Grid Object

You can use the grid object as a construction plane on all three orthographic planes in both directions. Here's how it works:

In the Parameters rollout of each grid object is a Display group with three options: XY Plane, YZ Plane, and ZX Plane. These determine which of the three planes of the grid object is displayed in the viewport.

When you activate a grid object, the displayed plane is the construction plane for all viewports. When you create a grid viewport, you can choose from one of six orthographic views (Front, Top, Left, and so on), or you can choose a special Display Planes grid viewport. (When you press the G key to create a grid viewport, the Display Planes type becomes the default.) The Display Planes type of grid viewport always displays the plane chosen by the three option buttons under Display. Thus, as you switch between XY Plane, YZ Plane, and ZX Plane, the view through the grid viewport switches accordingly, and objects created in that viewport are created on the displayed plane.

When you right-click a viewport label (or go to the Layout tab of the Viewport Configuration dialog), you can choose six additional types of grid viewports, based on the six orthographic views.

These are available by a cascading Grid menu that provides Left, Right, Front, Back, Top, Bottom, and Display Planes. Each of the orthographic directions is local to the grid object, regardless of its orientation in the scene.

When you choose a specific orthographic grid viewport (as opposed to the Display Planes viewport), the construction of objects in that viewport is on the plane specified in the viewport title regardless of the displayed plane of the grid object.

You can create more than one viewport based on the same grid object, with each using a different plane. For example, you can have a Grid (Front) viewport and a Grid (Top) viewport, as well as a Grid (Display Planes) viewport.

When you deactivate a grid object, its remaining viewports show the assigned orthographic view. Thus, a Grid (Front) viewport becomes a Front viewport, for example. A Grid (Display Planes) viewport always reverts to a Top view, regardless of the currently displayed plane.

Scaling Not Advised

As a rule, don't use scaling to resize a grid object. Scaling enlarges or reduces the apparent size of the grid object but has no effect on grid spacing. A sphere 20 units in radius created on a grid object appears smaller than another 20-unit sphere created on a scaled-up version of the same grid.

If you want to change the actual size of the grid object, select it and go to the Modify panel > Parameters rollout > Grid Size group, and change the Length and Width settings.

Locating Grid Tools

Grid tools are spread throughout the 3ds Max interface.

- In the Views menu, under Grids, are the commands to activate home and user grids. The same tools are available with a right-click anywhere in the viewport when a grid helper is selected.
- Choose Customize menu > Grid and Snap Settings to display the *Grid and Snap Settings dialog (page 2–41)*. Two tabs are devoted to Grid tools, one for home, another for user grids. You can change parameters of any active grid using the Modify panel.

See also

Viewing Grid Objects (page 2–6)

Procedures

To create a grid object:

1.  Click Create panel > Helpers > Standard > Object Type rollout > Grid.

   A Parameters rollout appears on the Create panel.
2. In a viewport, drag a rectangle and release the mouse button. This creates and selects a grid object, which appears in white wireframe, divided into four quadrants with coordinate axes at the center.

While the newly created grid object is still selected, you can change its settings on the Parameters rollout.

You can also create a grid object during object creation. Turn on AutoGrid, then press Alt during object creation. A grid is created at the same time as the object and remains displayed and active. See *AutoGrid (page 2–7)* for more information.

To activate a grid object:

A grid object requires activation before use. Standard selection doesn't activate it unless you turn on the option to do so (see *User Grids (page 2–51)*).

Only one grid can be active for construction at a time, whether it's the home grid or a grid object. Activating a user grid "deactivates" the home grid.

Activating a grid object enables options to reactivate the home grid on the Views menu > Grids submenu and the Quad menu.

If you have more than one grid object in your scene, you have to activate each one separately. Select the grid object you want to make active and follow the same procedure. Activating another grid object deactivates the current one.

1. Select a grid object.
2. Do one of the following:
   - From the Views menu, choose Grids > Activate Grid Object.
   - Right-click the selected grid object and choose Activate Grid from the quad menu.

     The grid object changes to show its internal grid structure. Except for its main axes, the home grid disappears in all viewports.

To return to the home grid, do one of the following:

- From the Views menu, choose Grids > Activate Home Grid.
- Right-click the selected grid object and choose Activate Home Grid from the quad menu.

  This deactivates the grid object and returns the home grid in all views.

  If you delete an activated grid object, the home grid also reactivates.

  You can assign a keyboard shortcut to Activate Home Grid in the *Keyboard panel (page 3–793)* of the Customize User Interface dialog. This is useful if you need to move back and forth between different grids.

To use a grid object as construction plane:

When activated, a grid object replaces the home grid as the frame of reference for creating objects.

An activated grid object creates a true plane in 3D space. No matter how small an activated grid object appears on screen, its XY plane is effectively infinite, just as if it were the XY plane of the home grid.

1. Activate the grid object.
2. Create any category of object.

   3ds Max creates the object directly on the plane of the grid object, with the object's Z axis perpendicular to the plane.

See *Align (page 1–462)* for details on aligning objects and grids.

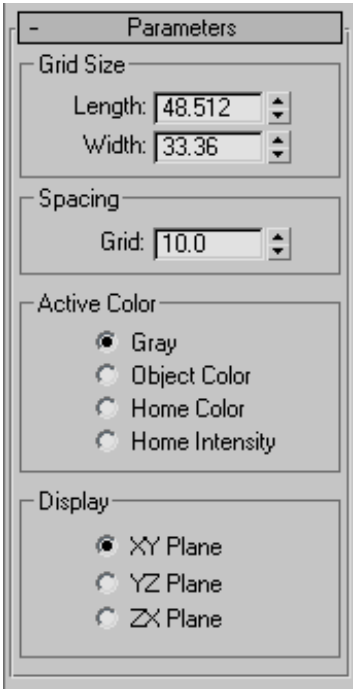
Like other objects in the software, you can move and rotate grid objects freely using standard transformation methods. These transforms, along with alignment, are essential in positioning a construction plane in 3D space.

To nudge a grid object up or down:

1. Activate the grid object.
2. Press the + or — keys on the numeric keypad to move the grid object up or down.

   The Grid Nudge Distance is controlled in the *Viewports panel (page 3–821)* of the Preferences dialog.

Interface



Grid Size group

Sets the overall size of the grid object. This size determines the extents of a viewport set to the grid object. It doesn't affect the useful limits of the grid, which extend infinitely.

Length/Width—Specifies the length and width of the grid.

Grid Spacing group

Grid—Specifies the size of the smallest square in the visible grid. This setting appears on the status line when the grid is activated.

Note: You can set Grid Spacing when a grid is selected, but you won't see the grid spacing until the grid is activated.

Active Color group

Determines the color used to draw the grid in viewports when it's not selected.

Gray—The active grid object is two shades of gray.

Object Color—The main grid lines use the assigned object color, while the secondary lines use a lighter intensity.

Home Color—The grid object uses the home grid color assigned via the *Customize User Interface dialog (page 3–792)*.

Home Intensity—The grid object uses the grid intensity settings assigned to the home grid in the *Customize User Interface dialog (page 3–792)*.

Display group

XY Plane, YZ Plane, ZX Plane—Determines which of the three planes of the grid object are displayed in the viewport.

Point Helper

Create menu > Helpers > Point

Point provides a specific location in 3D space that can be used as a reference or by other program functions.

Procedure

To create a point in space:

1. Click Point and check the type of display from the Parameters rollout.
2. Adjust the size of the point object by using the Size spinner in the Parameters rollout. Default=20.0
3. Left-click and drag anywhere in your scene.

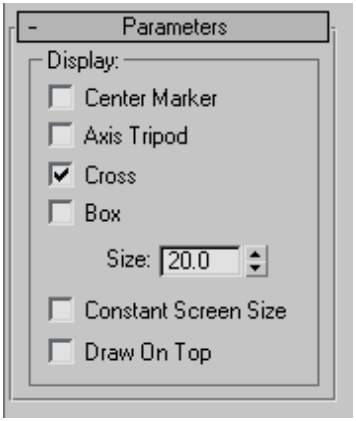
   The point follows the cursor until you release the mouse button, indicating the current location of the point object.

4. Move the cursor to where you want the point object and release the mouse button.

   The point object appears using the display setting you chose.

   You can move and rotate the point as needed using standard transformation methods.

Interface



Center Marker—Displays a small X marker at the center of the point object.

Axis Tripod—Displays a tripod axis indicating the position and orientation of the point object. The axis remains visible when the point object is no longer selected.

Cross—Displays an axis-aligned cross.

Box—Displays a small axis-aligned box at the center of the point object.

Size—Sets the size for the point object. Use this to minimize the point object, or increase its size to aid in locating it. Default=20.

Constant Screen Size—Keeps the size of the point object constant, regardless of how much you zoom in or out.

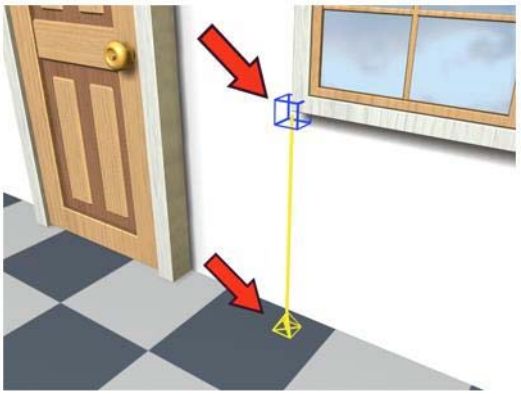
Draw On Top—Displays the point object on top of all other objects in the scene.

Note: You can also choose to activate multiple display settings.

Tape Helper

Create panel > Helpers > Standard > Object Type rollout > Tape

Create menu > Helpers > Tape Measure



Using a tape helper to measure the distance from floor to window frame.

Tape provides an on-screen "tape measure" for determining and setting distances. The tape is composed of two named objects, in the same way targeted lights and cameras are. By default, these are named Tape## and Tape##.Target. The tape and target icons are connected by a line representing the current distance between them.

You can drag a length (the default), or enter a specific length by turning on Specify Length. You can snap or align either end of the tape or link to objects in your scene. Deleting either end deletes the tape.

Note: When you add a Tape helper object, the software automatically assigns a *Look At controller (page 2–344)* to it, with the tape's target object assigned as the Look At target. You can use the controller settings on the Motion panel to assign any other object in the scene as the Look At target.

Selection

You can select a tape-measure object from either end, or from the middle. When you click the connecting line, both ends of the tape object become selected so you can move them as one. Note that the same is true of target cameras and lights.

Note: Line selection is available only when clicking with the mouse. Region selection doesn't work for this.

Tip: The Length readout of the tape object is displayed in the Modify panel only when you select the tape end of the object (as opposed to the target or link line). You can lock the Length display by activating the Pin Stack button while the tape end is selected, and then proceed to adjust any part of the tape object or any other object in your scene.

See also

Measure Distance Tool (page 2–15)

Procedures

To measure a distance:

1. Click Tape and drag from one point to create the tape, and then drag to a second point and release the mouse button to create the target.

   The distance between tape and target appears in the Length field. The number in the Length field is gray to indicate that this is a read-only measurement.

   Note: To snap the tape ends to object vertices, edges and so on, turn on *Snap (page 2–35)* with the appropriate snap settings activated.

2. Move either end to a new location. The line between stretches to the new distance, shown in the Length field.

   Note: For accurate results when moving a tape endpoint with Snap, you must turn on Pivot in the Snap Settings panel as well as the types of sub-object you want to snap to; see *Snaps (page 2–41)*. This is because the endpoints are actually pivots.

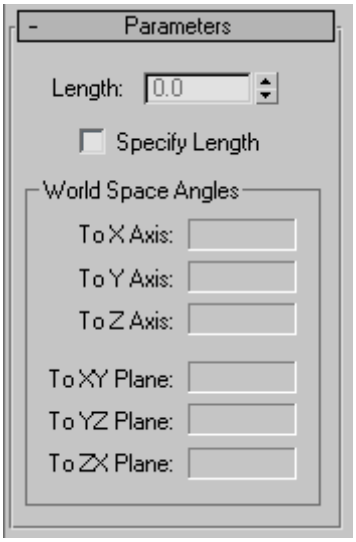
To set a distance:

- Turn on Specify Length in the Parameters rollout, enter a length, and then create the tape.

  The connector line is the length you specified, although you can still set the target to any offset by dragging. You can reorient this line as required.

  Note: Clearing Specify Length will adjust the connector line length to meet the target and the new length appears in the Length field.

Interface



Parameters rollout

Length—Specifies the length of the tape object. You must select the Specify Length check box to enable this option.

Specify Length—Activates the Length parameter setting. Otherwise the tape object's length will be specified by clicking and dragging in the viewport.

World Space Angles group

To X/Y/Z Axis—Displays the angle of the tape object relative to each axis in world space.

To XY/YZ/ZX Plane—Displays the angle of the tape object relative to each of these home planes.

Protractor Helper

Create panel > Helpers > Standard > Object Type rollout > Protractor

Create menu > Helpers > Protractor

Protractor measures the angle between a point and any two objects in your scene.

The names of the two objects appear above their respective buttons, and the angle formed at the protractor object between the pivot points of the two objects is displayed in the Angle field. Lines are drawn between the protractor and the two selected objects.

Note: The line to the first selected object uses the Target Line color, and the line for the second selected object uses the color of the protractor. The Target Line color is set with Customize User Interface dialog > *Colors panel (page 3–799)* > Elements drop-down list > Gizmos group. The Protractor color is set by choosing the Tape Object in Customize User Interface dialog > *Colors panel (page 3–799)* > Elements drop-down list > Objects group.

You can reselect the protractor at any time, enter the *Modify panel (page 3–758)*, and pick new objects to measure. You can move either of the two objects or the protractor itself to change the angle.

Note:  If you need to watch the Angle readout while moving either of the objects, first select the protractor, then click the Pin Stack button on the Modify panel. You can now select and move the other objects while the stack remains pinned to display the protractor angle.

Procedure

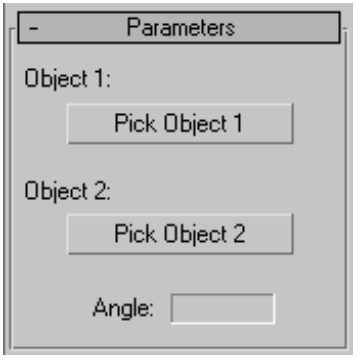
To use the protractor:

1. After selecting Protractor in the Create panel > Helpers section, click in the viewport to place and create the protractor icon. You can also drag to move it before releasing the mouse button.
2. Click the Pick Object 1 button, and then select one of the objects in the scene.

   The name of the selected object appears above the button.
3. Click Pick Object 2, and then select a second object in the scene.

   The Angle parameter shows the angle formed by lines from the protractor object to the two picked objects.

Interface



Parameters rollout

Pick Object 1—Click this and select the first of the two objects whose angle you want to measure.

Pick Object 2—Click this and select the second of the two objects whose angle you want to measure.

Angle—Displays the angle formed by the lines from the protractor to the two objects.

Compass Helper

Create panel > Helpers > Standard > Object Type rollout > Compass

Create menu > Helpers > Compass

Compass displays as a non-rendering compass rose, with indicators for North, South, East, and West. A compass is part of a *Daylight or Sunlight system (page 1–418)*; you create a compass automatically when you create sunlight. In a Sunlight system, the orientation of the compass indicates the orientation of the scene, relative to the path of the sun (the ecliptic).

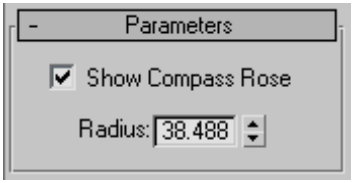
Use the Compass helper object if you want to create an orientation symbol in your scene, but don't need to create a sunlight system.

Procedure

To create a compass object:

1. Choose Compass in the Create panel > Helpers section.
2. Click and drag to define the radius of the compass.
3. Select and rotate the compass object to orient it.

Interface



Parameters rollout

Show Compass Rose—Displays the non-rendering compass rose in the viewport.

Radius—Specifies the radius of the compass object.

Manipulator Helper Objects

Create panel > Helpers > Manipulators > Object Type rollout

Create menu > Helpers > Manipulators

Manipulator helpers are objects you can create to manipulate other objects. They let you add customized controls to your scene that provide visual feedback in viewports. To control other objects, manipulators use the *parameter wiring feature (page 2–411)*.

Manipulators can be especially useful when you create a scene that will be used by more than one animator.

There are three kinds of manipulator helpers:

Cone Angle Manipulator (page 2–27)

Plane Angle Manipulator (page 2–29)

Slider Manipulator (page 2–31)

See also

Wire Parameters (page 2–411)

Parameter Wiring Dialog (page 2–412)

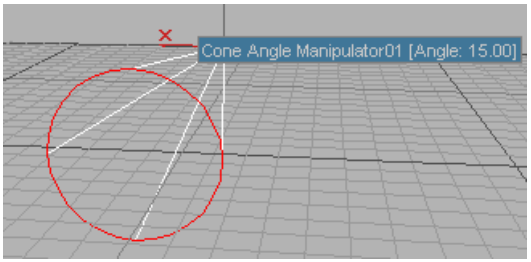
Select And Manipulate (page 2–15)

Cone Angle Manipulator

Create panel > Helpers > Manipulators > Object Type rollout > Cone Angle button

Create menu > Helpers > Manipulators > Cone Angle

The Cone Angle manipulator is a cone whose base you can adjust. By wiring its Angle value to a parameter of another object, you can create a custom control, with visual feedback, within a scene.



Cone angle manipulator

Cone angle manipulators are used by a spotlight's Hotspot and Falloff controls.

Procedures

To create a cone angle manipulator:

1.  In the Create panel, go to Helpers and choose Manipulators from the drop-down list.
2. Click to turn on Cone Angle.
3. Drag in a viewport, and then release the mouse.

   The cone angle manipulator is created along the negative Z axis of the viewport in which you drag. In other words, its initial position points away from you.

To change the angle of a cone angle manipulator:

1.  Turn on Select And Manipulate.

   Tip:  If you are still in the Create panel, turn on Select Object before using the manipulator. Otherwise, it is too easy to create a new one.
2. In a viewport, move the mouse over the manipulator.

   The manipulator turns red when the mouse is over it and it is available to be manipulated. Also, a tooltip appears, showing the name of the manipulator and its current value.
3. Drag the base of the cone angle manipulator's cone.

   The Angle value changes as you drag the manipulator. This value can range between 0.0 and 360.0.

To select a cone angle manipulator:

1.  Turn on Select Object.

    You can also use Select And Move, Select And Rotate, and the scale options on a cone angle manipulator.
2.  Turn off Select And Manipulate.

   If you don't turn off Select And Manipulate, you adjust the manipulator's Angle value without affecting its properties.
3. Adjust the properties of the selected manipulator.

To connect a cone angle manipulator so it controls another object:

1.  Select the cone angle manipulator.

   Note:  Select And Manipulate must be *off* before you can select the manipulator.
2. Choose Animation > Wire Parameters > Wire Parameters.

   You can also right-click the manipulator and choose Wire Parameters from the Transform (lower-right) quadrant of the quad menu.

   A pop-up menu appears over the manipulator.
3. In the pop-up menu, choose Object > Angle.
4. Drag to the object you want to manipulate, and click it.

   A pop-up menu appears over the object.

5. In the pop-up menu, choose Object, then the name of the parameter you want to manipulate.

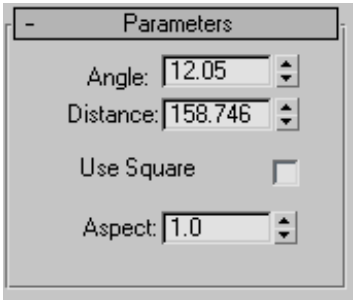
   The *Parameter Wiring dialog (page 2–412)* appears.

6. In the Parameter Wiring dialog, make sure the direction goes from the manipulator to the object (or both directions), and then click Connect.

7. Close the Parameter Wiring dialog.

   Now when you turn on Select And Manipulate and use the manipulator, the object updates under the manipulator's control.

Interface



Angle—The initial angle of the manipulator.

Distance—The length of the manipulator, in 3ds Max units. Default=the distance of mouse drag when the manipulator was created.

Use Square—When on, the base of the cone is square or rectangular, rather than circular. Default=off.

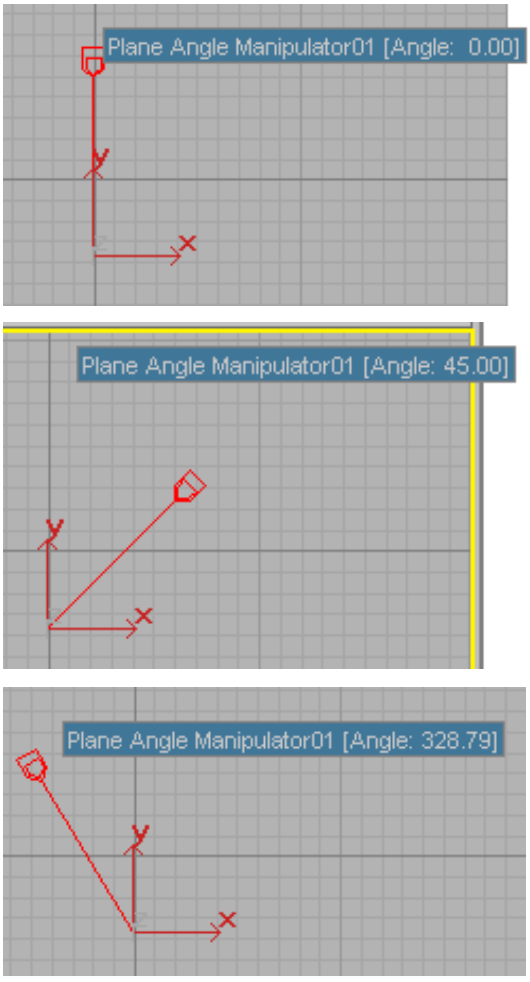
Aspect—When Use Square is on, adjusts the aspect ratio of the rectangular cone base. Default=1.0.

Plane Angle Manipulator

Create panel > Helpers > Manipulators > Object Type rollout > Plane Angle button

Create menu > Helpers > Manipulators > Plane Angle

The Plane Angle manipulator looks like a lever or joystick. By wiring its Angle value to a parameter of another object, you can create a custom control, with visual feedback, within a scene.



Plane angle manipulator at different angles

Procedures

To create a plane angle manipulator:

1.  In the Create panel, go to Helpers and choose Manipulators from the drop-down list.
2. Click to turn on Plane Angle.
3. Drag in a viewport, and then release the mouse.

   The plane angle manipulator is always created vertically, along the Y axis of the viewport in which you drag.

To change the angle of a plane angle manipulator:

1.  Turn on Select And Manipulate.

   Tip:  If you are still in the Create panel, turn on Select Object before using the manipulator. Otherwise, it is too easy to create a new one.
2. In a viewport, move the mouse over the manipulator.

   The manipulator turns red when the mouse is over it and it is available to be manipulated. Also, a tooltip appears, showing the name of the manipulator and its current value.
3. Drag the handle of the plane angle manipulator.

   The Angle value changes as you drag the manipulator. This value can range between 0.0 and 360.0.

To select a plane angle manipulator:

1.  Turn on Select Object.

    You can also use Select And Move, Select And Rotate, and the scale options on a plane angle manipulator.
2.  Turn off Select And Manipulate.

   If you don't turn off Select And Manipulate, you adjust the manipulator's Angle value without affecting its properties.
3. Adjust the properties of the selected manipulator.

To connect a plane angle manipulator so it controls another object:

1.  Select the plane angle manipulator.

   Note:  Select And Manipulate must be *off* before you can select the manipulator.
2. Choose Animation > Wire Parameters > Wire Parameters.

   You can also right-click the manipulator and choose Wire Parameters from the Transform (lower-right) quadrant of the quad menu.

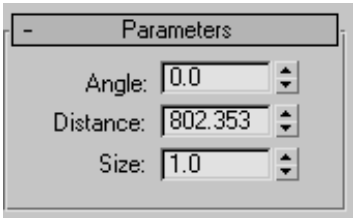
   A pop-up menu appears over the manipulator.
3. In the pop-up menu, choose Object > Angle.
4. Drag to the object you want to manipulate, and click it.

   A pop-up menu appears over the object.
5. In the pop-up menu, choose Object, then the name of the parameter you want to manipulate.

   The *Parameter Wiring dialog (page 2–412)* appears.
6. In the Parameter Wiring dialog, make sure the direction goes from the manipulator to the object (or both directions), and then click Connect.
7. Close the Parameter Wiring dialog.

   Now when you turn on Select And Manipulate and use the manipulator, the object updates under the manipulator's control.

Interface



Angle—The angle of the manipulator, from 0.0 to 360.0 (both values are perpendicular in the Y axis of the viewport where you created the manipulator, unless you have rotated the manipulator object). Default=0.0.

Distance—The length of the manipulator, in 3ds Max units. Default=the distance of mouse drag when the manipulator was created.

Size—The size of the manipulator's handle, in 3ds Max units. Default=1.0.

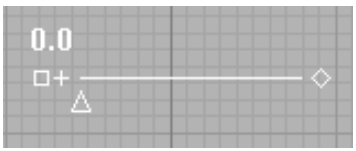
Slider Manipulator

Create panel > Helpers > Manipulators > Object Type rollout > Slider button

Create menu > Helpers > Manipulators > Slider

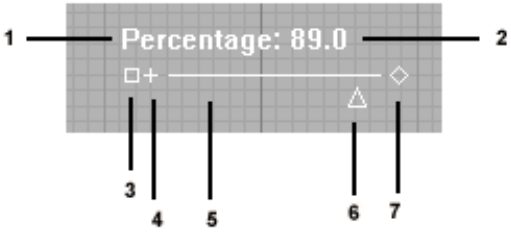
The Slider manipulator is a graphic control that appears in the active viewport. By wiring its Value to a parameter of another object, you can create a custom control, with visual feedback, within a scene.

Note: Because the slider manipulator is a custom control, its Value has no inherent unit. It takes on the unit of the parameter to which it is wired. When the minimum is 0.0 and the maximum is 100.0, the slider Value can represent a percentage.



Slider manipulator

Using a Slider Manipulator



Slider manipulator components:

1. Label

2. Value

3. Move

4. Show/hide

5. Slider bar

6. Adjust value

7. Change width

The slider appears in the same location in whichever viewport is active.

You can adjust the slider's value by dragging the triangle below the slider bar. The triangle turns red when you move the mouse over it. It represents the slider's value, and the value changes as the triangle is dragged. You can also make the following changes without going into the Modify panel:

- Dragging the small square at the left moves the slider.
- Dragging the diamond at the right changes the width of the slider.
- Clicking the plus sign (next to the move icon at the left) hides all of the slider except for the label and the move icon.

Like the Value triangle, these components also turn red when you move the mouse over them.

Procedures

To create a slider manipulator:

1.  In the Create panel, go to Helpers and choose Manipulators from the drop-down list.
2. Click to turn on Slider.
3. Click a viewport.

   The slider manipulator is created with its default width of 100 units. It will appear at the same viewport location in whichever viewport is active.

To change the value of a slider manipulator:

1.  Turn on Select And Manipulate.

   The slider turns green.

   Tip:  If you are still in the Create panel, turn on Select Object before using the manipulator. Otherwise, it is too easy to create a new one.
2. In a viewport, move the mouse over the triangle below the slider bar.

   The triangle turns red when the mouse is over it and it is available to be dragged.
3. Drag the value triangle of the slider manipulator.

   The slider's value display changes as you drag the triangle.

To move a slider manipulator:

1.  Turn on Select And Manipulate.
2. In the active viewport, move the mouse over the move icon, which is the small square at the left of the slider, below the sliders label (if there is one) and value display.

   The square turns red when the mouse is over it and it is able to be dragged.
3. Drag the box to move the slider.

   Unlike angle manipulators, transforms have no effect on sliders.

To connect a slider manipulator so it controls another object:

1.  Select the slider.

   Note:  Select And Manipulate must be *off* before you can select the slider.
2. Choose Animation > Wire Parameters > Wire Parameters.

   You can also right-click the manipulator and choose Wire Parameters from the Transform (lower-right) quadrant of the quad menu.

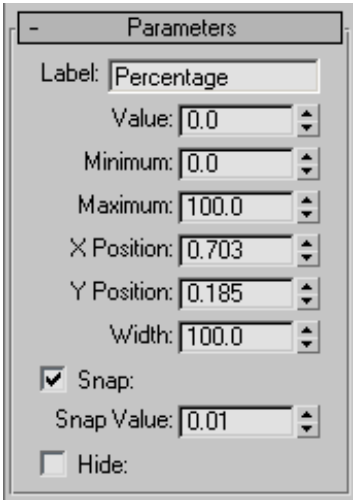
   A pop-up menu appears over the manipulator.
3. In the pop-up menu, choose Object > Value.
4. Drag to the object you want to manipulate, and click it.

   A pop-up menu appears over the object.
5. In the pop-up menu, choose Object, then the name of the parameter you want to manipulate.

   The *Parameter Wiring dialog (page 2–412)* appears.
6. In the Parameter Wiring dialog, make sure the direction goes from the slider to the object (or both directions), and then click Connect.
7. Close the Parameter Wiring dialog.

   Now when you turn on Select And Manipulate and use the slider, the object updates under the slider's control.

Interface



Label—The slider name that appears in viewports. Default=none.

Value—The value of the slider, based on the position of the slidable triangle. Default=0.0.

Minimum—The minimum possible value of the slider. Default=0.0.

Maximum—The maximum possible value of the slider. Default=100.0.

When the minimum is 0.0 and the maximum is 100.0, the slider Value can represent a percentage.

X Position—The slider's X location in the active viewport. Default=Viewport X location you clicked when you created the slider.

Y Position—The slider's Y location in the active viewport. Default=Viewport Y location you clicked when you created the slider.

Width—The slider's width, in 3ds Max units. Default=100.0.

Snap—When on, the slider "snaps" to incremental values determined by the Snap Value setting. Default=on.

Snap Value—The increment used by the slider when Snap is on. Default=0.01.

Hide—When on, hides all of the slider except for the label and the move and show/hide components. Default=off.

Grid Commands

Views menu > Grids

The Grids submenu provides functions to manipulate the *home grid (page 3–952)* and *grid objects (page 3–951)*.

The Snaps system uses grids both as visual references and as physical snap references. In addition, grids serve as construction planes, since the objects you create are built on the active grid in the active viewport.

As a default, the software uses a construction grid called the home grid. The home grid is the basic reference system, and is defined by three fixed planes on the world coordinate axes. It provides ready-to-use construction planes, intersecting at right angles through the origin. When the home grid is displayed, each *orthographic viewport (page 3–986)* is coplanar (parallel) with one of the grid faces.

A custom grid, or grid object, is a helper object that you can use in place of the home grid. While the home grid is three intersecting planes whose orientation is fixed in world space, a grid object is a free-floating, 2D grid that you can position and rotate anywhere, letting you construct objects at any angle in world space. You can also automatically create grids on the surface of objects by using *AutoGrid (page 2–7)*.

The Grids submenu contains the following commands:

Show Home Grid (page 2–34)

Activate Home Grid (page 2–34)

Activate Grid Object (page 2–34)

Align Grid to View (page 2–35)

Show Home Grid

Views menu > Grid > Show Home Grid

Right-click a viewport label. > Show Grid

Keyboard > G

Show Home Grid toggles the display of the *home grid (page 3–952)* in the active viewport.

Grids are used both as visual references and as physical snap references by the Snap system. In addition, grids are construction planes, since the objects you create are built on the active grid in the active viewport.

As a default, there is a home construction grid called home grid. The home grid is the basic reference system, and is defined by three fixed planes on the world coordinate axes. It provides ready-to-use construction planes, intersecting at right angles through the origin. When the home grid is displayed, each *orthographic viewport (page 3–986)* is coplanar (parallel) with one of the grid faces.

Procedure

To toggle the display of the home grid, do one of the following:

- Choose Views menu > Grids, and click Show Home Grid.
- Right-click a viewport label and click Show Grid on the viewport right-click menu.
- Press G on the keyboard.

Activate Home Grid

Activate a grid other than the Home grid. > Views menu > Grids > Activate Home Grid

This command activates the *home grid (page 3–952)* as the construction grid, and deactivates the active *grid object (page 3–951)*.

Procedure

To return to the home grid, do one of the following:

- Choose Views menu > Grids > Activate Home Grid.
- Right-click the selected grid object, and choose Activate Home Grid from the quad menu.

  This deactivates the grid object and activates the home grid in all views.

Activate Grid Object

Select a grid object. > Views menu > Grids > Activate Grid Object

This command activates the selected *grid object (page 3–951)* as the construction grid and deactivates the *home grid (page 3–952)*. This menu item is available only when a grid object is selected.

Procedure

To activate a grid object:

1. Select a grid object.
2. Do one of the following:
   - Choose Views menu > Grids > Activate Grid Object.
   - Right-click the selected grid object and choose Activate Grid from the quad menu.

     The grid object changes to show its internal grid structure. Except for its main axes, the home grid disappears in all viewports.

Align Grid to View

Activate a grid object. > Views menu > Grids > Align Grid to View

This command aligns a *grid object (page 3–951)* with the current view. The grid object is aligned so that it's coplanar with the view and oriented with the top of the grid at the top of the view.

Procedure

To align a grid object with a viewport:

1. Activate a grid object.
2. Click or right-click in a viewport.
3. Choose Views menu > Grids > Align Grid to View.

   The grid object rotates to align itself with the selected view. The grid object is now planar, or parallel, with the viewport.

Snap Commands

main toolbar > Snap buttons

Snaps help you precisely control the dimensions and placement of objects when you create them or transform them. There is also a Spinner snap for controlling the entry of values.

You can use objects within an XRef scene as snap references.

The snap commands listed below and described in this section are found on the main toolbar (except for Ortho and Polar; see their topics for details).



2D Snap, 2.5D Snap, 3D Snap (page 2–35)


Angle Snap Toggle (page 2–37)



Percent Snap Toggle (page 2–38)


Spinner Snap Toggle (page 2–38)


Ortho Snapping Mode (page 2–38)


Polar Snapping Mode (page 2–39)

2D Snap, 2.5D Snap, 3D Snap

Main toolbar > 2D Snap, 2.5D Snap, or 3D Snap on the Snaps Toggle flyout

Keyboard > S

The buttons on the Snaps Toggle flyout provide control over the range of 3D space where snaps are active.



2D, 2.5D, 3D Snaps flyout

Object Snapping

Object snapping lets you snap to specific portions of existing geometry during creation and transforms of objects or sub-objects. You can also snap to the grid, and you can snap to tangents, midpoints, pivot points, face centers and other options.

The mode you choose maintains its state when you switch levels.

Transforming Around Snap Points

When snapping is on and *Auto Key (page 3–717)* is off, rotations and scales occur about the snap point. For example, if you're using Vertex snapping and you're rotating a box, you can rotate it about any of its corner vertices. See *To use snaps to move an absolute distance: (page 2–36)* below.

When Auto Key is on and either Select And Rotate or Select And Scale is active, the Snaps Toggle button is disabled, and rotation and scaling take place about the pivot point of the object.

Relative and Absolute Snaps

The mouse snaps absolutely to the active snap types. Relative snapping is also available. By default, when you turn on snaps, only the Grid Points snap type is active. If you perform a Move operation, the cursor will snap to the grid, but you can select and pick up an object that's not aligned to the grid. When you move the object and snap it to grid points, the object's original position relative to the grid is maintained as an offset to each grid point, providing relative snapping.

To move the same object in an absolute fashion, add the Vertex snap type. Then, when you select the object to move, you select one of its vertices, and snap that vertex to any grid point, resulting in an absolute snap.

You can also achieve variations on relative snaps by turning on the Selection Lock Toggle button in the status line.

A wide variety of snap types are available from the *Snaps dialog (page 2–41)*, that you use to activate different snap types as you work.

Procedures

To turn snap off during a transform:

- Press S to toggle snapping off when it gets in the way. Press S again to turn it back on.

To use snaps to move a relative distance:

1. Turn on Snaps with the S key, or by clicking the Snaps Toggle button.
2.  Lock your selection set with the SPACEBAR, or by clicking Lock Selection on the status bar.
3. Wherever you click in the viewport, the snap will stay relative to the distance of your cursor to the object.

To use snaps to move an absolute distance:

1. Turn on Snaps Toggle.
2. Click the selection. Snap uses the point you click for the snap source.
3. Move to whatever target snap you desire. The object will snap to an absolute distance.

To rotate a box around a vertex using snaps:

1. Make sure the Auto Key button is off.
2. Select the box.
3.  On the Main toolbar, click Select And Rotate.
4. Turn Snaps on by pressing S on the keyboard.
5. On the Customize menu, choose Grid And Snap Settings. Turn on Vertex and turn off Grid Points.
6.  Lock your selection set by clicking Lock Selection on the status bar.
7. On the toolbar, choose Use Transform Coordinate Center (hold the mouse down on Use Pivot Point Center to expose the flyout).
8. Move your cursor over any vertex in the box. The blue snap cursor will appear, and then you can rotate the box around that vertex.

Interface

There are three snap modes:


2D Snap—The cursor snaps only to the active construction grid, including any geometry on the plane of that grid. The Z axis, or vertical dimension, is ignored.


2.5D Snap—The cursor snaps only to the vertices or edges of the projection of an object onto the active grid.

Suppose you create a *grid object (page 3–951)* and make it active. You then position the grid object so you can see through the grid to a cube further off in 3D space. Now with 2.5D set, you can snap a line from vertex to vertex on the distant cube, but the line is drawn on the active grid. The effect is like holding up a sheet of glass and drawing the outline of a distant object on it.


3D Snap—This is the default. The cursor snaps directly to any geometry in 3D space. 3D snapping lets you create and move geometry in all dimensions, ignoring the construction plane.

Right-click this button to display the *Grid and Snap Settings dialog (page 2–41)*, which lets you change snap categories and set other options.

Angle Snap Toggle 

Main toolbar > Angle Snap Toggle

Keyboard > A

Angle Snap Toggle determines the incremental rotation for a number of features, including the standard Rotate transform. As you rotate an object (or group of objects), the object moves around a given axis in the increment you set.

Angle Snap Toggle also affects Pan/Orbit camera controls, *FOV (page 3–937)* and Roll camera settings, and *Hotspot/Falloff (page 3–954)* spotlight angles.

Procedures

To turn angle snap on:

- Click Angle Snap Toggle on the Main toolbar. When turned on, angle snap affects all rotational transforms.

To rotate an object an even number of degrees:

- Turn on Angle Snap Toggle and rotate the object.

  The rotations take place in 5-degree increments as a default.

To rotate an object a precise degree of rotation, do one of the following:

- Click Select And Rotate, then right-click it to display the *Transform Type-In dialog (page 1–431)*. Enter the exact rotation you want.
- Right-click Angle Snap Toggle and click the Options tab on the dialog that is displayed. Set the Angle value in the General group to the precise degree of rotation you need, then rotate the object. Rotation snaps to the angle increment you specified.

Interface

The angle increment is set on the Options panel of the *Grid and Snap Settings dialog (page 2–41)*. Right-click the Angle Snap Toggle button to display the Options panel of the Grid and Snap Settings dialog. The default is 5 degrees.

Percent Snap Toggle 

Main toolbar > Percent Snap Toggle

Keyboard > Ctrl+Shift+P

Percent Snap Toggle increments scaling of objects by the specified percentage.

Interface

The snap percent increment is set in the Grid And Snap Settings dialog. The default is 10 percent. Right-click the Percent Snap Toggle button to display the *Grid and Snap Settings dialog (page 2–41)*.

This is a general-purpose snap system that applies to any operation involving a percentage, such as scaling or squashing.

Spinner Snap Toggle 

Main toolbar > Spinner Snap Toggle

Spinner Snap Toggle sets the single-click increment or the decrement value for all of the spinners in 3ds Max.

Procedure

To set and toggle spinner snap:

1. Do one of the following:
   - Choose Customize menu > Preferences.
   - Right-click the Spinner Snap Toggle button on the main toolbar.

     Either method opens the Preference Settings dialog > General panel. The two controls for spinner snap are in the Spinners group on this panel.
2. Set a value in the Snap field.
3. Turn on Use Snap.

   When you exit the dialog, the Spinner Snap Toggle button is turned on.
4. As you work, use the Spinner Snap Toggle button to toggle the alternate setting.

Interface

The amounts for Spinner Snap are controlled by settings on the *General panel (page 3–815)* of the Preferences dialog. Default=1.0.

Ortho Snapping Mode 

When drawing a *Line spline (page 1–270)* or *Wall object (page 1–223)*, Ortho Snapping Mode constrains line creation to the horizontal or vertical directions relative to the active grid. That is, with Ortho on, you can draw only lines that are parallel to lines of the active grid. Ortho is particularly useful for drawing plans where all lines must be at 90 degrees to each other, such as a house plan. To constrain line-drawing to other angle increments, use *Polar Snapping Mode (page 2–39)* instead.

Important: By default, Ortho Snapping Mode is not directly accessible in the 3ds Max user interface. To add it as a button on the Snaps toolbar, follow *this procedure (page 2–39)*.

The following provisions apply to using Ortho:

- Use of Ortho is mutually exclusive of use of the *Polar Snapping Mode toggle (page 2–39)*; only one can be active at a time.
- Ortho Snapping applies primarily to the creation of line splines and wall objects. It can be used while creating other objects, but results are likely to be unsatisfactory.

Ortho mode displays a compass which gives a readout of the current angle of the input relative to the positive direction of the local X axis.

Note: If you hold down the Alt key in Ortho mode, the next point becomes doubly constrained by both the previous point and the first point of the current object. This allows you to close splines precisely. In this mode, two compasses are displayed; one each at the first *and* previous points.

Procedures

To add Ortho Snapping Mode to the Snaps toolbar:

1. Open the Snaps toolbar, if necessary. To do so, right-click an empty part of the main toolbar, such as the area directly below one of the drop-down lists, and choose Snaps.
2. Drag the right end of the toolbar to the right to make room for a new button.
3. Choose Customize menu > Customize User Interface.

   This opens the Customize User Interface dialog.
4. On the dialog, click the Toolbars tab.
5. Scroll down the Action list on the dialog to the Ortho Snapping Mode item. You can jump to the O section by clicking any item in the list and then pressing O on the keyboard.
6. Drag the Ortho Snapping Mode item from the list to the empty section of the Snaps toolbar. This adds the button to the toolbar.
7. Close the Customize User Interface dialog.

   3ds Max automatically saves the toolbar in its revised state and makes the new button a permanent part of the user interface.

To use Ortho:

1. Turn on the Ortho Snapping Mode button.

   The button appears depressed.
2. Begin to draw a Line spline.

   An orange compass appears where you place the first point, along with a red number indicating the angle of the current line segment with the positive direction of the local X axis.
3. Move the mouse cursor around in the viewport.

   The line jumps to 90-degree angle increments on the local X axis, while the compass reading updates to show the current angle.
4. Click to place the next vertex.
5. Repeat steps 3 and 4 until you're ready to complete the shape.
6. Do either of the following:
   - To finish the shape at the most recent vertex without closing it, right-click anywhere.
   - To finish the shape by closing it, position the mouse cursor close to the first point and then click. A small dialog opens asking if you want to close the spline; click Yes.

Polar Snapping Mode 

When drawing a *Line spline (page 1–270)* or *Wall object (page 1–223)*, Polar Snapping Mode constrains line creation to angle increments determined by the Angle Snap setting relative to the active grid. To change the Angle Snap setting, right-click the Angle Snap Toggle button on the main toolbar and in the Options panel > General group, edit the Angle setting.

Important: By default, Polar Snapping mode is not directly accessible in the 3ds Max user interface. To add it as a button on the Snaps toolbar, follow *this procedure (page 2–40)*.

Polar is particularly useful for drawing plans where angles between all lines must conform to specific angle increments, such as 45 degrees. If all lines must be at 90-degree angles to each other, use the *Ortho Snapping Mode toggle (page 2–38)* instead.

The following provisions apply to using Polar:

- Use of Polar is mutually exclusive of use of the Ortho toggle; only one can be active at a time.

- Polar applies primarily to the creation of line splines and wall objects. It can be used while creating other objects, but results are likely to be unsatisfactory.

Polar mode displays a compass that provides a readout of the current angle of the input relative to the positive direction of the local X axis.

Note: If you hold down Alt in Polar mode, the next point becomes doubly constrained by both the previous point and the first point of the current object. This allows you to close splines precisely. In this mode, two compasses are displayed; one each at the first *and* previous points.

Procedures

To add Polar Snapping Mode to the Snaps toolbar:

1. Open the Snaps toolbar, if necessary. To do so, right-click an empty part of the main toolbar, such as the area directly below one of the drop-down lists, and choose Snaps.
2. Drag the right end of the toolbar to the right to make room for a new button.
3. Choose Customize menu > Customize User Interface.

   This opens the Customize User Interface dialog.
4. On the dialog, click the Toolbars tab.
5. Scroll down the Actions list on the dialog to the Polar Snapping Mode item. You can jump to the P section by clicking any item in the list and then pressing P on the keyboard.
6. Drag the Polar Snapping Mode item from the list to the empty section of the Snaps toolbar. This adds the button to the toolbar.
7. Close the Customize User Interface dialog.

   3ds Max automatically saves the toolbar in its revised state and makes the new button a permanent part of the user interface.

To use Polar:

1. Turn on the Polar Snapping Mode button.

   The button appears depressed.
2. Begin to draw a line spline.

   An orange compass appears where you place the first point, along with a red number indicating the angle of the current line segment with the positive direction of the local X axis.
3. Move the mouse cursor around in the viewport.

   The line jumps to specific angle increments from the X-axis, while the compass reading updates to show the current angle. You set the angle increment in the Grid And Snap Settings dialog > *Options panel (page 2–46)* > General group, which you can access by right-clicking the Angle Snap Toggle button on the main toolbar.
4. Click to place the next vertex.
5. Repeat steps 3 and 4 until you're ready to complete the shape.
6. Do any of the following:
   - To finish the shape at the most recent vertex without closing it, right-click anywhere.
   - To finish the shape by closing it, position the mouse cursor close to the first point and then click. A small dialog opens asking if you want to close the spline; click Yes.
   - To finish the shape by closing it while constraining the line to the polar snap, first position the mouse cursor close to the first point, press and hold Alt to constrain the mouse by both the previous point and the first point, and then click. This vertex is automatically placed at the current angle increment from the first point, so that you need only click the first point to close the shape.

Grid and Snap Settings

Customize menu > Grid and Snap Settings

Right-click any snap button

This command displays the Grid and Snap Settings dialog. This modeless dialog establishes settings and options for snaps, the *home grid (page 3–952)*, and user-defined grids.

Controls on the Grid And Snap Settings dialog determine which snap settings are used when you activate snaps by clicking 3D Snap Toggle. Adjusting any of these snap settings does not automatically turn on snaps.

You can use objects within an XRef scene as snap references.

The Grid And Snap Settings dialog contains tabs for:

Snap Settings (page 2–41)

Snap Options (page 2–46)

Home Grid Settings (page 2–49)

User Grids Settings (page 2–51)

Procedure

To change grid and snap settings:

1. Choose Customize menu > Grid And Snap Settings and click the appropriate tab.
2. Choose the type of snap you want (Standard or NURBS).
3. Select the snap settings.

Snap Settings

Customize menu > Grid And Snap Settings > Grid And Snap Settings dialog > Snaps tab

Main toolbar > Right-click a snap button

Keyboard > Hold Shift +right-click > Snaps quadrant

Keyboard > S (toggles snaps on and off)

Snapping gives you additional control when creating, moving, rotating, and scaling objects by causing the cursor to "jump" to specific portions of existing geometry and other scene elements during creation and transformation of objects or sub-objects. The controls in this dialog set the snap strength and other characteristics such as the snap target.

Note: Snap functionality includes several features that enhance ease of use. See *the Snap Options topic (page 2–48)* for details.

You can specify the portion of the geometry where you will snap. For example, when Vertex is the active snap type, creating and transforming objects snaps to the vertices of existing geometry. You can specify any combination of active snap types to provide multiple snapping points. For example, if Vertex and Midpoint are active, snapping occurs at both vertices and edge midpoints.

The default snap type is Grid Points.

Note: Snapping is not on by default. You can toggle snapping by pressing the S key at any time, even in the middle of a transform. In this way you can combine snapping with free positioning.

Snaps work at sub-object levels. For example, you can use snaps to position a gizmo to the object on which you're working, or snap it to other objects in the scene.

You must activate a viewport in order to use snaps. Also, the Z-axis constraints don't apply to the

home grid or grid objects, since grids don't have a Z axis.

Settings are stored in the *3dsmax.ini (page 1–18)* file. The state of the snap settings persists from session to session.

Snaps and Axis Constraints

Snaps take precedence over axis constraints. If you highlight an axis constraint, such as Restrict to X, you can move the object only in X. But if you then turn on snaps, Restrict to X is suspended and not used.

The Snaps Toolbar



The most common Snaps settings are available from an optional toolbar. To toggle display of the Snaps toolbar, right-click an empty area of the main toolbar, such as the section under the Reference Coordinate System drop-down, and choose Snaps. The toolbar buttons are shown next to the relevant commands, below and in the *Snap Options topic (page 2–46)*.

The same settings are also available from the snap quad menu, available with Shift+right-click.

Snaps: Relative and Absolute

3ds Max offers two different types of snap behavior. You can use snaps to move a selection to a snap point, an absolute snap. You can also use snaps to move a selection a relative distance from a snap point. This is called a relative snap. For further information, see *To use relative snaps: (page 2–42)*.

Procedures

To set grid and snap settings:

1. Turn on the 3D Snap Toggle button to activate snaps.
2. Choose Customize > Grid And Snap Settings to display the Grid and Snap Settings dialog.
3. In the Snaps tab, select one or more of the types of snaps you want active.
4. Create an object or transform an object.

   Snap markers appear when the mouse cursor is over existing geometry or on a grid, depending on the active snap types. Each snap type has a different display; clicking when the snap-specific display is visible snaps to that spot.

To display the Snaps shortcut menu:

- Hold Shift and right-click anywhere in any viewport. The quad menu that opens gives you access to various snap settings including Snaps Use Axis Constraints and Snap To Frozen Objects.

To use absolute snaps:

- Click the object with Snaps on.

To use relative snaps:

1. Lock the selection set.
2. Click in the viewport.

   The distance from your cursor to the selection set is used as the relative snap distance. The object snaps that relative distance away from the snap point.

To use both constraint and snaps, do one of the following:

- In the Grid and Snap Settings dialog > Options tab > Translation group, turn on Use Axis Constraints.

- Hold Shift and right-click in the viewport, and then choose Options > Transform Constraints from the Snap quadrant.

Example: To use 3D snaps and rotation transformations together:

1. Create a box.
2. Select the box and choose Lock Selection.
3. Turn on 3D Snaps and click Rotate on the toolbar.
4.  From the Use Center flyout on the main toolbar, choose *Use Transform Coordinate Center (page 1–447)*.
5. Activate the Perspective viewport and move the cursor over the grid.

   A blue icon displays when the cursor passes over a grid point.
6. When the blue icon displays, click and drag to rotate the box around the selected grid point.

   You can rotate around anything you can snap to.

To turn snaps on and off during an operation:

- Use the S keyboard shortcut to turn snap on and off.

  Tip: You can select something with snap off, and then turn snap on to snap it to a snap target. Alternately you might want to snap to something, then position it freely wherever you want.

Interface

Use these check boxes on the Snaps tab to turn on any combination of snap settings.

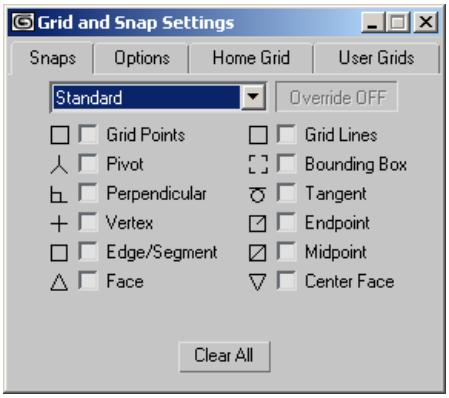
After setting snaps, close the dialog using the Close button in the dialog's upper-right corner. Do not click the Clear All button, or you'll turn off all the snaps.

Override—This label changes to display the temporary snap type used by the Override system. For more information, see *Snap Override (page 2–45)*.

Clear All—Turns off all of the Snaps check boxes.

Note: The layout of the Grid And Snap Settings dialog is generated at runtime. Because of this, it might appear slightly different than the illustrations shown here.

Standard snaps



These are the standard snap types used for grids, mesh, and shape objects. Non-grid snap types, when active, take priority over Grid Points and Grid Lines snaps: if the mouse is equally near a grid point and some other snap type, it will choose the other snap type.

Note: The button images shown below are from the *Snaps toolbar (page 2–42)*.

 **Grid Points**—Snaps to grid intersections. This snap type is on by default. Keyboard shortcut=Alt+F5.

Grid Lines—Snaps to any point on a grid line.

 **Pivot**—Snaps to pivot points of objects. Keyboard shortcut=Alt+F6.

Bounding Box—Snaps to one of the eight corners of an object's bounding box.

Perpendicular—Snaps to the perpendicular point on a spline, relative to the previous point.

Tangent—Snaps to a tangent point on a spline, relative to the previous point.


Vertex—Snaps to vertices of mesh objects or objects that can be converted to editable meshes. Snaps to segments on splines. Keyboard shortcut=Alt+F7.

Endpoint—Snaps to the end points of edges on meshes or spline vertices.

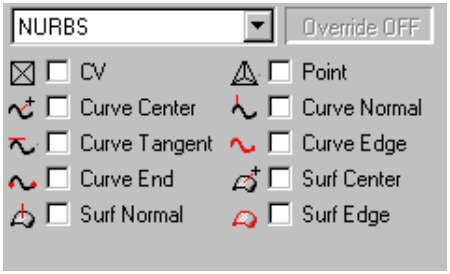

Edge/Segment—Snaps anywhere along edges (visible or invisible) or spline segments. Keyboard shortcut=Alt+F10.


Midpoint—Snaps to the middle of edges on meshes and spline segments. Keyboard shortcut=Alt+F9.


Face—Snaps anywhere on the surface of a face. Back faces are culled, so they have no effect. Keyboard shortcut=Alt+F11.

Center Face—Snaps to the center of triangular faces.

NURBS Snaps



These options snap to objects or sub-objects in a *NURBS model (page 3–980)*.

The NURBS snaps settings are aids for creating and transforming objects, and are not constraints. The software does not maintain the relationship between the NURBS object and other objects you create or transform.

CV—Snaps to a *CV sub-object (page 3–926)* in a *NURBS curve (page 3–980)* or *NURBS surface (page 3–981)*.

Point—Snaps to a *point sub-object (page 3–995)* in a NURBS model.

Curve Center—Snaps to the center of a NURBS curve.

The center of a NURBS curve is calculated parametrically, and might not be the same as the curve's apparent visual center.

Curve Normal—Snaps to a point normal to a NURBS curve.

This snap operates only while you are creating a new object that requires two or more clicks to create.

Curve Tangent—Snaps to a point tangent to a NURBS curve.

This snap operates only while you are creating a new object that requires two or more clicks to create.

Curve Edge—Snaps to the edge of a NURBS curve (the current object moves or is created to lie along the curve).

Curve End—Snaps to the end of a NURBS curve.

Surf Center—Snaps to the center of a NURBS surface.

The center of a NURBS curve is calculated parametrically, and might not be the same as the curve's apparent visual center.

Surf Normal—Snaps to a point on a NURBS surface normal to previous point.

This snap operates only while you are creating a new object.

Surf Edge—Snaps to the edge of a NURBS surface.

Snap Override

Any viewport > Hold Shift and right-click. > Standard > Choose snaps type.

Customize menu > Grid and Snap Settings > Snaps tab

Snap Override lets you supersede all the currently selected snap types and temporarily use only one, or none, of the snap types currently selected on the Grid and Snap Settings dialog. For example, you might be creating a spline while snapping to grid points, but then need to snap one of its vertices to the midpoint of an object.

Tip: You can choose the override by cycling through active snap types via a keyboard shortcut instead of using the quad menu. Press Alt+S repeatedly to choose the override snap type. If the Grid and Snap dialog is displayed, the "Override OFF" label changes to display the selected snap type.

Procedures

To use Snap Override (keyboard shortcut method):

1. Activate all the snap types you want to use or override with.
2. Make sure the Grid And Snap Settings dialog is open to the Snaps panel, so you can easily track the override type. Alternatively, if you're using only snap types available on the *Snaps toolbar (page 2–42)*, you can keep the toolbar open instead.
3. While creating or moving an object with Snap on, press Alt+S repeatedly.

   The Grid And Snap Settings dialog > "Override OFF" label changes to display the first available override snap type, chosen from the active snap types. If the Snaps toolbar is open, and the override snap type is available on the toolbar, the override snap type is visible there as well.
4. Press Alt+S repeatedly.

   The software cycles through all the active snap types for the override type.

   The benefit of this method is that it's fast and easy; you don't have to use the quad menu. But if you want to override with a snap type that isn't currently active, it's necessary to use the standard override method, which is described in the next procedure.

To use Snap Override (standard method):

1. While creating or moving an object with Snap on, hold down Shift and right-click in a viewport.
2. Select one of the snap types from the sub-menus in the Snap Override quadrant to make it the *only* active snap type.

   If the Grid And Snap Settings dialog is open to the Snaps panel, the "Override OFF" label changes to display the active snap override type. When you complete the mouse action, "Override OFF" is again displayed in the dialog, and the previously active snap types are active again.

To use Snap Override during a drag operation:

1. Left-click, press Shift, and then right-click to display the menu.
2. Release the left mouse button, and then left-click to select the snap you want.
3. Release the Shift key, right-click, and continue the drag operation (the geometry remains locked to the mouse).
4. Left-click to complete the operation.

Interface

In addition to the available snap types, the snap quad menu contains these items:

Snap Options quadrant—Lets you set the following options:

- **Grid and Snap Settings:** Toggles display of the Grid And Snap Settings dialog.
- **Snaps To Frozen Objects:** Turn this on to enable snapping to frozen objects. Default=off.
- **Snaps Use Axis Constraints:** Turn this on to use the current transform constraints. For example, if you're moving a vertex with Restrict To XY Plane on and want to snap the vertex to a point removed on the Z axis, turn this off, if necessary. Default=off.

Snap Override quadrant—Lets you set the following options:

NURBS/Standard—Let you choose a snap type for one-time use. This applies to the next snap only.

None—Turns off all snap types for the next mouse action. (This item is unavailable if the Snap Toggle is off.)

[last]—Displays the last snap override type you chose from the NURBS or Standard submenu, letting you easily reuse that type.

Snap Toggles quadrant—Lets you set six of the most common snap types on an ongoing basis. These remain in effect until you turn them off; it's the same as setting them on the Grid And Snap Settings dialog.

Snap Options

Customize menu > Grid and Snap Settings > Grid and Snap Settings dialog > Options panel

Main toolbar > Right-click 3D Snap Toggle > Grid and Snap Settings dialog > Options panel

The Options panel of the Grid And Snap Settings dialog lets you set options related to snapping.

Procedures

To set snap options:

1. Choose Customize menu > Grid And Snap Settings.
2. Click the Options tab.
3. Set the options as desired.

Example: To use Snap Preview:

The Snap Preview Radius setting lets you preview a snap before it actually happens. The Snap Radius setting determines when actual snapping occurs, providing the same functionality as Snap Strength in previous versions.

This procedure also demonstrates how to cycle through available snap/preview points with the keyboard shortcut.

1. Reset the software, and then maximize the Perspective viewport.

   To maximize the Perspective viewport, make sure it's active (yellow border), and then press Alt+W.

2. Zoom in a bit so the grid squares are relatively large.

   Snap and Snap Preview measure in pixels, so if the grid points are relatively far apart, it's easy to see the difference between Snap Preview and Snap.

3.  On the main toolbar, click the Snaps Toggle button to turn on snapping, and then right-click the button to open the Grid And Snap Settings dialog to the Snaps panel.

   The default Snaps setting is Grid Points only; if this is not the case, make it so. A grid point is the intersection of two grid lines.

4. On the dialog, click the Options tab.

   In the General group, Snap Preview Radius is set to 30 pixels and Snap Radius is set to 20 pixels. You can leave this dialog open while you use snaps.

5.  On the Create panel, click Shapes, and then click Object Type rollout > Line.

   You're now in line-creation mode.

6. Click anywhere in the viewport to start creating a line-type spline object.

7. Move the cursor around the viewport.

   As you do so, a rubber-band line connects the mouse cursor to the start point.

8. Position the mouse cursor near a grid point, but not too close, so that the snap cursor (cyan box + crosshairs) appears at the grid point, but the line endpoint remains attached to the mouse cursor.

   This shows the snap preview; you'll see its purpose in the next step.

9. Click the left mouse button.

   The second line vertex is created not where you clicked but at the grid point indicated by the snap preview. Snap preview lets you see where a snap will occur before it actually happens.

10. Move the cursor toward another grid point.

    At the Snap Preview Radius distance (30 pixels), the snap cursor appears at the grid point.

11. Continue moving the cursor toward the grid point.

    After you move it 10 pixels closer (the difference between Snap Preview Radius and Snap Radius), the line endpoint jumps from the mouse cursor to the grid point. This is the legacy snap functionality.

12. Click to accept the snap.

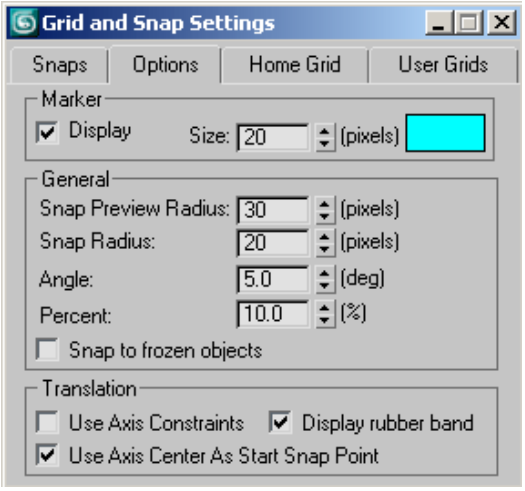
13. Complete the line, and then zoom out so the grid points are closer together. Start another line, and then position the mouse cursor in the center of a grid square, so the preview snap point appears at one of the grid points.

14. On the keyboard, press and hold Alt+Shift and then press S repeatedly to cycle through the remaining nearby grid points. If nothing happens, try zooming out more or increasing the Snap Preview Radius value.

    By default, the software uses the closest snap point as the preview location. This keyboard shortcut lets you choose other qualifying snap points.

Interface

Note: The layout of the Grid And Snap Settings dialog is generated at runtime. Because of this, it might appear slightly different than the illustration shown here.



Marker group

Provides settings that affect the visual display of the snap points.

Display—Toggles the display of the snap guides. When off, the snaps still function, but there's no display.

Size—Sets the size, in pixels, of the snap "hit" point. This is the small icon that indicates either the source or target snap point.

Color—Click the color swatch to display the Color Selector, where you can set the color of the snap display.

General group

Snap functionality includes several features that enhance ease of use. When using Snap, if the cursor comes within distance of a potential snap point that's less than the *Snap Preview Radius* distance but greater than the *Snap Radius* distance, the snap cursor jumps to that point as a preview of where the snap will occur, but no actual snapping occurs.

To use the preview point as the snap point, click or release the mouse button, depending on the context. If, instead, you then continue to move the cursor toward the potential snap point so that it comes within a distance equal to or less than the Snap Radius value, the snap takes place.

Tip: For best results, keep the Snap Preview Radius value 10 pixels or more higher than the Snap Radius value. This lets you preview any snap before it actually occurs.

Snap Preview Radius—When the cursor is a distance from a potential snap-to point between the Snap Preview Radius value and the Snap Radius value, the snap marker jumps to the closest potential snap-to point, but no snap occurs. Default=30.

When a snap-to preview point is highlighted, release or click the mouse button (depending on the context) to snap the current selection to that location. Alternatively, before using the displayed snap point, you can cycle though any other available preview and snap points by pressing and holding Alt+Shift and then press S repeatedly.

Note: The Snap Preview Radius value should generally be set higher than the Snap Radius value, so that previewing occurs before snapping. If you attempt to set the Snap Preview Radius value lower than the Snap Radius value, the software lowers the latter so that the two are equal. This effectively turns off previewing, so that only snapping is in effect.

Snap Radius—Sets the size of the area around the cursor, in pixels, within which snapping occurs automatically. Default=20.

Note: The Snap Radius value should generally be set lower than the Snap Preview Radius value, so that previewing occurs before snapping. If you attempt to set the Snap Radius value higher than the Snap Preview Radius value, the software raises the latter so that the two are equal. This effectively turns off previewing, so that only snapping is in effect.

Angle—Sets the increment at which objects are rotated about a given axis (degrees).

Percent—Sets the percentage increment for scale transforms.

 **Snap to Frozen Objects—**When on, snapping to frozen objects is enabled. Default=off. This option is also available from the Snaps shortcut menu, accessed when you hold Shift and right-click in any viewport, as well as from the *Snaps toolbar (page 2–42)*. Keyboard shortcut= Alt+F2 .

Translation group

 **Use Axis Constraints—**Constrains the selected object to move only along the axes specified on the *Axis Constraints toolbar (page 3–687)*. When turned off (the default), the constraints are ignored, and snapped objects can be translated in any dimension (assuming 3D snapping is used). This is also available from the Snaps shortcut menu, accessed when you hold Shift and right-click in any viewport, as well as from the *Snaps toolbar (page 2–42)*. Keyboard shortcut= Alt+F3 or Alt+D .

Display rubber band—When on and you move an selection, a rubber-band line appears between the original location and the mouse position. Use this visual aid for greater accuracy when fine-tuning Default=on.

Use Axis Center As Start Snap Point—Sets the center of the transform axis of the current selection set to be the initial snap point if there's no other start snap point detected by the snap system. This option works both at the object and sub-object levels.

This lets you, for example, snap the geometric center of a multi-object selection to a grid point.

Home Grid Settings

Customize menu > Grid And Snap Settings > Home Grid panel

Main toolbar > Right-click 3D snap toggle. > Grid And Snap Settings > Home Grid panel

The Home Grid panel of the Grid And Snap Settings dialog sets the spacing and other characteristics of the *home grid (page 3–952)*. Choosing useful home grid settings can simplify the construction process. The home grid provides a visual reference to use when creating objects in a scene.

In the software, grids have these primary uses:

- An aid in visualizing space, scale, and distance.
- Construction planes where you create and align objects in your scene.
- A reference system for using grid snaps.

Procedures

To set grid spacing for unit measure:

1. From the menu bar, choose Customize menu > Grid And Snap Settings.
2. Click the Home Grid tab.
3. Adjust the value for Grid Spacing, which is in current units.

   For example, if you have units set to centimeters, you might make one grid space equal to 1.000 (one unit, or one centimeter in this case).

To set major grid divisions for multiple units:

1. From the menu bar, choose Customize menu > Grid And Snap Settings.
2. Click the Home Grid tab.
3. Adjust the Major Lines Every Nth Grid Line value, which is the number of grid squares between major lines. The minimum is 2.

For example, if you use a grid spacing of one centimeter, you might use a value of 10 so the major grid divisions represent one decimeter.

In perspective viewports, you can set a fixed size for the displayed home grid. If Inhibit Perspective View Grid Resize is turned off however, the grid size adjusts as you zoom in or out.

To set view update options:

1. From the menu bar, choose Customize menu > Grid And Snap Settings.
2. Click the Home Grid tab.
3. Under Dynamic Update, choose either Active Viewport (the default) or All Viewports.

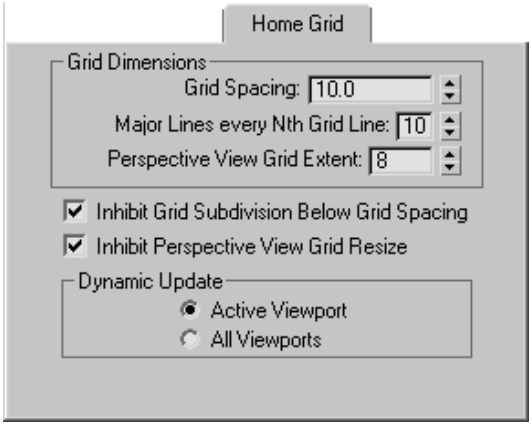
To allow subdivision below grid spacing:

1. From the menu bar, choose Customize menu > Grid And Snap Settings.
2. Click the Home Grid tab.
3. Turn off Inhibit Grid Subdivision Below Grid Spacing.

   When you turn off this box, you can zoom indefinitely "deep" into any plane of the home grid. Each grid square subdivides into the same number of smaller grid spaces.

Interface

Note: The layout of the Grid And Snap Settings dialog is generated at runtime. Because of this, it might appear slightly different than the illustration shown here.



Grid Dimensions group

Grid spacing—Grid spacing is the size of the grid's smallest square. Use this spinner to adjust the spacing (which is in current units), or enter the value directly.

For example, if you have units set to centimeters, you might make one grid space equal to 1.000 (one unit, or one centimeter in this case).

Major Lines every Nth Grid Line—The home grid displays heavier or "major" lines to mark groups of grid squares. Use spinner to adjust the value, which is the number of grid squares between major lines, or you can enter the value directly, the minimum is 2.

For example, if you use a grid spacing of one centimeter, you might use a value of 10 so the major grid divisions represent one decimeter.

Perspective View Grid Extent—Sets the size of the home grid in the Perspective viewport.

This value is specified in terms of the Grid Spacing value, and represents the length of half the grid along an axis. This means that if Grid Spacing=10.0 and Perspective View Grid Extent=7, you will have a grid that is 140 x 140 units in size.

Inhibit Grid Subdivision Below Grid Spacing—Causes 3ds Max to treat the grid as a fixed set of lines when you zoom in on the home grid. In effect, the grid stops at the grid space setting. If you keep zooming, the fixed grid is lost from view. Zooming out is not affected. When you zoom out, the home grid expands indefinitely to maintain the major grid divisions. Default=on.

When this is turned off, you can zoom indefinitely into any plane of the home grid. Each grid square subdivides into the same number of smaller grid spaces.

For a grid spacing of one centimeter and a major division of 10, the next level down subdivides into millimeter spaces, and so on.

Inhibit Perspective View Grid Resize— Causes 3ds Max to treat the grid in the Perspective viewport as a fixed set of lines when you zoom in *or* out. In effect, the grid maintains one size, no matter how much you zoom. Default=on.

When this is turned off, the grid in the Perspective viewport will subdivide to adjust its size when you zoom in or out.

Dynamic Update group

Dynamic Update—By default, only the active viewport updates as you change values for Grid Spacing and Major Lines every Nth. The other viewports update after you have completed changing the values. Choose All Viewports to have all viewports update as you change the values.

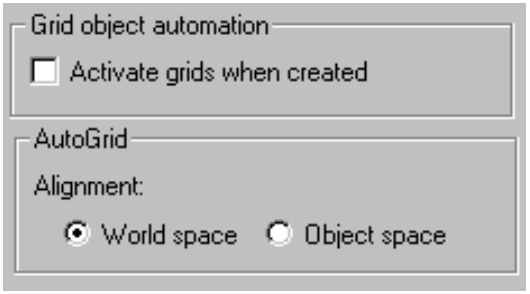
User Grids Settings

Customize menu > Grid And Snap Settings > Grid And Snap Settings dialog > User Grids tab

The User Grids panel controls automatic activation of *grid objects (page 2–20)* and settings for *AutoGrid (page 2–7)*.

Interface

Note: The layout of the Grid And Snap Settings dialog is generated at runtime. Because of this, it might appear slightly different than the illustration shown here.



Grid object automation group

Determines if the software automatically makes grids active upon creation.

Activate grids when creating—Turn on to automatically activate the grid you created.

Tip: If not turned on, you can activate the grid by selecting the grid, then right-clicking and choosing Activate Grid.

AutoGrid group

AutoGrid allows you to automatically create grids on the surface of objects. The task of creating objects that are aligned to other object surfaces has been simplified. Turn on AutoGrid in the Object Type rollout of a command panel. When you turn on AutoGrid, the software uses the setting in the User Grids dialog for world or object space.

World space—Aligns grids to world space.

Object space—Aligns grids to object space.

Drawing Aid Utilities

The utilities listed here are additional aids to drawing and precision. The Measure utility returns the measurements of a selected shape or 3D object. The Rescale World Units utility rescales either a selection or an entire scene.

Use the *Utilities panel (page 3–778)* to open a drawing aid utility and make it active.

Measure Utility (page 2–52)

Rescale World Units Utility (page 2–53)

Measure Utility

Utilities panel > Utilities rollout > Measure button

The Measure utility provides measurements of a selected object or shape.

See also

Measure Distance Tool (page 2–15)

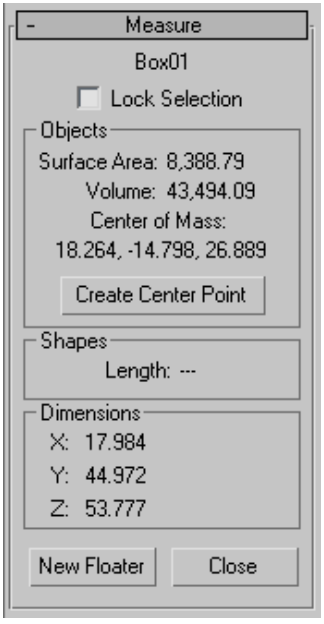
Procedure

To measure an object:

1. On the Utilities panel, click the Measure button.
2. Select the object you want to measure.

   If both an object and a shape are in the selection, information is displayed for both types. If several objects are in a selection, the sum of their measurements is displayed.

Interface



Text display—Displays the name of the object in the current selection. If more than one object is in the selection, Multiple Objects Selected is displayed.

Lock Selected—Prevents the displayed data from changing when you change selection. For example, you might need to select and manipulate another object that affects the currently selected object.

Objects group

Displays information about renderable mesh objects.

Surface Area—Displays the total surface area of all objects in the selection, in units squared.

Volume—Displays the total volume of all objects in the selection, in units cubed. Note that objects with "holes" caused by missing faces can result in inaccurate volume values. When an object has one or more holes, an asterisk appears beside the Volume number.

Center Of Mass—Displays the world coordinates of the location of the center of mass for the object or the center of mass of the selected objects.

Create Center Point—Creates a point helper object at the center of mass.

Shapes group

Displays information about shape objects.

Length—Displays the sum of the length of all splines in all selected shapes.

Dimensions group

Displays the dimensions of the object, as they appear in world space. For example, if it were a box with the created dimensions of 15 x 10 x 25, and that box were scaled 200 percent, then this group would report dimensions of 30 x 20 x 50.

Button set

New Floater—Launches a modeless Measure dialog that displays the same information found under the Objects, Shapes, and Dimensions groups on the Utilities panel. In addition, you can expand the dialog horizontally, in case the values are too long to be viewed in the default dialog size.

Tip: While the Measure floater is displayed, you can view the length of a spline while you're creating it.

Close—Closes the utility.

Rescale World Units Utility

Utilities panel > Utilities rollout > More button > Utilities dialog > Rescale World Units

This utility rescales the world units of either the entire scene or selected objects in the scene. The Rescale World Units dialog controls scale factor and whether it is applied to the entire scene or the current selection.

Procedure

To rescale an object:

1. Select an object in the viewport.
2. On the Utilities panel, click the More button, and on the Utilities dialog, choose Rescale World Units from the list.
3. On the Rescale World Units rollout, click the Rescale button.
4. Set the Scale Factor to the desired scale. For example, setting the scale to a value of 2.0 doubles the size of the object that it's applied to.
5. In the Affect group, choose the Selection option.
6. Click OK to apply the scale.

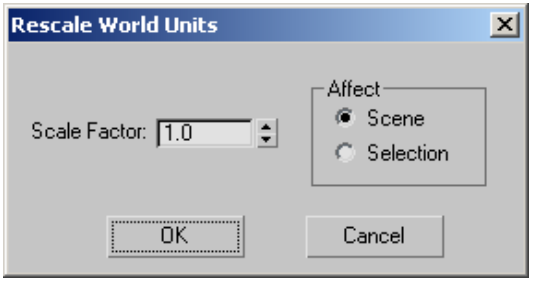
Interface

Rescale World Units rollout



Rescale—After selecting the objects you want to rescale, click this to display the Rescale World Units dialog.

Rescale World Units dialog



Scale Factor—Specifies the scaling factor.

Affect—

- **Scene—**Applies the scale to the entire scene.
- **Selection—**Applies the scale only to the current selection.

OK—Applies the scale.

Cancel—Cancels the operation.

11 Space Warps and Particle Systems

Space warps and particle systems are additional modeling tools. Space warps are "force fields" that deform other objects, creating the effect of ripples, waves, blowing wind, and so on. Particle systems generate particle sub-objects for the purpose of simulating snow, rain, dust, and so on. (You use particle systems primarily in animations.)

See also

Space Warp Objects (page 2–55)

Particle Flow (page 2–109)

Non-Event-Driven Particle Systems (page 2–237)

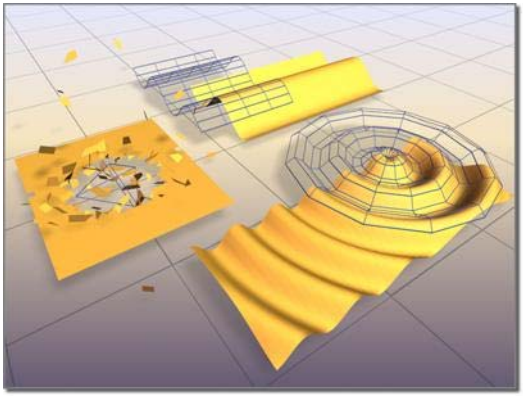
Space Warps

Create panel > Space Warps

Space warps are nonrenderable objects that affect the appearance of other objects. Space warps create force fields that deform other objects, creating the effect of ripples, waves, blowing wind, and so on.

Space warps behave somewhat like modifiers, except that space warps influence world space, rather than object space, as geometric modifiers do.

When you create a space warp object, viewports show a wireframe representation of it. You can transform the space warp as you do other 3ds Max objects. The position, rotation, and scale of the space warp affect its operation.



Surfaces deformed by space warps

Left: Bomb

Right: Ripple

Rear: Wave

A space warp affects objects only when the objects are *bound to it (page 2–58)*. The warp binding appears at the top of the object's modifier stack. A space warp is always applied after any transforms or modifiers.

When you bind multiple objects to a space warp, the space warp's parameters affect all the objects equally. However, each object's distance from the space warp or spatial orientation to the warp can change the warp's effect. Because of this spatial effect, simply moving an object through warped space can change the warp's effect.

You can also use multiple space warps on one or more objects. Multiple space warps appear in an object's stack in the order you apply them.

Note: You can take advantage of the AutoGrid feature to orient and position new space warps with respect to existing objects. See *AutoGrid (page 2–7)* for details.

Space Warps and Supported Objects

Some types of space warps are designed to work on deformable objects, such as geometric primitives, meshes, patches, and splines. Other types of warps work on particle systems such as Spray and Snow.

Five space warps (Gravity, PBomb, Wind, Motor, and Push) can work on particle systems and also serve special purposes in a dynamics simulation. In the latter case, you do not bind the warps to objects, but rather assign them as effects in the simulation.

On the Create panel, each space warp has a rollout labeled Supports Objects Of Type. This rollout lists the kind of objects you can bind to the warp.

Basics of Using Space Warps

Follow these general steps to use space warps:

1. Create the space warp.
2. Bind objects to the space warp.

    Click Bind To Space Warp on the main toolbar (available from the Select And Link flyout), and then drag between the space warp and the object.

   The space warp has no visible effect on your scene until you bind an object, system, or selection set to it.
3. Adjust the space warp's parameters.
4. Transform the space warp with Move, Rotate, or Scale. The transforms often directly affect the bound object.

You can animate space warp parameters and transforms. You can also animate space warp effects by animating transforms of an object bound to the warp.

Particle Leakage and Deflector Space Warps

A deflector is a space warp that acts as a barrier to particles in particle systems. Occasionally stray particles can leak through a deflector under the following circumstances:

- When a particle happens to hit the deflector too near the end or beginning of a time interval, and numerical error in the solution doesn't report a hit
- When a particle hits too near the edge of a face referenced by a UDeflector, and neither face finds it
- When a particle is moving quickly and first appears too close to the deflector, so the very first update cycle within the particle system takes it past the deflector without the deflector ever seeing it

Often this isn't a problem because the particles bounce off solid objects, so you don't see the errant particles. When it does cause problems, you can use a planar deflector instead of a UDeflector, a collection of planars to approximate the mesh, or a simple mesh to replace the planar. The solutions vary, so where one has a glitch the other might work just fine.

If particles are moving fast and the deflector is in a particular position (for example, it might be too close to the emitter) many particles can leak through defectors. Sometimes you can fix this leakage by changing the particle system's Subframe Sampling setting or the particle speed. Other times you must reposition the deflector farther away from the emitter.

In addition, particles imbued with *bubble motion (page 2–270)* can leak through deflectors, particularly when set to high amplitudes. To avoid this, use alternative methods of implementing bubble-like motion, such as varying speed (see *Particle Generation Rollout (page 2–260)*), setting a higher angle for the stream spread with spray systems (see *Spray Particle System (page 2–244)*, *Super Spray Particle System (page 2–249)*), or using larger particles with an animated texture map.

Space Warp Categories

There are four categories of space warps, available via the list on the Create panel's Space Warps category.

Forces

These space warps are used to affect particle systems and dynamics systems. All of them can be used with particles, and some can be used with dynamics. The Supports Objects of Type rollout indications which systems each space warp supports.



Motor Space Warp (page 2–61)

Push Space Warp (page 2–59)

Vortex Space Warp (page 2–63)

Drag Space Warp (page 2–66)

Path Follow Space Warp (page 2–71)

PBomb Space Warp (page 2–68)

Displace Space Warp (page 2–76)

Gravity Space Warp (page 2–73)

Wind Space Warp (page 2–75)

Deflectors

These space warps are used to deflect particles or to affect dynamics systems. All of them can be used with particles and with dynamics. The Supports Objects of Type rollout indicates which systems each space warp supports.

PDynaFlect Space Warp (page 2–81)

POmniFlect Space Warp (page 2–78)

SDynaFlect Space Warp (page 2–85)

SOmniFlect Space Warp (page 2–84)

UDynaFlect Space Warp (page 2–86)

UOmniFlect Space Warp (page 2–85)

SDeflector Space Warp (page 2–87)

UDeflector Space Warp (page 2–89)

Deflector Space Warp (page 2–90)

Geometric/Deformable

These space warps are used to deform geometry.

FFD(Box) Space Warp (page 2–91)

FFD(Cyl) Space Warp (page 2–95)

Wave Space Warp (page 2–100)

Ripple Space Warp (page 2–102)

Displace Space Warp (page 2–76)

Conform Space Warp (page 2–103)

Bomb Space Warp (page 2–105)

Modifier-Based

These are space-warp versions of object modifiers (see *Modify Panel (page 3–758)*). Read more about these in *Modifier-Based Space Warps (page 2–107)*.

Bend Modifier (page 1–560)

Noise Modifier (page 1–743)

Skew Modifier (page 1–790)

Taper Modifier (page 1–863)

Twist Modifier (page 1–876)

Stretch Modifier (page 1–836)

Procedure

To create a space warp:

1.  On the Create panel, click Space Warps.

   The Space Warps panel is displayed.
2. Choose a category of space warp from the list.
3. On the Object Type rollout, click a space warp button.

   Tip: You can take advantage of the AutoGrid feature to orient and position new space warps with respect to existing objects. For details, see *AutoGrid (page 2–7)*.
4. Drag in a viewport to create the space warp.

   See the topics for the various space warps for further details.

Bind to Space Warp 

Main toolbar > Bind to Space Warp

Use the Bind to Space Warp button to attach the current selection to a *space warp (page 3–1014)* or vice versa.

Procedure

To bind the current selection to a space warp:

1. Select an object
2.  Click Bind to Space Warp.
3. Drag a line from the selected object to the space warp object. You can also press H to select the space warp by name.

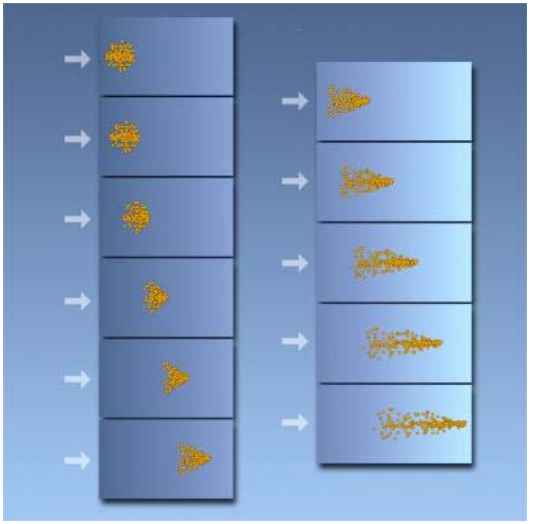
   The space warp object flashes for a moment to show that the bind was successful.

Forces

Push Space Warp

Create panel > Space Warps > Forces > Object Type rollout > Push

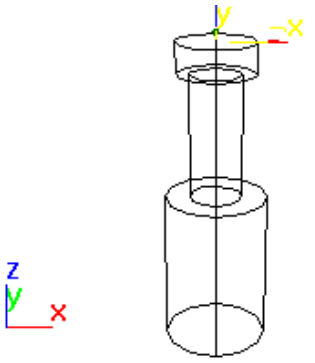
Create menu > Space Warps > Forces > Push



Push disperses a cloud of particles

The Push space warp applies a force to either *particle systems (page 2–108)* or *dynamics systems (page 2–686)*. The effect is slightly different, depending on the system.

- **Particles:** Applies a uniform, unidirectional force in a positive or negative direction. A positive force moves in the direction of the pad on the hydraulic jack. The breadth of the force is infinite, perpendicular to the direction; you can limit it using the Range option.
- **Dynamics:** Provides a point force (also called a point load) away from the pad of the hydraulic jack icon. A negative force pulls in the opposite direction. In dynamics, applying a force is the same as pushing something with your finger.



Push viewport icon

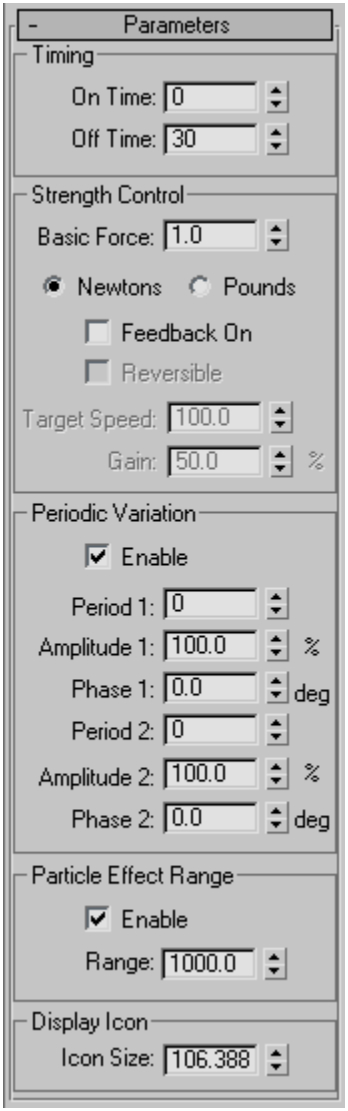
Procedure

To create a Push space warp:

1.  On the Create panel, click Space Warps. Choose Forces from the list, and then on the Object Type rollout, click Push.
2. Drag in a viewport to define the size.

   The Push warp appears as a hydraulic jack icon.

Interface



Timing group

On Time/Off Time—The numbers of the frames in which the space warp begins and ends its effect. Because Push moves the particles to which it's applied over time, no keyframes are created.

Strength Control group

Basic Force—The amount of force exerted by the space warp.

Newtons/Pounds—This option specifies the units of force used by the Basic Force spinner.

A pound is about 4.5 Newtons, and one newton is one kilogram-per-second-squared. When Push is applied to particle systems, these values have only subjective meaning because they depend on the built-in weighting factors and time scaling used by the particle system. However, when used in a dynamics system, the value listed is precisely the value used.

Feedback On—When on, the force varies depending on the speed of the affected particles relative to the specified Target Speed. When off, the force remains constant, regardless of the speed of the affected particles.

Reversible—When on, if the particle's speed exceeds the Target Speed setting, the force is reversed. Available only if you turn on Feedback On.

Target Speed—Specifies the maximum speed in units per frame before the Feedback takes effect. Available only if you turn on Feedback On.

Gain—Specifies how quickly the force adjusts to approaching the target speed. If set to 100 percent, the correction is immediate. If set lower, a slower and "looser" response occurs. Available only if you turn on Feedback On.

Note: Setting Gain above 100 percent can result in over-correction, but is sometimes necessary to overcome damping from other system settings, such as IK damping.

Periodic Variation group

These settings introduce variations into the force by affecting the Basic Force value randomly. You can set two waveforms to produce a noise effect.

Enable—Turns on the variations.

Period 1—The time over which the noise variation makes a full cycle. For example, a setting of 20 means one cycle per 20 frames.

Amplitude 1—The strength of the variation (in percent). This option uses the same types of units as the Basic Force spinner.

Phase 1—Offsets the variation pattern.

Period 2—Provides an additional variation pattern (a second wave) to increase the noise.

Amplitude 2—The strength of the variation of the second wave (in percent). This option uses the same types of units as the Basic Force spinner.

Phase 2—Offsets the variation pattern of the second wave.

Particle Effect Range group

Lets you restrict the Push effect's range to a specific volume. This affects particle systems only; it has no effect on dynamics.

Enable—When on, limits the range of the effect to a sphere, displayed as a tri-hooped sphere. The effect falls off increasingly as the particles near the boundary of the sphere.

Range—Specifies the radius of the range of the effect, in units.

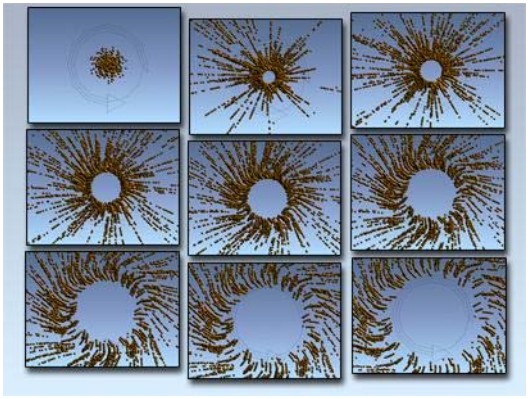
Display Icon group

Icon Size—Sets the size of the Push icon. This is for display purposes only, and does not alter the Push effect.

Motor Space Warp

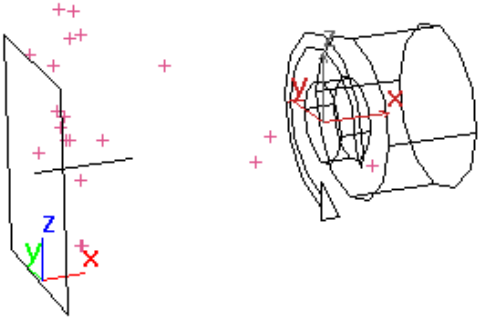
Create panel > Space Warps > Forces > Object Type rollout > Motor

Create menu > Space Warps > Forces > Motor



Motor disperses a cloud of particles

The Motor space warp works like *Push (page 2–59)*, but applies rotational torque to the affected particles or objects rather than a directional force. Both the position and orientation of the Motor icon affect particles, which swirl around the Motor icon. When used in dynamics, the position of the icon relative to the affected object has no effect, but the orientation of the icon does.



Motor viewport icon (with particle system on the left)

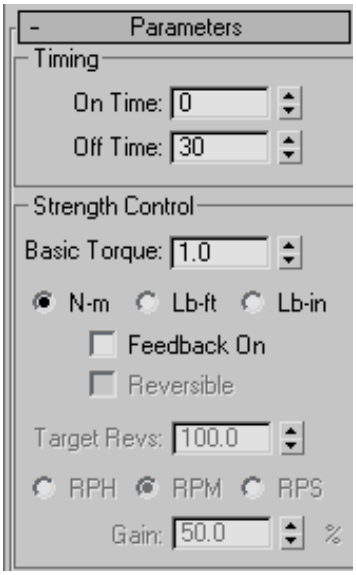
Procedure

To create a motor space warp:

1.  On the Create panel, click Space Warps > Forces > Motor. Choose Forces from the list, then on the Object Type rollout, click Motor.
2. Click and drag in a viewport to define the size.

   The Motor warp appears as a box-shaped icon with an arrow indicating the direction of the torque.

Interface



Timing group

On Time/Off Time—The numbers of the frames in which the space warp begins and ends its effect. Because Motor moves the objects to which it's applied over time, no keyframes are created.

Strength Control group

Basic Torque—The amount of force exerted by the space warp.

N-m/Lb-ft/Lb-in—Specify the unit of measure for the Basic Torque setting, using common world measurements of torque. N-m stands for Newton meters, Lb-ft stands for pound-force feet, and Lb-in stands for pound-force inches.

Feedback On—When on, the force varies depending on the speed of the affected objects relative to the specified Target Speed. When off, the force remains constant, regardless of the speed of the affected objects.

Reversible—When on, if the object's speed exceeds the Target Speed setting, the force is reversed. Available only if you turn on Feedback On.

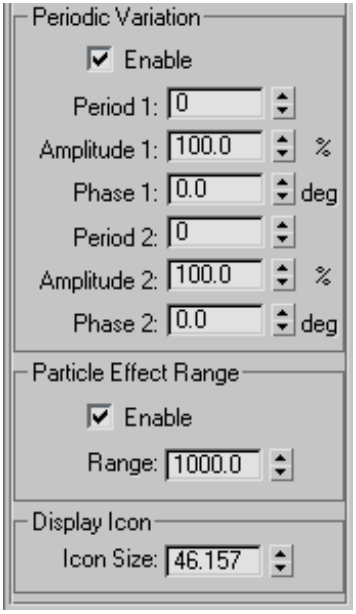
Target Revs—Specifies the maximum revolutions before the feedback takes effect. Speed is specified in units traveled per frame. Available only if you turn on Feedback On.

RPH/RPM/RPS—Specifies the units of measure for Target Revs in revolutions per hour, minute, or second. Available only if you turn on Feedback On.

Gain—Specifies how quickly the force adjusts to approaching the target speed. If set to 100%, the correction is immediate. If set lower, a slower and "looser" response occurs. Available only if you turn on Feedback On.

Note: Setting Gain above 100% can result in over-correction, but is sometimes necessary to overcome damping from other system settings, such as IK damping.

Periodic Variation group



These settings introduce variations into the force by affecting the Basic Torque value randomly. You can set two waveforms to produce a noise effect.

Enable—Turn on to enable the variations.

Period 1—The time over which the noise variation makes a full cycle. For example, a setting of 20 means one cycle per 20 frames.

Amplitude 1—The strength of the variation (in percent). This option uses the same types of units as the Basic Torque spinner.

Phase 1—Offsets the variation pattern.

Period 2—The next two spinners provide an additional variation pattern to increase the noise.

Amplitude 2—The strength of the variation of the second wave in (percent). This option uses the same types of units as the Basic Torque spinner.

Phase 2—Offsets the variation pattern of the second wave.

Particle Effect Range group

Lets you restrict the Motor effect's range to a specific spherical volume. This affects particles systems only; it has no effect on dynamics.

Enable—When on, limits the range of the effect to a sphere, displayed as a tri-hooped sphere. The effect falls off increasingly as the particles near the boundary of the sphere.

Range—Specifies the radius of the range of the effect, in units.

Display Icon group

Icon Size—Sets the size of the Motor icon. This is for display purposes only, and does not alter the Motor effect.

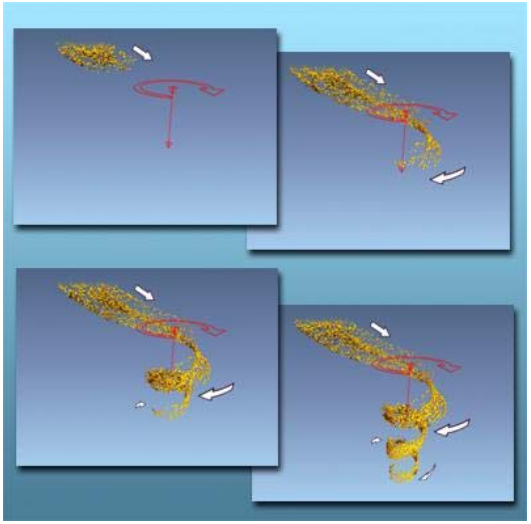
Vortex Space Warp

Create panel > Space Warps > Forces > Object Type rollout > Vortex

Create menu > Space Warps > Forces > Vortex

The Vortex space warp applies a force to *particle systems (page 2–108)*, spinning them through a whirling vortex, and then moving them down a long, thin spout or *vortex well*. Vortex is useful for creating black holes, whirlpools, tornadoes, and other funnel-like objects.

The space warp settings let you control the vortex shape, the well characteristics, and rate and range of particle capture. The shape of the vortex is also affected by particle system settings, such as speed.



Particle stream caught in a vortex

Procedure

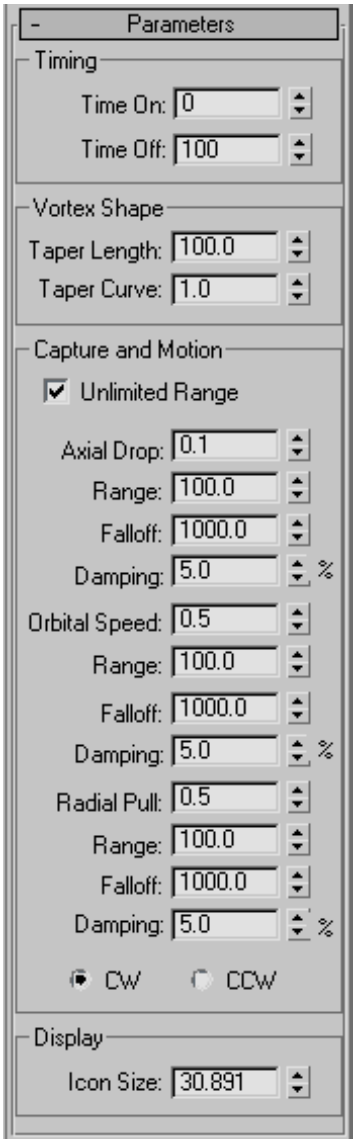
To create a Vortex space warp:

1.  On the Create panel, click Space Warps. Choose Forces from the list, and then click Vortex.
2. Determine which world axis you want the vortex to spiral around, and then drag in the appropriate viewport to create the space warp. For example, if you want the vortex to spin around the vertical world axis, create the space warp in the Top viewport. You can rotate the warp later to change the vortex direction, and animate the warp orientation.

   The Vortex warp appears as a curved-arrow icon in the plane you drag in, with a second, perpendicular arrow indicating the axis of rotation as well as the direction of the well. This second axis is called the *drop axis*.

   Note: The position of the space warp plays an important role in the results. The vertical position affects the shape of the vortex, and the horizontal position determines its location. If you want the particles to spiral around the particle emitter, place both at the same location.

Interface



Timing group

Time On/Time Off—The frame numbers at which the space warp becomes active and becomes inactive.

Vortex Shape group

Taper Length—Controls the length of the vortex, as well as its shape. Lower settings give you a "tighter" vortex, while higher settings give you a "looser" vortex. Default=100.0.

Taper Curve—Controls the shape of the vortex. Low values create a vortex with a wide, flared mouth, while high values create a vortex with nearly vertical sides. Default=1.0. Range=1.0 to 4.0.

Capture and Motion group

This group contains basic settings for Axial Drop, Orbital Speed, and Radial Pull, with Range, Falloff, and Damping modifiers for each.

Unlimited Range—When on, Vortex exerts full damping strength over an unlimited range. When off, the Range and Falloff settings take effect.

Axial Drop—Specifies how quickly particles move in the direction of the drop axis.

Range—The distance from the center of the Vortex icon, in system units, at which Axial Damping has its full effect. Takes effect only when Unlimited Range is turned off.

Falloff—Specifies the distance beyond the Axial Range within which Axial Damping is applied. Axial Damping is strongest at the Range distance, decreases linearly out to the limit of the Axial Falloff, and has no effect beyond that. Takes effect only when Unlimited Range is turned off.

Damping—Controls the degree to which particle motion parallel to the drop axis is restrained per frame. Default=5.0. Range=0 to 100.

For subtle effects, use values of less than 10%. For more overt effects, try using higher values that increase to 100% over the course of a few frames.

Orbital Speed—Specifies how quickly the particles rotate.

Range—The distance from the center of the Vortex icon, in system units, at which Orbital Damping has its full effect. Takes effect only when Unlimited Range is turned off.

Falloff—Specifies the distance beyond the Orbital Range within which Orbital Damping is applied. Orbital Damping is strongest at the Range distance, decreases linearly out to the limit of the Orbital Falloff, and has no effect beyond that. Takes effect only when Unlimited Range is turned off.

Damping—Controls the degree to which orbital particle motion is restrained per frame. Smaller values produce a wide spiral, while larger values produce a thin spiral. Default=5.0. Range=0 to 100.

Radial Pull—Specifies the distance from the drop axis at which the particles rotate.

Range—The distance from the center of the Vortex icon, in system units, at which Radial Damping has its full effect. Takes effect only when Unlimited Range is turned off.

Falloff—Specifies the distance beyond the Radial Range within which Radial Damping is applied. Radial Damping is strongest at the Range distance, decreases linearly out to the limit of the Radial Falloff, and has no effect beyond that. Takes effect only when Unlimited Range is turned off.

Damping—Controls the degree to which Radial Pull is restrained per frame. Default=5.0. Range=0 to 100.

CW/CCW—Determines whether particles rotate clockwise or counterclockwise.

Display group

Icon Size—Specifies the size of the icon.

Drag Space Warp

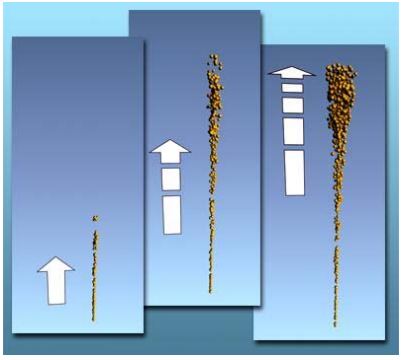
Create panel > Space Warps > Forces > Object Type rollout > Drag

Create menu > Space Warps > Forces > Drag

The Drag space warp is a particle motion damper that reduces particle velocity by a specified amount within a specified range. The damping can be applied linearly, spherically, or cylindrically. Drag is useful for simulating wind resistance, transfers into dense media (like water), impacts with force fields, and other, similar situations.

With each damping type, you can control the damping effect along several vectors. The damping is also affected by particle system settings, such as speed.

Note: To create a more uniform drag effect, default values for all directional parameters are now 5.0%.



Drag slows down a stream of particles.

Procedure

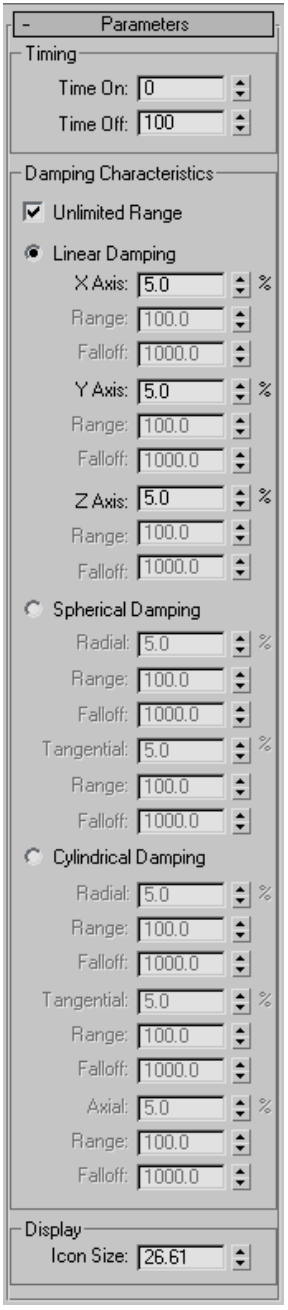
To create a Drag space warp:

1.  On the Create panel, click Space Warps. Choose Forces from the list, and then click Drag.
2. Drag in a viewport to create the space warp.

   The space warp appears initially as a box within a box, indicating that it's using the default Linear Damping mode.
3. To apply the damping spherically or cylindrically, choose Spherical Damping or Cylindrical Damping in the command panel.

   Note: The position and orientation of the space warp plays an important role for all three damping types.
4. Change the settings for the current damping type as necessary.

Interface



Timing group

Time On/Time Off—The frame numbers at which the space warp becomes active and becomes inactive.

Damping Characteristics group

This group lets you choose Linear Damping, Spherical Damping, or Cylindrical Damping, plus a set of parameters for each.

Unlimited Range—When on, Drag exerts full damping strength over an unlimited range. When off, the Range and Falloff settings for the current damping type take effect.

- **Linear Damping—**Motion for each particle is separated into vectors for the space warp's local X, Y, and Z axes. The area over which damping is exerted for each vector is an infinite plane whose thickness is determined by the corresponding Range setting.

  X Axis/Y Axis/Z Axis—Specifies the percentage of particle motion along the local Drag space warp axis that's affected by the damping.

  Range—Sets the thickness of the "range plane," or the infinite plane perpendicular to the specified axis. Takes effect only when Unlimited Range is turned off.

  Falloff—Specifies the distance beyond the X, Y, or Z Range within which Linear Damping is applied. Damping is strongest at the Range distance, decreases linearly out to the limit of the Falloff, and has no effect beyond that. While Falloff takes effect only beyond the Range, it is measured from the center of the icon, and always has a minimum value equal to the Range value. Takes effect only when Unlimited Range is turned off.

- **Spherical Damping—**When Drag operates in Spherical Damping mode, its icon is a sphere within a sphere. Particle motion is broken up into radial and tangential vectors. Damping

is applied for each vector within a spherical volume whose radius is set by the Range setting, when Unlimited Range is off.

Radial/Tangential—Radial specifies the percentage of particle motion toward or away from the center of the Drag icon that's affected by the damping. Tangential specifies the percentage of particle motion across the body of the Drag icon that's affected by the damping.

Range—Specifies the distance from the center of the Drag icon, in system units, within which damping is in full effect. Takes effect only when Unlimited Range is turned off.

Falloff—Specifies the distance beyond the Radial/Tangential Range within which Linear Damping is applied. Damping is strongest at the Range distance, decreases linearly out to the limit of the Falloff, and has no effect beyond that. While Falloff takes effect only beyond the Range, it is measured from the center of the icon, and always has a minimum value equal to the Range value. Takes effect only when Unlimited Range is turned off.

- **Cylindrical Damping—**When Drag operates in Spherical Cylindrical mode, its icon is a cylinder within a cylinder. Particle motion is broken up into radial, tangential, and axial vectors. Damping is applied within a spherical volume for the radial and tangential vectors and on a planar basis for the axial vector.

  Radial/Tangential/Axial—Damping controls the percentage of particle motion toward or away from the center of the circular portion of the icon (Radial), across the radial vector (Tangential), or along the length of the icon's long axis (Axial) that's affected by the damping, on a per-frame basis.

  Range—Specifies the distance from the center of the Drag icon, in system units, within which Radial and Axial damping are in full effect. Range also specifies the thickness of the infinite plane that governs the range of Axial damping. Takes effect only when Unlimited Range is turned off.

  Falloff—Specifies the distance beyond the Radial/Tangential/Axial Range within which Linear Damping is applied. Damping is strongest at the Range distance, decreases linearly out to the limit of the Falloff, and has no effect beyond that. While Falloff takes effect only beyond the Range, it is measured from the center of the icon, and always has a minimum value equal to the Range value. Takes effect only when Unlimited Range is turned off.

Display group

Icon Size— Specifies the size of the icon.

PBomb Space Warp

Create panel > Space Warps > Forces > Object Type rollout > PBomb

Create menu > Space Warps > Forces > PBomb

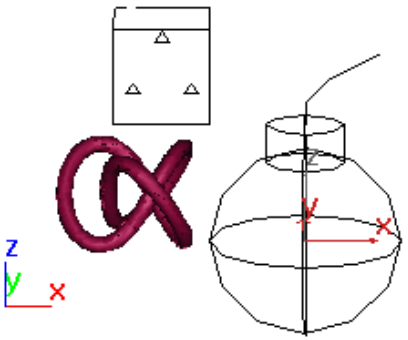
The PBomb space warp creates an impulse wave to blow up a particle system, as distinguished from the *Bomb space warp (page 2–105)*, which blows up geometry. The PBomb is particularly well suited to the Particle Array (PArray) system with Particle Types set to Object Fragments. This space warp also applies an impulse as a dynamics effect.

The general usage is as follows:

- Create a particle system (recommended: a *PArray system (page 2–256)* set to Object Fragments).
-  Create a PBomb and use *Bind to Space Warp (page 2–58)* to bind it to a non-event-driven particle system. Be sure to bind the PBomb to the particle system and not to the distribution object. Or, if using Particle

Flow, use a *Force operator (page 2–204)* to apply the space warp to the particle system.

- Adjust the parameters of both the PBomb and the particle system.



Right: PBomb viewport icon

Above: PArray particle system

Below: Torus knot used as the PArray's distribution object



Effect of blowing up the torus knot

Procedure

Example: To use PBomb with PArray:

You can use the Particle Bomb bound to a Particle Array to blow an object into fragments. The following steps demonstrate the basic setup.

Begin by binding a particle array to an object, and setting parameters.

1. Create the object you want to blow up.
2. Create a PArray and use Pick Object to assign the object to blow up as the PArray object-based emitter.
3. In the Basic Parameters rollout > Viewport Display group, choose Mesh to display the fragments as mesh objects in the viewports.
4. On the Particle Generation rollout of PArray, set Speed and Divergence to 0.0. This prevents PArray from moving the particles, letting PBomb do the work.
5. On the Particle Generation rollout, set Life to the length of the active time segment, so that the fragments appear during the entire animation.
6. In the Particle Type rollout > Particle Types group, choose Object Fragments. In the Object Fragment Controls group, choose Number of Chunks, and set the Minimum to about 50, depending on your distribution object and the effect you want.

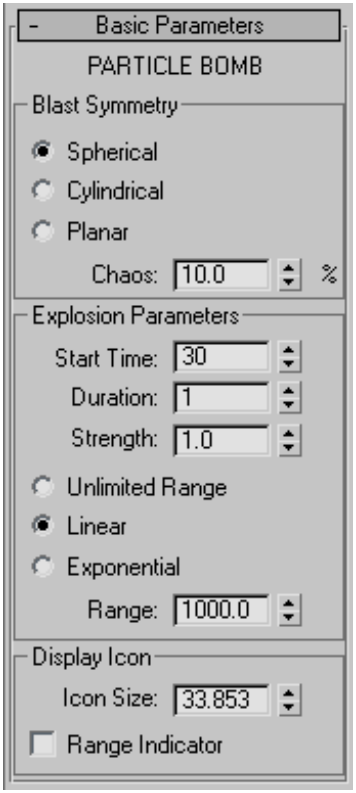
Now create a particle bomb, and bind it to the particle array.

1. In the Space Warps panel > Particles & Dynamics category, click the PBomb button and drag in a viewport to create the PBomb icon.
2. Use Bind to Space Warp to bind the PArray icon to the PBomb icon. (Don't bind the distribution object by mistake.)
3. Select the PBomb icon and move to the Modify panel.
4. In the Explosion Parameters group, set Blast Symmetry to Spherical, Start Time to 10, Duration to 1, and Strength to 1.0. Choose Linear, if it's not already chosen.
5. Drag the time slider between frames 9 and 20 to see the effect.
6. Go to frame 12 and try out various settings. Notice that increasing Strength expands the explosion effect at the current frame. If you

decrease Range enough, the bomb no longer affects all or part of the object (depending on the placement of the PBomb icon). Test the three Blast Symmetry settings by placing the bomb in the center of the object and then seeing the different blast patterns.

Once you get an explosive effect you like, you can return to the PArray settings, add spin or thickness to the fragments, and so on.

Interface



Blast Symmetry group

These options specify the shape, or pattern of the blast effect.

Spherical—The blast force radiates outward from the PBomb icon in all directions. The icon looks like a spherical anarchist's bomb.

Cylindrical—The blast force radiates outward from and normal to the central axis, or core of the cylindrical icon. The icon looks like a stick of dynamite with a fuse.

Planar—The blast force radiates both up and down, perpendicular to the plane of the planar icon. The icon looks like a plane with arrows pointing up and down along the direction of the blast force.

Chaos—The blast forces vary for each particle or each frame, an effect similar to Brownian motion, with a rate of change in the direction of force equal to the rendering interval rate. Note: This setting is effective only when the Duration spinner is set to 0.

Explosion Parameters group

Start Time—The frame number at which the impulse forces are first applied to the particles.

Duration—The number of frames, beyond the first, over which the forces are applied. This value should typically be a small number, such as between 0 and 3.

Strength—The change in velocity along the blast vector, in units per frame. Increasing Strength increases the speed with which the particles are blown away from the bomb icon.

Unlimited Range—The effects of the bomb icon reach all bound particles throughout the scene. This option ignores the Range setting (which specifies the distance of the PBomb effect).

Linear—The impulse forces decay linearly between the full Strength setting to a value of 0 at the specified Range setting.

Exponential—The impulse forces decay exponentially between the full Strength setting to a value of 0 at the specified Range setting.

Range—The maximum distance, in units, over which the PBomb icon affects the bound particle system. If the Range is large enough to reach only

a portion of the particle system, only that part of the system is affected.

If you turn on Range Indicator (see following), the extent of the range is indicated by a tri-hooped sphere. If you choose Unlimited Range, this parameter has no effect.

Display Icon group

These options affect the visual display of the PBomb icon.

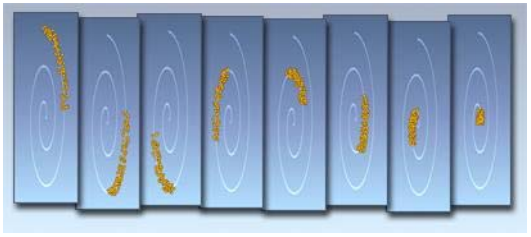
Icon Size—Alters the overall size of the PBomb icon.

Range Indicator—Displays a wireframe sphere that indicates the volume of the particle bomb's influence. If you choose Unlimited Range, turning this on has no effect.

Path Follow Space Warp

Create panel > Space Warps > Forces > Object Type rollout > Path Follow

Create menu > Space Warps > Forces > Path Follow



Particles following a spiral path

The Path Follow space warp forces particles to follow a spline path.

Procedures

To create a Path Follow space warp:

1.  On the Create panel, click Space Warps. Choose Forces from the list, then on the Object Type rollout, click Path Follow.
2. Click and drag in a viewport to define the size.

   The Path Follow warp appears as a box-shaped icon containing curved lines that indicate hypothetical paths.

To set up and use Path Follow:

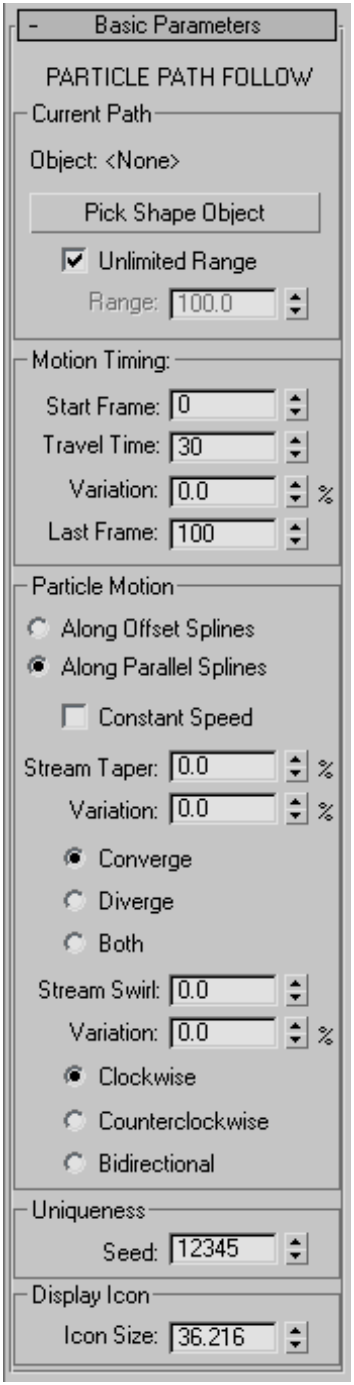
1. Create a shape consisting of a single spline. (You can use shapes with multiple splines, but the software uses only one spline for the path.)
2. Create a particle system and set its parameters to emit particles over a range of frames.
3. On the Create panel, click Path Follow, and drag in a viewport to create the Path Follow icon (a cube with wavy lines). The display and position of this icon don't affect the particles.
4. On the Basic Parameters rollout, click Pick Shape Object and select the spline you created earlier.
5.  Bind the Path Follow icon to the particle system.
6. Adjust the Path Follow parameters to create the particle movement you want.

   Note: You can further animate the particle effect by animating the spline vertices.

   Tip: To adjust the position of the path or particle system while maintaining access to the Path Follow parameters, turn on Modifier Stack rollout > Pin Stack while the Path Follow space warp is selected.

If the particles don't follow the emitter after it's moved, then change any PathFollow parameter, and the motion will be corrected.

Interface



Current Path group

Lets you choose the path for the particles, and specify the range of influence of the Space Warp.

Object—Displays the name of the currently assigned path.

Pick Shape Object—Click this, and then click a shape in the scene to select it as a path. You can use any shape object as a path; if you select a multiple-spline shape, only one the lowest-number spline is used. You can also use NURBS curves as paths.

Unlimited Range—When off, the range of influence of the space warp is limited to the value set in the Distance spinner. When on, the space warp influences all bound particles in the scene, regardless of their distance from the path object.

Range—Specifies the range of influence when Unlimited Range is off. This is the distance between the path object and the particle system. The position of the Path Follow space warp's icon is ignored.

Motion Timing group

These controls affect how long particles are influenced by Path Follow.

Start Frame—The frame at which Path Follow begins to influence the particles.

Travel Time—The time each particle takes to traverse the path.

Variation—The amount by which each particle's travel time can vary.

Last Frame—The frame at which Path Follow releases the particles and no longer influences them.

Particle Motion group

The controls in this area determine the motion of particles.

Along Offset Splines—The distance between the particle system and the path alter the effect of the particle motion. If the first vertex of the spline is at the birthplace of the particle, the particle follows the spline path. If you move the path away from the particle system, the particles are affected by the offset.

Along Parallel Splines—Particles follow a copy of the selected path, parallel to the particle system. In this mode, the position of the path relative to the particle system does not matter. The orientation of the path, however, affects the particle stream.

Constant Speed—When on, all particles travel at the same speed.

Stream Taper—Causes particles to converge or diverge toward the path over time, or to simultaneously converge and diverge. You specify the effect by choosing Converge, Diverge, or Both (see following). This provides a tapering effect over the length of the path.

Variation—The amount by which Stream Taper can vary for each particle.

Converge—When Stream Taper is greater than 0, the particles move in toward the path as they follow the path. The effect is that the stream tapers from larger to smaller over time.

Diverge—Provides the opposite effect of Converge. The particles diverge from the path over time.

Both—Splits the particle stream, causing some particles to converge and others to diverge.

Stream Swirl—Specifies the number of turns by which particles spiral about the path. In conjunction with Stream Taper, alters the diameter of the spiral.

Stream Swirl is generally more effective when you choose Along Offset Splines.

Variation—The amount by which each particle can vary from the Spiral value.

Clockwise—Particles spiral in a clockwise direction.

Counterclockwise—Particles spiral in a counterclockwise direction.

Bidirectional—The stream is split so that particles spiral in both directions.

Uniqueness group

Provides a seed number for unique generation of the particle pattern.

Seed—Specifies the seed number for the current Path Follow.

Display Icon group

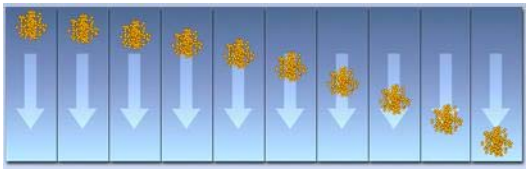
Affects the display of the Path Follow icon.

Icon Size—Specifies the size of the Path Follow icon. Does not alter the Path Follow effect.

Gravity Space Warp

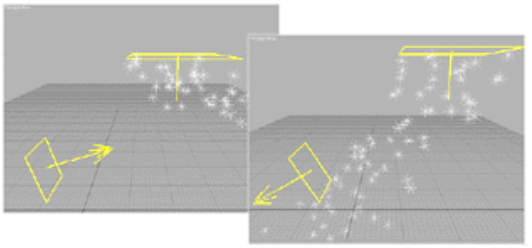
Create panel > Space Warps > Forces > Object Type rollout > Gravity

Create menu > Space Warps > Forces > Gravity



Particles falling because of gravity

The Gravity space warp simulates the effect of natural gravity on particles generated by a particle system. Gravity is directional. Particles moving in the direction of the gravity arrow accelerate. Particles moving against the arrow decelerate. In the case of spherical gravity, motion is toward the icon. Gravity can also be used as an effect in dynamics simulations. See *Dynamics Utility (page 2–686)*.



Gravity effect on snow

Procedure

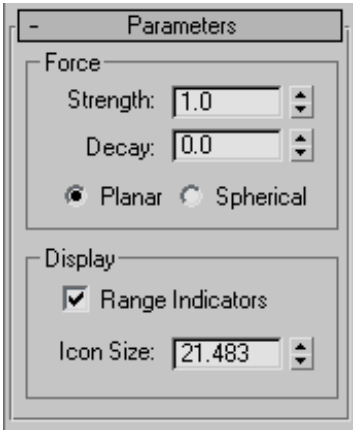
To create gravity:

1.  On the Create panel, click Space Warps. Choose Forces from the list, then on the Object Type rollout, click Gravity.
2. Drag in a viewport.

   The Gravity icon appears. For planar gravity (the default), the icon is a wireframe square with a direction arrow on one side. For spherical gravity, the icon is a wireframe sphere.

   The initial direction of planar gravity is along the negative Z axis of the construction grid that is active in the viewport where you drag. You can rotate the gravity object to change the direction.

Interface



Force group

Strength—Increasing Strength increases the effect of gravity; that is, how objects move in relation to the Gravity icon's direction arrow. Strength less than 0.0 creates negative gravity, which repels particles moving in the same direction and attracts particles moving in the opposite direction. When Strength is set to 0.0, the Gravity space warp has no effect.

Decay—When Decay is set to 0.0, the Gravity space warp has the same strength throughout world space. Increasing the Decay value causes gravity strength to diminish as distance increases from the position of the gravity warp object. Default=0.0.

Planar—Gravity effect is perpendicular to the plane of the Gravity warp object throughout the scene.

Spherical—Gravity effect is spherical, centered on the Gravity warp object. This choice is effective for creating water fountain or planetary effects.

Display group

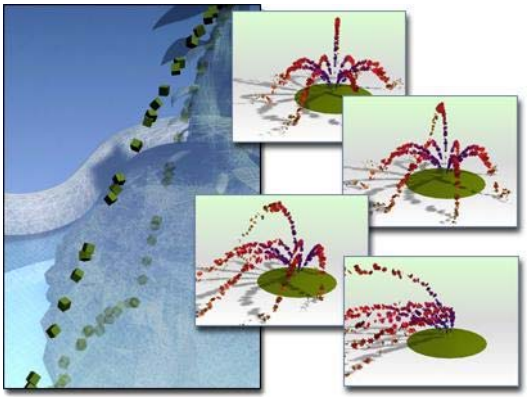
Range Indicators—When on, and when the Decay value is greater than 0.0, icons in the viewports indicate the range at which the force of gravity is half the maximum value. For the Planar option, the indicators are two planes; for use the Spherical option, the indicator is a double-hooped sphere.

Icon Size—Size of the Gravity warp object icon, in active units. You set the initial size when you drag to create the Gravity object. This value does not change the gravity effect.

Wind Space Warp

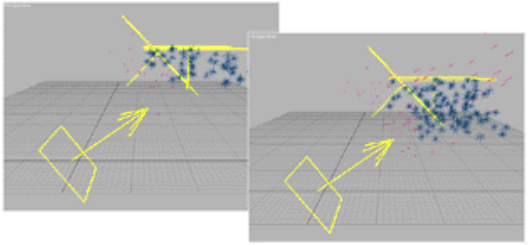
Create panel > Space Warps > Forces > Object Type rollout > Wind

Create menu > Space Warps > Forces > Wind



Wind changing the direction of the spray of a fountain

The Wind space warp simulates the effect of wind blowing particles generated by a particle system. Wind is directional. Particles moving in the direction of the wind arrow accelerate. Particles moving against the arrow decelerate. In the case of spherical wind, motion is toward or away from the icon. Wind is similar in effect to the Gravity space warp, but has added parameters for turbulence and other features characteristic of wind in the natural world. Wind can also be used as an effect in dynamics simulations. See *Dynamics Utility (page 2–686)*.



Wind effect on snow and spray

Procedure

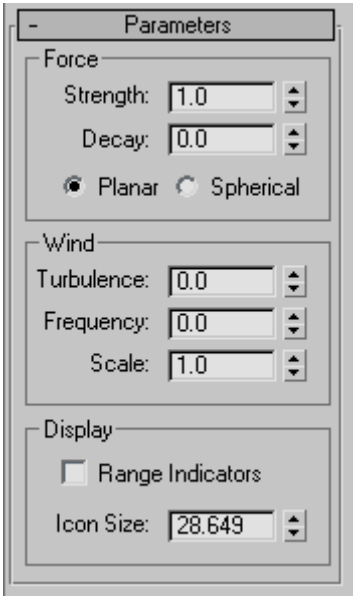
To create wind:

1.  On the Create panel, click Space Warps. Choose Forces from the list, then on the Object Type rollout, click Wind.
2. Drag in a viewport.

   The wind icon appears. For planar wind (the default), the icon is a wireframe square with a direction arrow coming out of one side. For spherical wind, the icon is a wireframe sphere.

   The initial direction of planar wind is along the negative Z axis of the construction grid that is active in the viewport where you drag. You can rotate the wind object to change the direction.

Interface



Force group

These settings are comparable to the Gravity parameters.

Strength—Increasing Strength increases the wind effect. Strength less than 0.0 creates a suction. It

repels particles moving in the same direction and attracts particles moving in the opposite direction. When Strength is 0.0, the Wind warp has no effect.

Decay—When Decay is set to 0.0, the Wind warp has the same strength throughout world space. Increasing the Decay value causes wind strength to diminish as distance increases from the position of the Wind warp object. Default=0.0.

Planar—Wind effect is perpendicular to the plane of the Wind warp object, throughout the scene.

Spherical—Wind effect is spherical, centered on the Wind warp object.

Wind group

These settings are specific to the Wind space warp.

Turbulence—Causes particles to change course randomly as the wind blows them. The greater the value, the greater the turbulence effect.

Frequency—When set greater than 0.0, causes turbulence to vary periodically over time. This subtle effect is probably not visible unless your bound particle system generates a large number of particles.

Scale—Scales the turbulence effect. When Scale is small, turbulence is smoother and more regular. As Scale increases, turbulence grows more irregular and wild.

Display group

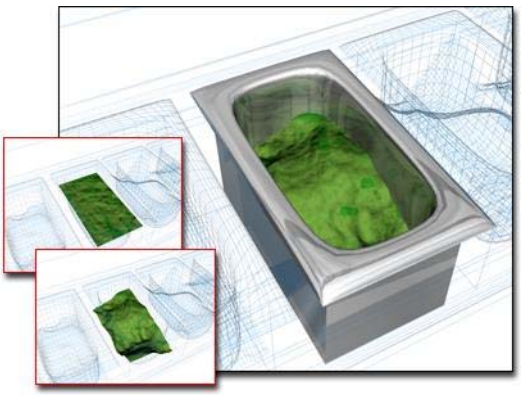
Range Indicators—When the Decay value is greater than zero, icons appear in the viewports that represent the range at which the force of wind is half the maximum value. When you use the Planar option, the indicators are two planes; when you use the Spherical option, the indicator is a double-hooped sphere.

Icon Size—Size of the Wind warp object icon, in active units. You set the initial Icon Size value when you drag to create the wind object. This value does not change the wind effect.

Displace Space Warp

Create panel > Space Warps > Geometric/Deformable > Object Type rollout > Displace

Create menu > Space Warps > Geometric/Deformable > Displace



Displace used to change the surface in the container

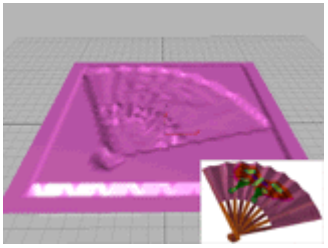
The Displace space warp acts as a force field to push and reshape an object's geometry. Displace affects both geometry (deformable objects) and particle systems. There are two basic ways to use the Displace space warp:

- Apply the gray scale of a bitmap to generate the displacement amount. Black areas of the 2D image are not displaced. Whiter areas push outward, causing a 3D displacement of geometry.
- Apply displacement directly by setting displacement Strength and Decay values.

The Displace space warp works similarly to the Displace modifier, except that, like all space warps, it affects world space rather than object space. Use the Displace modifier when you need to create detailed displacement of a small number of objects. Use the Displace space warp to displace

particle systems, a large number of geometric objects at once, or an object relative to its position in world space.

For geometry, the detail of the displacement depends on the number of vertices. Use the Tessellate modifier to tessellate faces you want to show in greater detail.



Displace space warp on a patch and the bitmap it uses

Procedures

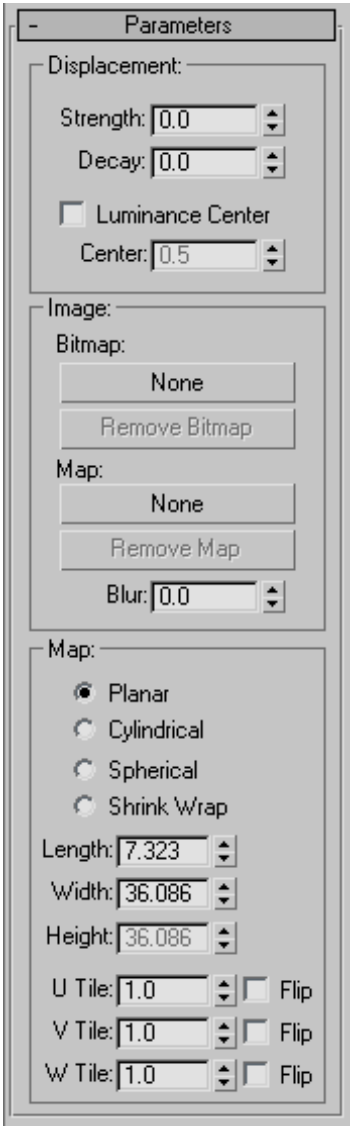
To create a Displace space warp:

1.  On the Create panel, click Space Warps. Choose Geometric/Deformable from the list, and then on the Object Type rollout, click Displace.
2. Drag in a viewport to create the Displace warp object, which appears as a wireframe. Its shape depends on the active mapping parameter settings. Regardless of the mapping, a single drag creates the space warp.
3. Bind the space warp to an appropriate object.

To assign a bitmap to a displace space warp:

1. Select the Displace warp object.
2. In the Parameters rollout > Displacement group, click the Bitmap button (labeled "None" by default). Use the selection dialog to choose a bitmap.
3. Set the Strength value. Vary the strength of the field to see how the bitmap displaces the object's geometry.

Interface



Displacement group

These are the basic controls for Displace space warps.

Strength—When set to 0.0, the Displace warp has no effect. Values greater than 0.0 displace object geometry or particles away from the position of

the Displace space warp object. Values less than 0.0 displace geometry toward the warp. Default=0.0

Decay—By default, the Displace warp has the same strength throughout world space. Increasing Decay causes displacement strength to diminish as distance increases from the position of the Displace warp object. Default=0.0

Luminance Center—By default, the Displace space warp centers the luminance by using medium (50%) gray as the zero displacement value. Gray values greater than 128 displace in the outward direction (away from the Displace warp object) and gray values less than 128 displace in the inward direction (toward the Displace warp object). You can adjust the default using the Center spinner. With a Planar projection, the displaced geometry is repositioned above or below the Planar gizmo. Default=0.5. Range=0 to 1.0.

Image group

These options let you choose a bitmap and map to use for displacement.

Bitmap—(Labeled "None" by default.) Click to assign a bitmap or map from a selection dialog. After you choose a bitmap or map, this button displays the bitmap's name.

Remove Bitmap—Click to remove the bitmap or map assignment.

Blur—Increase this value to blur or soften the effect of the bitmapped displacement.

Map group

This area contains mapping parameters for a bitmapped Displace warp. The mapping options are comparable to those options used with mapped materials. The four mapping modes control how the Displace warp object projects its displacement. The warp object's orientation controls where in the scene the displacement effect will appear on bound objects.

Planar—Projects the map from a single plane.

Cylindrical—Projects the map as if it were wrapped around the cylinder.

Spherical—Projects the map from a sphere, with singularities at the top and bottom of the sphere, where the bitmap edges meet at the sphere's poles.

Shrink Wrap—Truncates the corners of the map and joins them all at a single pole, creating one singularity.

Length, Width, Height—Specify the dimensions of the bounding box of the space warp gizmo. Height has no effect on planar mapping.

U/V/W Tile—The number of times the bitmap repeats along the specified dimension. The default value of 1.0 maps the bitmap once; a value of 2.0 maps the bitmap twice, and so on. Fractional values map a fractional portion of the bitmap in addition to copies of the whole map. For example, a value of 2.5 maps the bitmap two and one-half times.

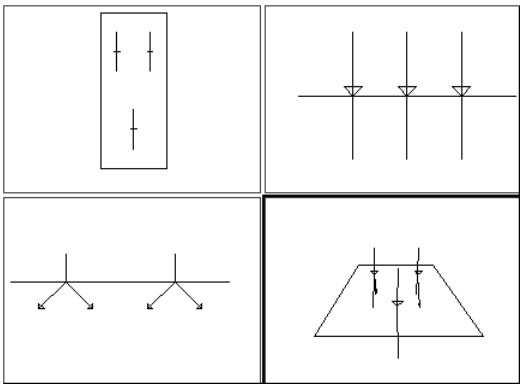
Flip—Reverses the orientation of the map along the corresponding U, V, or W axis.

Deflectors

POmniFlect Space Warp

Create panel > Space Warps > Deflectors > Object Type rollout > POmniFlect

POmniFlect is a planar version of the omniflector type of space warp. It provides enhanced functionality over that found in the original Deflector space warp, including refraction and spawning capabilities.



POmniFlect viewport icon

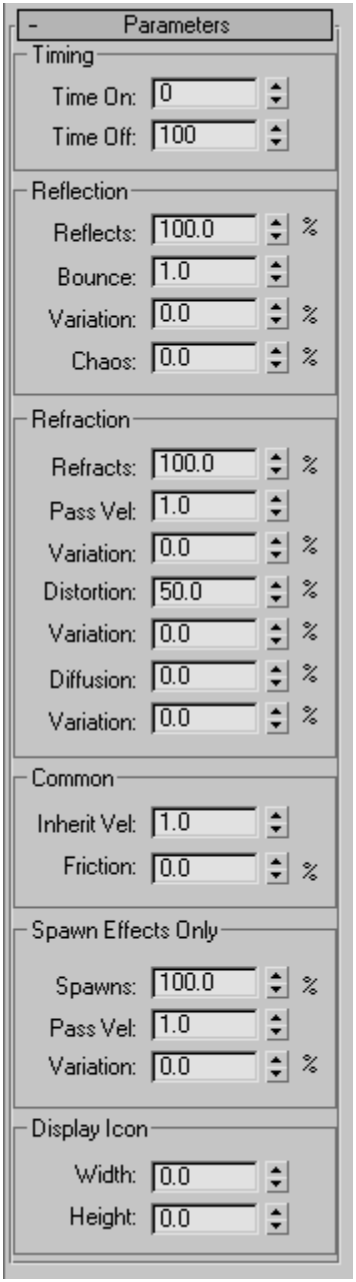
Procedure

To create a POmniFlect space warp:

1.  On the Create panel, click Space Warps. Choose Deflectors from the list, then on the Object Type rollout, click POmniFlect.
2. Drag in a viewport to create the planar icon.

   Note: Because particles bounce off the icon, the size of the icon affects particle deflection.
3. Apply the deflector to the particle system using the appropriate method:
   - If using *Particle Flow (page 2–109)*, specify the deflector in the *Collision test (page 2–212)* or *Collision Spawn test (page 2–215)* parameters.
   -  If using a *non-event-driven particle system (page 2–237)*, *bind (page 2–58)* the particle system to the deflector icon.
4. Position the POmniFlect icon to interrupt the particle stream.
5. Adjust the POmniFlect parameters as necessary.

Interface



Timing group

The two spinners specify the start frame and end frame of the deflection effect.

Time On/Off—Time On specifies the frame at which the deflection begins, and Time Off specifies the frame at which the deflection ends.

Reflection group

These options affect the *reflection* of particles from the space warp. The POmniFlect can reflect or refract particles, or perform a combination of the two.

Reflects—Specifies the percentage of particles to be reflected by the POmniFlect. See also Refracts, later in this topic, for methods of combining the two effects.

Bounce—This is a multiplier that specifies how much of the initial speed of the particle is maintained after collision with the POmniFlect. Using the default setting of 1.0 causes the particle to rebound with the same speed as it collides. A real-world effect would usually be less than 1.0. For a "flubber" effect, set greater than 1.0.

Variation—Specifies the variation of Bounce applied to the range of particles. For example, a Variation of 50% applied to a Bounce setting of 1.0 would result in randomly applied Bounce values ranging from 0.5 to 1.5.

Chaos—Applies a random variation to the bounce angle. When set to 0.0 (no chaos), all particles bounce off the POmniFlect surface perfectly (like banking pool balls). A non-zero setting causes the deflected particles to scatter.

Refraction group

These settings are similar to those in the Reflection group, but these affect the *refraction* of particles as they pass through the POmniFlect, causing the direction of the particles to change.

Refracts—Specifies the percentage of particles *not already reflected* that will be refracted by the POmniFlect.

Note: The Refracts value affects only those particles not already reflected because the reflected particles are processed before the refracted particles. Thus, if you set Reflects to 50% and Refracts to 50%, you would not get a 50/50 split of particles. Rather, half the particles would be reflected, and then *half the remainder* (25% of the total) would be refracted. The remaining particles either pass through without being refracted or are passed on to Spawn Effects.

To get a 50/50 split of reflection and refraction, set Reflects to 50% and Refracts to 100%.

Pass Vel—Specifies how much of a particle's initial speed is maintained after passing through the POmniFlect. The default setting of 1 retains the initial speed is retained, so there's no change. A setting of 0.5 reduces the speed by half.

Variation—Specifies the variation of Pass Velocity applied to the range of particles.

Distortion—Controls the angle of refraction. A value of 0 means there's no refraction. A value of 100% sets the angle of the particles to be parallel with the POmniFlect surface. A value of −100% sets the angle perpendicular to the surface. The Distortion effect is reversed when particles strike the POmniFlect from the back side.

Note: Distortion and Refraction do not work properly when particles strike the POmniFlect surface at exactly 90 degrees. In this case, any positive Distortion settings cause a scattering of particles, while negative values have no effect.

Variation—Specifies a range of variation of the Distortion effect.

Diffusion—Applies a diffusion effect to the refraction by randomly modifying each particle's

Distortion angle by the Diffusion angle. This effectively scatters the particles into a hollow cone.

Variation—Specifies a range of variation of the Diffusion value.

Common group

Friction—The amount by which particles are slowed as they move along the deflector surface. A value of 0% means they're not slowed at all. A value of 50% means they're slowed to half their original speed. A value of 100% means they stop moving when they strike the surface. Default=0%. Range=0% to 100%.

Tip: To have particles slide along a deflector surface, set Bounce to 0. Also, unless influenced by a force such as Wind or Gravity, particles meant to slide should strike the surface at an angle other than 90 degrees.

Inherit Vel—(Velocity Inheritance) Determines how much of a moving POmniFlect's speed is applied to reflected or refracted particles.

For example, if Inherit Vel is 1.0, particles with no motion that are hit by a moving PomniFlect inherit the speed of the POmniFlect at the point of collision.

Spawn Effects Only group

These settings affect only particles set to Spawn On Collision that do not either reflect or refract from the omniflector. See *Particle Spawn Rollout (page 2–271)*. The Spawns percentage spinner works like the Reflects and Refracts percentage spinners, but is the third in line to be processed. Thus, if either Reflects or Refracts are set to 100%, no particles are affected by these settings. Note also that reflective or refractive particles spawn upon collision, regardless of the settings in this group.

Spawns—Specifies the percentage of particles that can use spawn effects.

Pass Vel—Specifies how much of the particle's initial speed is maintained after passing through the POmniFlect.

Variation—Specifies the variation of the Pass Velocity setting applied to the range of particles.

Display Icon group

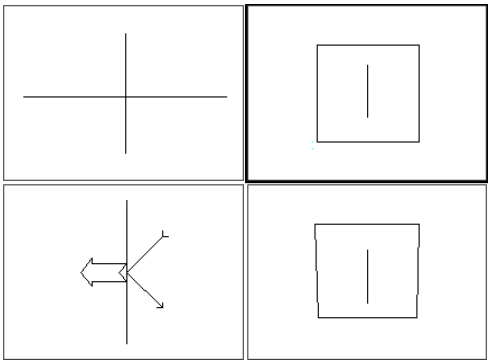
Width/Height—Specify the width and height of the POmniFlect icon. This is for display purposes only and does not influence the deflector effect.

PDynaFlect Space Warp

Create panel > Space Warps > Deflectors > Object Type rollout > PDynaFlect

PDynaFlect (planar dynamics deflector) is a planar version of the dynaflector, a special class of space warp that lets particles affect objects in a dynamics situation. For example, if you want a stream of particles to strike an object and knock it over, like the stream from a firehose striking a stack of boxes, use a dynaflector.

Note: You use dynaflectors in the same way as omniflectors; that is, you can use them as deflectors alone, with no dynamics simulation. Because they take physics into account, dynaflectors are slower than omniflectors. Therefore, it is recommended that you use dynaflectors only when a dynamics simulation is involved.



PDynaflect viewport icon

Procedure

Many dynaflector parameters are the same as those for omniflectors. However, the procedure of associating the dynaflector with both the particle system and the object to be affected is more complex than the simple binding used by omniflector.

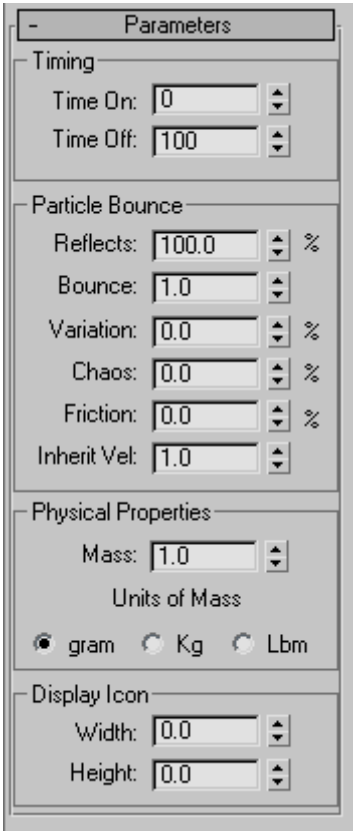
To create a particle/dynamics system:

1. Create or load a scene containing a *non-event-driven particle system (page 2–237)* and an object to be affected by the particles (hereafter called "the object"). Position and orient the two as desired so that the particles strike the object.
   - The icon for the particle system must be visible in at least one non-disabled viewport.
   - On the Basic Parameters rollout of the particle system, set Percentage of Particles to 100%. Otherwise, only those particles displayed in the viewport will be calculated.
2.  On the Create panel, click Space Warps. Choose Deflectors from the drop-down list, and then click one of the dynaflector buttons (PDynaFlect, SDynaFlect, or UDynaFlect). Drag in a viewport to create the deflector.
3. Do either of the following:
   - If using a PDynaFlect or SDynaFlect, position the space warp icon where the particles strike (or will strike) the object. Resize and orient it as needed.

      Also, link the PDynaFlect or SDynaFlect as a child of the object. This step isn't absolutely necessary, but allows the deflector to follow the object as it moves.
   - If using a UDynaFlect, go to the Modify panel > Parameters rollout, click the Pick Object button, and then select the object. The position of the space warp icon doesn't matter, and it's not necessary to link the two.
4.  Bind the particle system to the dynaflector using the toolbar *Bind to Space Warp button (page 2–58)*.
5.  Go to the Utilities panel and open the Dynamics utility.
6. On the Dynamics rollout, click New to create a new simulation.
7. Click Edit Object List, and add the object to the dynamics simulation.
8. Click Edit Object. On the Edit Object dialog, click Assign Object Effects and assign the dynaflector to the object.

   Note: Because only those particles displayed in viewports are used in the dynaflector calculations, solving the dynamics involves three requirements:
9. On the Dynamics rollout, turn on the Update Display w/Solve check box above the Solve button.
10. Solve the simulation.

Interface



Timing group

The two spinners specify the start frame and end frame of the deflection effect.

Time On/Off—Time On specifies the frame at which the deflection begins, and Time Off specifies the frame at which the deflection ends.

Particle Bounce group

These settings affect the reflection of particles from the space warp.

Reflects—Specifies the percentage of particles to be reflected by the PDynaFlect.

This affects both the particles and the dynamics reaction of the object struck by the particles. The more particles that strike the affected object, the more force applied to that object. If set to 0.0, the particles have no effect on the object.

Bounce—This is a multiplier that specifies how much of the initial speed of the particle is maintained after collision with the PDynaFlect.

Using the default setting of 1.0 causes the particle to rebound with the same speed as it collides. A real-world effect would usually be less than 1.0. For a "flubber" effect, set greater than 1.0.

Variation—Specifies the variation of Bounce applied to the range of particles.

For example, a Variation setting of 50% applied to a Bounce setting of 1.0 would result in randomly applied Bounce values ranging from 0.5 to 1.5.

Chaos—Applies a random variation to the bounce angle.

When you set Chaos to 0.0 (no chaos), all particles bounce off the PDynaFlect surface perfectly (like banking pool balls). A non-zero setting causes the deflected particles to scatter.

Friction—The amount by which particles are slowed as they move along the deflector surface. A value of 0% means they're not slowed at all. A value of 50% means they're slowed to half their original speed. A value of 100% means they stop moving when they strike the surface. Default=0%. Range=0% to 100%.

Tip: To have particles slide along a deflector surface, set Bounce to 0. Also, unless influenced by a force such as Wind or Gravity, particles meant to slide should strike the surface at an angle other than 90 degrees.

Inherit Vel—(Velocity Inheritance) Determines how much of a moving PDynaFlect's speed is applied to reflected or refracted particles.

For example, if Inherit Vel is 1.0, particles with no motion that are hit by a moving PDynaFlect inherit the speed of the PDynaFlect at the point of collision.

Physical Properties group

These options let you set the mass of each particle.

Mass—Specifies the mass based on the chosen unit.

gram—One gram equals 1/1000 kg or 22/1000 Lbm at a gravity of 1.0.

Kg—One kg (kilogram) equals 1000 grams or 2.2 Lbm at a gravity of 1.0.

Lbm—One Lbm (pounds-mass) equals 5/11 kg or 454 5/11 grams at a gravity of 1.0. (Pounds-mass, the amount of mass in one pound of weight, is dependent on gravity. For pounds-mass values at a gravity other than 1.0, multiply the pounds-mass value by the gravity factor.)

Display Icon group

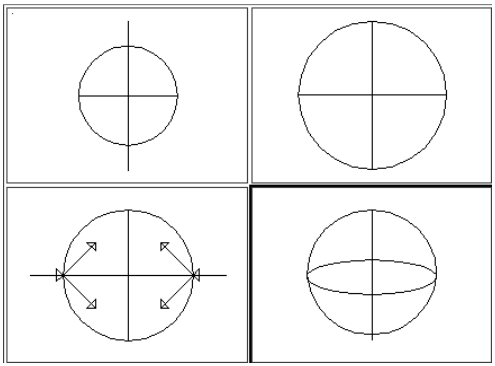
Width/Height—Specify the width and height of the PDynaFlect icon. This is for display purposes only and does not influence the deflector effect.

SOmniFlect Space Warp

Create panel > Space Warps > Deflectors > Object Type rollout > SOmniFlect

SOmniFlect is the spherical version of the omniflector type of space warp. It provides more options than the original SDeflector. Most settings are the same as those in *POmniFlect (page 2–78)*. The difference is that this space warp provides a spherical deflection surface rather than the planar surface. The only settings that are different are in the Display Icon area, in which you set the Radius, instead of the Width and Height.

Note: Opposite sides of the deflectors reverse the distortion effect. Thus, a refracted particle passing through the SOmniFlect hits its outside surface first, and then its inside surface. A positive Distortion value warps the particles toward the perpendicular; then, as the particles pass through the inside surface, the same positive Distortion value warps them toward the parallel.



SOmniFlect viewport icon

Procedure

To create an SOmniFlect space warp:

1.  On the Create panel, click Space Warps. Choose Deflectors from the list, then on the Object Type rollout, click SOmniFlect.
2. Drag in a viewport to create the spherical icon.

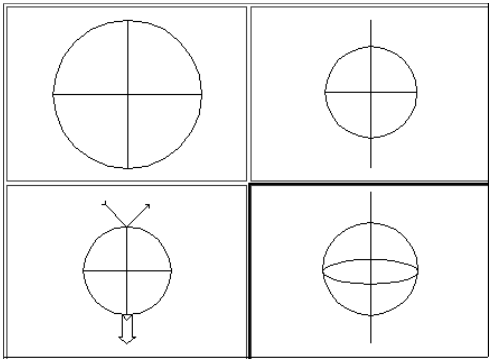
   Note: Because particles bounce off the icon, the size of the icon affects particle deflection.
3. Apply the deflector to the particle system using the appropriate method:
   - If using *Particle Flow (page 2–109)*, specify the deflector in the *Collision test (page 2–212)* or *Collision Spawn test (page 2–215)* parameters.
   -  If using a *non-event-driven particle system (page 2–237)*, *bind (page 2–58)* the particle system to the deflector icon.
4. Position the SOmniFlect icon to interrupt the particle stream.

5. Adjust the SOmniFlect parameters as necessary.

SDynaFlect Space Warp

Create panel > Space Warps > Deflectors > Object Type rollout > SDynaFlect

The SDynaFlect space warp is a spherical dynamics deflector. It's like the *PDynaFlect (page 2–81)* warp, except that it's spherical, and its Display Icon spinner specifies the icon's Radius value.



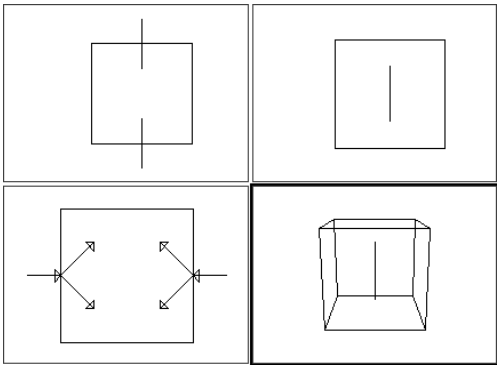
SDynaFlect viewport icon

UOmniFlect Space Warp

Create panel > Space Warps > Deflectors > Object Type rollout > UOmniFlect

UOmniFlect, the universal omniflector, provides more options than the original UDeflector. This space warp lets you use any other geometric object as a particle deflector. The deflections are face accurate, so the geometry can be static, animated, or even morphing or otherwise deforming over time.

Note: Some particle "leakage" can occur with this deflector, particularly when you use many particles and a complex deflector object. To avoid this, perform a test render to check for leaking particles, and then add POmniFlects to catch the strays.



UOmniFlect viewport icon

Procedure

To use the UOmniFlect space warp:

To use a universal omniflector, you need a minimum of three objects in the scene:

- The particle system
- The UOmniFlect space warp
- The object used as the deflector

1. Add or select an object used as the deflector.
2. Create a particle system whose particles intersect the deflector object.
3.  On the Create panel, click Space Warps. Choose Deflectors from the list, and then click UOmniFlect.
4. Click and drag in a viewport to place the space warp icon.

   Note: The size and position of the UOmniFlect icon do not alter the effect.
5. Apply the deflector to the particle system using the appropriate method:
   - If using *Particle Flow (page 2–109)*, specify the deflector in the *Collision test (page 2–212)* or *Collision Spawn test (page 2–215)* parameters.

-  If using a *non-event-driven particle system (page 2–237)*, *bind (page 2–58)* the particle system to the deflector icon.

6.  On the Modify panel > Parameters rollout > Object-Based OmniFlector group, click Pick Object, and then select the object to use as a deflector.
7. Adjust the particle system and UOmniFlect parameters as necessary.

Interface

Parameters rollout

The settings for the UOmniFlect are the same as those for *POmniFlect (page 2–78)*, with the following additions:

Object-Based OmniFlector group

Lets you choose the object to use as a deflector.

Item—Displays the name of the selected object.

Pick Object—Click this, and then select any renderable object to use as a deflector.

Display Icon group

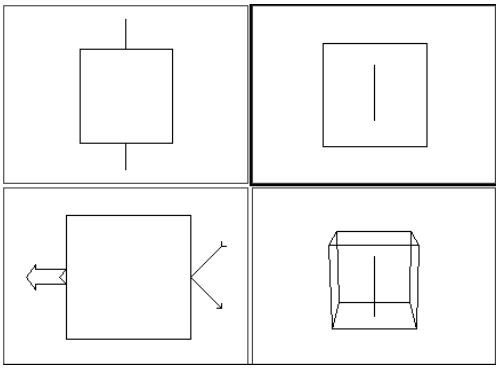
Icon Size—Specifies the size of one side of the square UOmniFlect icon.

UDynaFlect Space Warp

Create panel > Space Warps > Deflectors > Object Type rollout > UDynaFlect

The UDynaFlect space warp is a universal dynamics deflector that lets you use the surface of any object as both the particles deflector and the surface that reacts dynamically to the particle impact.

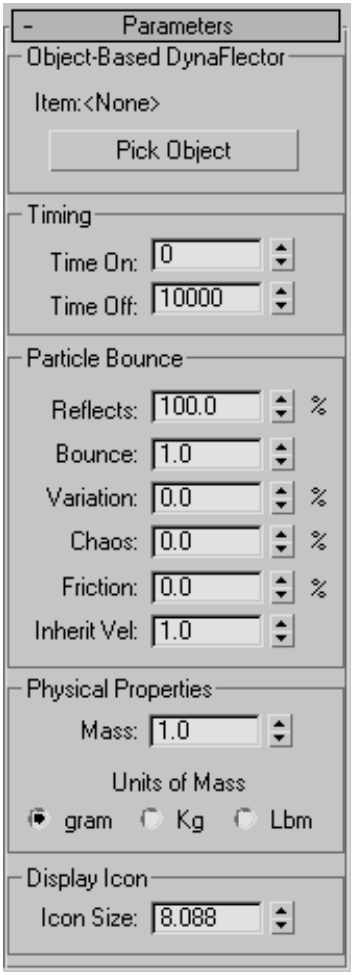
The procedures and options for using UDynaFlect are the same as for *PDynaFlect (page 2–81)*, with the following changes and additions.



UDynaFlect viewport icon

Note: When you use UDynaFlect, you must indicate the object to be affected with the Pick Object button. Linking is not necessary.

Interface



Object-Based DynaFlector group

Lets you choose the object to use as a deflector.

Item—Displays the name of the selected object.

Pick Object—Click this, and then select any renderable object to use as a deflector.

Display Icon group

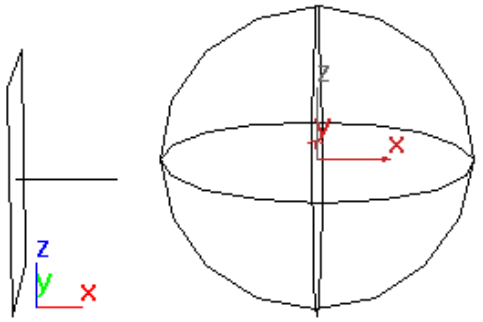
Icon Size—Specifies the size of the UDynaFlect icon.

SDeflector Space Warp

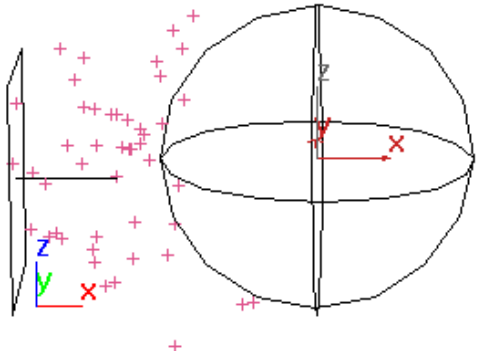
Create panel > Space Warps > Deflectors > Object Type rollout > SDeflector

Create menu > Space Warps > Deflectors > SDeflector

The SDeflector space warp serves as a spherical deflector of particles.



SDeflector viewport icon (with particle system on the left)



SDeflector repelling particles

Procedure

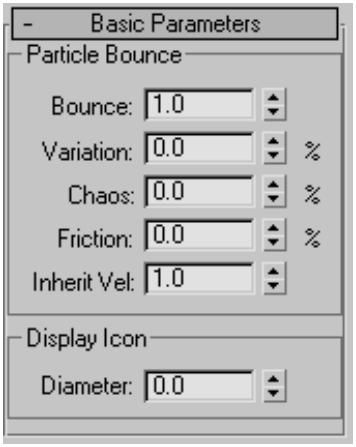
To create an SDeflector:

1.  On the Create panel, click Space Warps. Choose Deflectors from the list, then on the Object Type rollout, click SDeflector.
2. Drag in a viewport to create the spherical icon.

Note: Because particles bounce off the perimeter of the spherical icon, the size of the icon affects particle deflection.

3. Apply the deflector to the particle system using the appropriate method:
   - If using *Particle Flow (page 2–109)*, specify the deflector in the *Collision test (page 2–212)* or *Collision Spawn test (page 2–215)* parameters.
   -  If using a *non-event-driven particle system (page 2–237)*, *bind (page 2–58)* the particle system to the deflector icon.
4. Position the SDeflector icon to interrupt the particle stream.
5. Adjust the SDeflector parameters as necessary.

Interface



Particle Bounce group

These settings determine how the deflector affects the bound particles.

Bounce—Determines the speed with which particles bounce off the deflector. At 1.0, the particles bounce at the same speed as they approach. At 0, they don't deflect at all.

Variation—The amount by which each particle can vary from the Bounce setting.

Chaos—The amount of variation from the perfect angle of reflection (found when Chaos is set to 0.0). 100% induces a variation in reflection angle of up to 90 degrees

Friction—The amount by which particles are slowed as they move along the deflector surface. A value of 0% means they're not slowed at all. A value of 50% means they're slowed to half their original speed. A value of 100% means they stop moving when they strike the surface. Default=0%. Range=0% to 100%.

Tip: To have particles slide along a deflector surface, set Bounce to 0. Also, unless influenced by a force such as Wind or Gravity, particles meant to slide should strike the surface at an angle other than 90 degrees.

Inherit Vel—(Velocity Inheritance) When the value is greater than 0, the motion of the deflector affects particles as well as the other settings. For example, to animate the SDeflector passing through a passive array of particles, turn up this value to affect the particles.

Display Icon group

This option affects the display of the icon.

Diameter—Specifies the diameter of the SDeflector icon. This setting also alters the effect of the deflection, because particles bounce off the perimeter of the icon. The scale of the icon also affects the particles.

UDeflector Space Warp

Create panel > Space Warps > Deflectors > Object Type rollout > UDeflector

Create menu > Space Warps > Deflectors > UDeflector



Particles scatter when they strike a UDeflector object

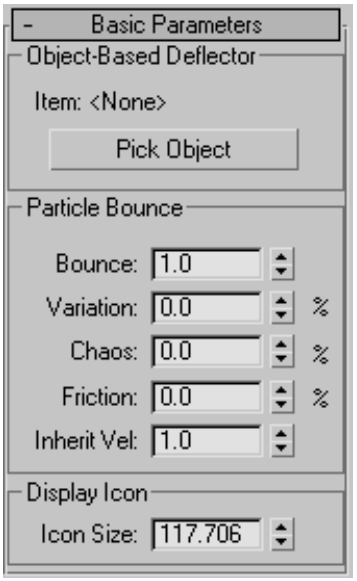
The UDeflector is a universal deflector that lets you use any object as a particle deflector.

Procedure

To create a UDeflector:

1. On the Create panel, click Space Warps. Choose Deflectors from the list, then on the Object Type rollout, click UDeflector.
2. In a viewport, drag out a rectangle to add a UDeflector warp to the scene.
3. On the command panel, click the Pick Object button and select an object to be a particle deflector.
4. Apply the deflector to the particle system using the appropriate method:
   - If using *Particle Flow (page 2–109)*, specify the deflector in the *Collision test (page 2–212)* or *Collision Spawn test (page 2–215)* parameters.
   - If using a *non-event-driven particle system (page 2–237)*, *bind (page 2–58)* the particle system to the deflector icon.
5. Position the UDeflector icon to interrupt the particle stream.
6. Adjust the UDeflector parameters as necessary.

Interface



Object-Based Deflector group

Specifies the object to use as a deflector.

Item—Displays the name of the selected object.

Pick Object—Click this, and then click any renderable mesh object to be used as a deflector.

Particle Bounce group

Bounce—Determines the speed with which particles bounce off the deflector. At 1.0, the particles bounce at the same speed as they approach. At 0, they don't deflect at all.

Variation—The amount by which each particle can vary from the Bounce setting.

Chaos—The amount of variation from the perfect angle of reflection (found when Chaos is set to 0.0). 100% induces a variation in reflection angle of up to 90 degrees.

Friction—The amount by which particles are slowed as they move along the deflector surface. A value of 0% means they're not slowed at all. A value of 50% means they're slowed to half their original speed. A value of 100% means they stop moving when they strike the surface. Default=0%. Range=0% to 100%.

Tip: To have particles slide along a deflector surface, set Bounce to 0. Also, unless influenced by a force such as Wind or Gravity, particles meant to slide should strike the surface at an angle other than 90 degrees.

Inherit Vel—(Velocity Inheritance) When greater than 0, the motion of the deflector affects particles as well as the other settings. For example, to animate the SDeflector passing through a passive array of particles, turn up this value to affect the particles.

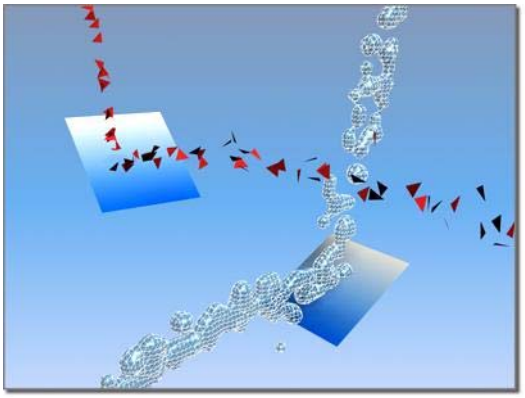
Display Icon group

Icon Size—This spinner displays and lets you change the size of the icon.

Deflector Space Warp

Create panel > Space Warps > Deflectors > Object Type rollout > Deflector

Create menu > Space Warps > Deflectors > Deflector



Two streams of particles striking two deflectors

The Deflector space warp acts as a planar shield to repel the particles generated by a particle system. For example, you can use Deflector to simulate pavement being struck by rain. You can combine a Deflector space warp with a Gravity space warp to produce waterfall and fountain effects.

See also

SDeflector Space Warp (page 2–87)

UDeflector Space Warp (page 2–89)

Procedure

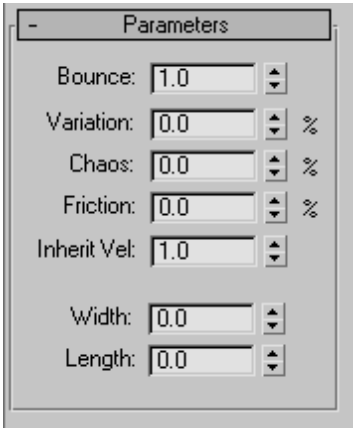
To create a deflector:

1.  On the Create panel, click Space Warps. Choose Deflectors from the list, then on the Object Type rollout, click Deflector.
2. Drag in a viewport to define the deflection area.

   The deflector appears as a wireframe rectangle.
3. Apply the deflector to the particle system using the appropriate method:

- If using *Particle Flow (page 2–109)*, specify the deflector in the *Collision test (page 2–212)* or *Collision Spawn test (page 2–215)* parameters.
-  If using a *non-event-driven particle system (page 2–237)*, *bind (page 2–58)* the particle system to the deflector icon.

Interface



The deflector's effect is controlled mainly by its size and orientation in the scene, relative to the particle system that is bound to it. You can also adjust how strongly the deflector deflects particles.

Bounce—Controls the speed at which particles bounce off the deflector. At a setting of 1.0, particles bounce off the deflector at the same speed they struck it. At 0.0, particles do not bounce at all. At values between 0.0 and 1.0, particles bounce off the deflector at a speed reduced from their initial speed. At values greater than 1.0, particles bounce off the deflector at a speed greater than their initial speed. Default=1.0.

Variation—The amount by which each particle can vary from the Bounce setting.

Chaos—The amount of variation from the perfect angle of reflection (found when Chaos is set to 0.0). 100% induces a variation in reflection angle of up to 90 degrees

Friction—The amount by which particles are slowed as they move along the deflector surface. A value of 0% means they're not slowed at all. A value of 50% means they're slowed to half their original speed. A value of 100% means they stop moving when they strike the surface. Default=0%. Range=0% to 100%.

Tip: To have particles slide along a deflector surface, set Bounce to 0. Also, unless influenced by a force such as Wind or Gravity, particles meant to slide should strike the surface at an angle other than 90 degrees.

Inherit Vel—(Velocity Inheritance) When the value is greater than 0, the motion of the deflector affects particles as well as the other settings. For example, if you want an animated SDeflector passing through an array of particles to affect the particles, turn up this value.

Width—Sets the deflector's width.

Length—Sets the deflector's length.

Geometric/Deformable

FFD(Box) Space Warp

Create panel > Space Warps > Geometric/Deformable > Object Type rollout > FFD(Box)

Create menu > Space Warps > Geometric/Deformable > FFD(Box)

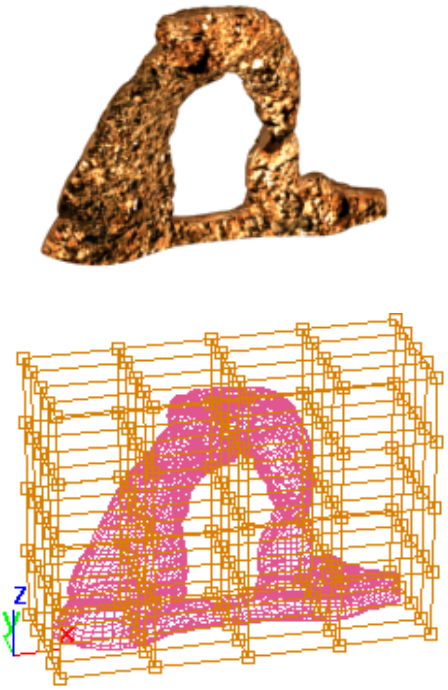
Free-form deformations (FFDs) provide a method of deforming an object by adjusting the control points of a lattice. The offset position of the control points to the original lattice source volume causes the distortion of the affected object.

The FFD(Box) space warp is a box-shaped lattice FFD object similar to the original FFD modifiers. This FFD is available as both an object modifier and a space warp. For information on the object-modifier version, see *FFD (Box/Cylinder) Modifiers (page 1–685)*.

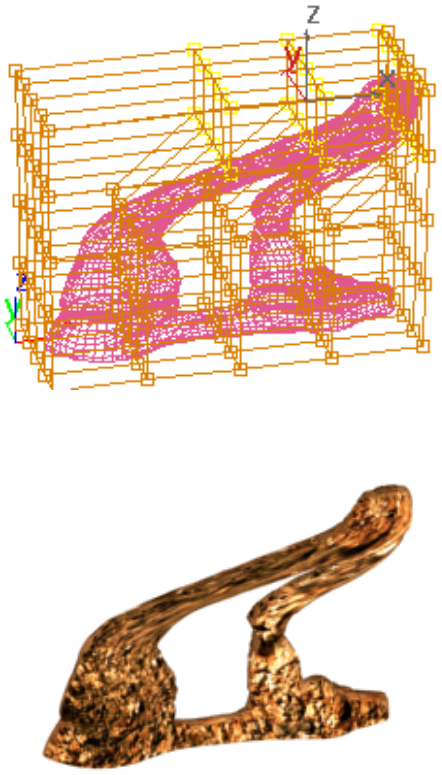
You create FFD space warps as separate objects similarly to the way you create standard primitives: by dragging the mouse in the viewport. The result is a lattice of control points. The source lattice of an FFD modifier is fitted to the geometry it's assigned to in the stack. This might be a whole object or a sub-object selection of faces or vertices.

Because FFD space warps are separate objects, they carry their own adjustable dimension parameters among the creation parameters.

You can apply object modifiers to space warp objects. For example, you can use the Linked XForm modifier with a space-warp FFD.



Object and object surrounded by an FFD lattice



Moving control points in the lattice deforms the object.

Procedure

To use the FFD(box) space warp:

1.  On the Create panel, click Space Warps. Choose Geometric/Deformable from the list, then on the Object Type rollout, click FFD(Box).
2. Drag in a viewport to create the base. Release the mouse button, and then move the mouse to define the height of the FFD lattice. Click to finish the lattice.
3.  Bind the lattice to the object you want to deform.
4. Determine the relative placement of the lattice to the object.

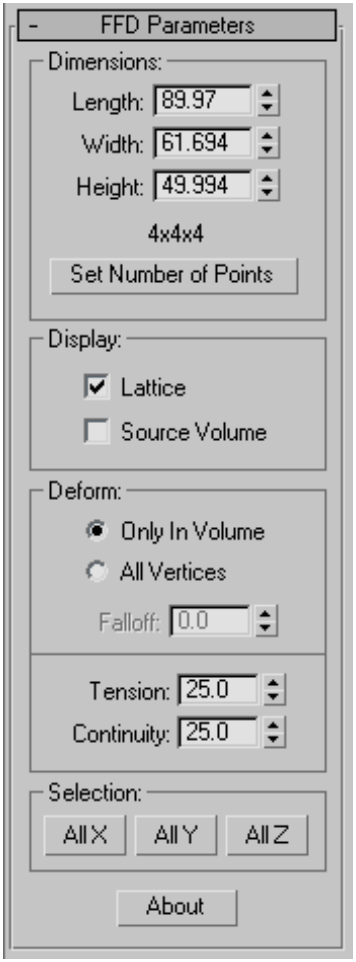
   If the lattice is to be outside of the object, turn on All Vertices. To affect only those vertices inside the lattice, choose Only In Volume, and position the lattice accordingly.
5.  In the modifier stack display, choose Control Points as the sub-object level for FFD(box).
6. Adjust the control points.

   Note: The distortion effect of an FFD modifier is based on the positional offset of the control points from their original positions in the source volume. If you don't move control points, there is no effect on the target object. Keep this in mind when using space-warp version of the FFD.

If you're using the Deform group > All Vertices option, once you've distorted the object you can set the Falloff value to adjust how much the lattice affects the object, based on distance. This is particularly useful if the lattice is animated to approach or move away from the target object. When Falloff is set to 0, all the vertices are affected, regardless of distance.

Note: When you're at the base-parameters level of an FFD modifier in the Stack, the Show End Result button is turned off and spring-loaded, as it is in an Edit Mesh modifier.

Interface



This rollout lets you set the size and resolution of the lattice, and how it displays and deforms.

Dimensions group

These options let you adjust the unit dimensions of the source volume, and specify the number of control points in the lattice. Note that the point dimensions are displayed beside the modifier name in the Stack list.

Length, Width, Height—These three spinners display and let you adjust the length, width, and height of the lattice. To create the space warp, you drag the mouse in the same way that you would to create a standard Box primitive. Note that these spinners don't exist in the object-modifier version of the FFD.

Label—Displays the current number of controls points in the lattice (for example: 4x4x4).

Set Number of Points—Displays a dialog containing three spinners labeled Length, Width, and Height, plus OK/Cancel buttons. Specify the number of control points you want in the lattice, and then click OK to make the change.

Note: Make changes to the dimensions before you adjust the positions of the lattice control points. When you change the number of control points with this dialog, any adjustments you've already made to the control points are lost; however, you can undo this use of the dialog.

Display group

These options affect the display of the FFD in the viewports.

Lattice—When turned on, lines are drawn connecting the control points to make a grid. Although viewports can become cluttered when these extra lines are drawn, they help to visualize the lattice.

Source Volume—When on, the control points and lattice are displayed in their unmodified state. This display is helpful when you're adjusting the source volume to affect specific vertices that lie within or without it. See the All Vertices and Only in Volume options, later in this topic.

Deform group

These options provide controls that specify which vertices are affected by the FFD.

Only In Volume—Only vertices that lie inside the source volume are deformed. Vertices outside the source volume are not affected.

All Vertices—All vertices are deformed regardless of whether they lie inside or outside the source volume, depending on the value in the Falloff spinner. The deformation outside the volume is a continuous extrapolation of the deformation inside the volume. The deformation can be extreme for points far away from the source lattice.

Falloff—This spinner, enabled only when you choose All Vertices, determines the distance from the lattice that the FFD effect will decrease to zero. When this spinner is set to 0, it's effectively turned off, and there is no falloff; that is, all vertices are affected regardless of their distance from the lattice. The units of the Falloff parameter are specified relative to the size of the lattice: A falloff of 1 means that the effect will go to 0 for points that are a lattice width/length/height away from the lattice (depending on which side they are on).

Tension/Continuity—Lets you adjust the tension and continuity of the deformation splines. Although you can't see the splines in an FFD, the lattice and control points represent the structure that controls the splines. As you adjust the control points, you alter the splines (which move through each point). The splines, in turn, deform the geometry of the object. By altering the tension and continuity of the splines, you alter their effect on the object.

Selection group

These options provide additional methods of selecting the control points. You can toggle the state of any combination of the three buttons to select in one, two, or three dimensions at once.

All X, All Y, All Z—When one of these buttons is on and you select a control point, all control points along the local dimension specified by the button are selected as well. By turning on two buttons, you can select all control points in two dimensions.

About—Displays a dialog with copyright and licensing information.

FFD(Cyl) Space Warp

Create panel > Space Warps > Geometric/Deformable > Object Type rollout > FFD(Cyl)

Create menu > Space Warps > Geometric/Deformable > FFD(Cyl)

Free-form deformations (FFDs) provide a method of deforming an object by adjusting the control points of a lattice. The offset position of the control points to the original lattice source volume causes the distortion of the affected object.

The FFD(Cyl) space warp uses a cylindrical array of control points in its lattice. This FFD is available as both an object modifier and a space warp. For information on the object modifier version, see *FFD (Box/Cylinder) Modifiers (page 1–685)*.

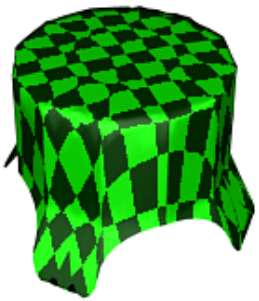
You create FFD space warps as separate objects similarly to the way you create standard primitives: by dragging the mouse in the viewport. The result is a lattice of control points. The source lattice of an FFD modifier is fitted to the geometry it's assigned to in the stack. This might be a whole object or a sub-object selection of faces or vertices.

Because FFD space warps are separate objects, they carry their own adjustable dimension parameters among the creation parameters.

You can apply object modifiers to space warp objects. For example, you can use the Linked XForm modifier with a space-warp FFD.

Procedures

Example: Create an animated tablecloth:



Finished tablecloth using FFD (Cyl) space warp

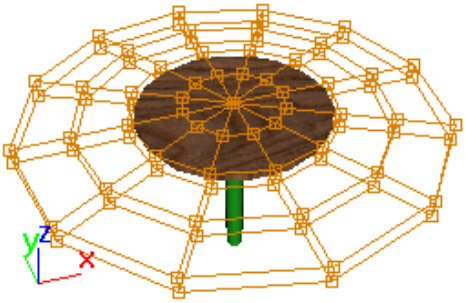
This example shows how to use the FFD(Cyl) space warp to create a tablecloth that flies in and drapes itself over a table.

Begin by creating the table and tablecloth.

1. Create a table from two cylinders. Make the table top with a radius of **30** units, and a height of **2** units. Make the "table stand" cylinder with a radius of **3** and a height of **60**.
2. Make a tablecloth from a box **100** units square and **0.5** units in height. Increase Length and Width Segments to **30**, and keep Height Segments at **1**.
3. Position the tablecloth so it's level with or slightly above the table top, and a little less than **100** units to the left of the table edge, as seen from the Top view.
4. Apply a nice wood grain to the table, and a checker to the tablecloth. (Set the checker tiling to about 15x15, and choose any color for the checkers.)

Now, set up a cylindrical FFD space warp that will form the drape of the tablecloth over the table.

1. From the Create panel > Space Warps > Geometric/Deformable, choose FFD(Cyl).
2. In the Top viewport, create a cylindrical FFD space warp, centered on the table top, with a radius of **45** and a height of **5**.
3. Click the Set Number of Points button and, in the Set FFD Dimensions dialog, set Side points to **12**, Radial points to **5**, and Height points to **2**.
4. Move the entire FFD lattice up until it's just over the surface of the table, as seen from the Front viewport.

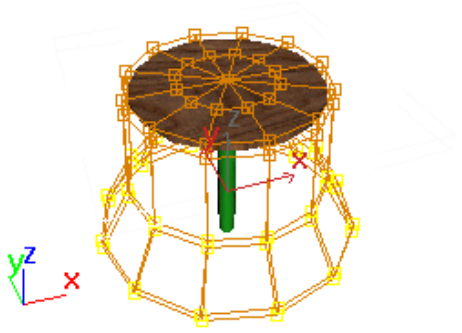


Next, adjust the control points of the lattice to drape over the table.

1. Zoom Extents All Selected.
2. On the Modify panel, in the stack display (below "Modifier List"), click the FFD(cyl) item so it turns yellow. This means you've enabled direct access to the FFD space warps control point sub-objects.
3. In the FFD Parameters rollout > Selection group, turn on All X. This lets you select control points around the perimeter of the FFD cylinder.
4. In the Top viewport, use the *Select and Move tool (page 1–439)* and region-select the two visible control points of the two outer rings of control points at the nine-o'clock position. (This is easier shown than described. You can actually region-select any number of vertices in the two outer concentric rings of vertices.

Because All X is on, all control points for the two rings will be selected.)

5.  On the status bar, click the Selection Lock Toggle button to lock the selection. In the Front viewport, drag the selected points halfway down the height of the table.
6. Unlock the selection, and then region-select any control point in the outer ring to select all points in the outer ring.
7. Lock the selection and, in the Front viewport, drag the outer ring of points down to the floor.



You now have a truncated cone shape over the table.

8. Unlock the selection. In the Top viewport, again select all the control points in the two outer rings.
9. Scale the selected control points in, until the radius of the inner ring of points is slightly larger than the table top.
10. Select only the outer ring of points, and scale them so their radius is slightly larger than the inner ring.

    If necessary, move the points, as seen from the Front viewport, down to touch the floor again.

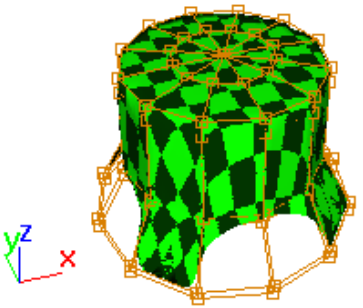
    You now have a lattice of control points draped over the table.

In this next series of steps, you'll bind the tablecloth to the FFD lattice, and then animate it to move over the table.

1. In the stack display, click the FFD(cyl) entry again to exit the sub-object level. It turns gray.
2.  In the main toolbar, click the Bind to Space Warp button, and drag between the FFD and the tablecloth.
3. Select the FFD lattice.
4. Choose Deform group > All Vertices.

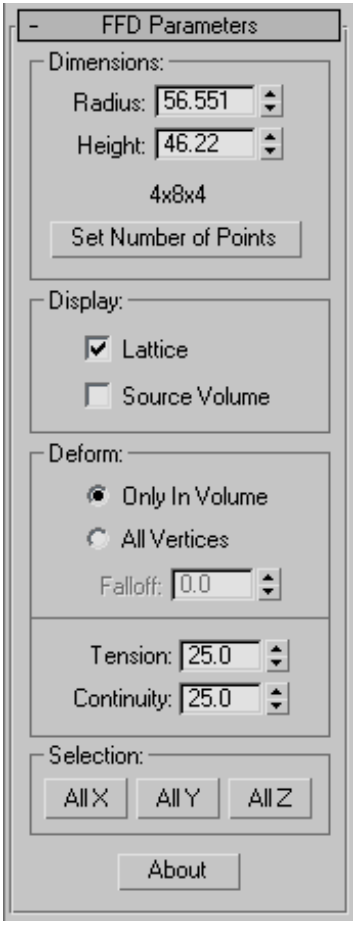
    The tablecloth is immediately deformed because all vertices are now affected, including those outside the lattice volume, and Falloff is set to 0. A falloff value of 0 means that the distance of the vertices from the lattice doesn't matter. Any number greater than 0, however, limits the effect.
5. Set the Falloff spinner to **0.4**.

    No longer influenced by the FFD space warp, the tablecloth returns to its square shape.
6. Turn on Auto Key, and go to frame 100.
7. In the Top viewport, select the tablecloth, and move it until it's centered over the table.
8. As the tablecloth nears the table, it droops down to the floor, sweeps up and over the table, and finally drapes itself over the table.



9. As you move the time slider back and forth and examine the animation, you might find that the bottom of the tablecloth is deformed to the point where it's hanging below the floor. To fix this, turn off Auto Key, go to frame 100, select the lower ring of control points, and move them up until the tablecloth is at the height you want it. You can also adjust the position of the other control points to create drapes, and so on.
10. On the Display command panel > Hide by Category rollout, turn on Space Warps to hide the FFD space warp.
11. Set up appropriate lights and a camera, and play your animation.

Interface



This rollout lets you set the size and resolution of the lattice, and how it displays and deforms.

Dimensions group

These options let you adjust the unit dimensions of the source volume, and specify the number of control points in the lattice. Note that the point dimensions are displayed beside the modifier name in the Stack list.

Radius, Height—These two spinners display and let you adjust the length, width, and height of the lattice. To create the space warp, drag the mouse in the same way that you would to create a standard

Cylinder primitive. Note that these spinners don't exist in the object-modifier version of the FFD.

Label—Displays the current number of controls points in the lattice (for example: 4x8x4).

Set Number of Points—Displays a dialog containing three spinners labeled Side, Radial, and Height, plus OK/Cancel buttons. Specify the number of control points you want in the lattice, and then click OK to make the change.

Side—The number of control points around the perimeter of the lattice.

Radial—The number of control points, radially, from the center to the outer perimeter of the lattice.

Height—The number of control points along the height of the lattice.

Note: Make changes to the dimensions before you adjust the positions of the lattice control points. When you change the number of control points with this dialog, any adjustments you've already made to the control points are lost; however, you can undo this use of the dialog.

Display group

These options affect the display of the FFD in viewports.

Lattice—When on, lines are drawn connecting the control points to make a grid. Although viewports can become cluttered when these extra lines are drawn, they help to visualize the lattice.

Source Volume—When on, the control points and lattice are displayed in their unmodified state. This display is helpful when you're adjusting the source volume to affect specific vertices that lie within or without it. See the All Vertices and Only in Volume options, later in this topic.

Deform group

These options provide controls that specify which vertices are affected by the FFD.

Only In Volume—When on, only vertices that lie inside the source volume are deformed. Vertices outside the source volume are not affected. This is the default choice.

All Vertices—When on, all vertices are deformed regardless of whether they lie inside or outside the source volume, depending on the value in the Falloff spinner. The deformation outside the volume is a continuous extrapolation of the deformation inside the volume. Note that the deformation can be extreme for points far away from the source lattice.

Falloff—This spinner, enabled only when you choose All Vertices, determines the distance from the lattice that the FFD effect will decrease to zero. When this spinner is set to 0, it's effectively turned off, and there is no falloff; that is, all vertices are affected regardless of their distance from the lattice. The units of the Falloff parameter are specified relative to the size of the lattice: A falloff of 1 means that the effect will go to 0 for points that are a lattice width/length/height away from the lattice (depending on which side they are on).

Tension/Continuity—Lets you adjust the tension and continuity of the deformation splines. Although you can't see the splines in an FFD, the lattice and control points represent the structure that controls the splines. As you adjust the control points, you alter the splines (which move through each point). The splines, in turn, deform the geometry of the object. By altering the tension and continuity of the splines, you alter their effect on the object.

Selection group

These options provide additional methods of selecting control points. You can toggle any combination of the three buttons to select in one, two, or three dimensions at once.

All X, All Y, All Z—When one of these buttons is on and you select a control point, all control points

along the local dimension specified by the button are selected as well. By turning on two buttons, you can select all control points in two dimensions.

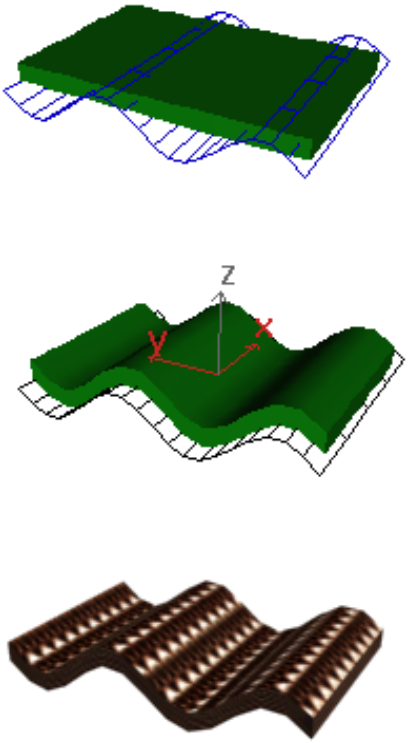
About—Displays a dialog with copyright and licensing information.

Wave Space Warp

Create panel > Space Warps > Geometric/Deformable > Object Type rollout > Wave

Create menu > Space Warps > Geometric/Deformable > Wave

The Wave space warp creates a linear wave through world space. It affects geometry and behaves the same as the *Wave modifier (page 1–957)*. Use the Wave space warp when you want the wave to affect a large number of objects, or to affect an object relative to its position in world space.



Using a wave to deform a box

See also

Ripple Space Warp (page 2–102)

Procedure

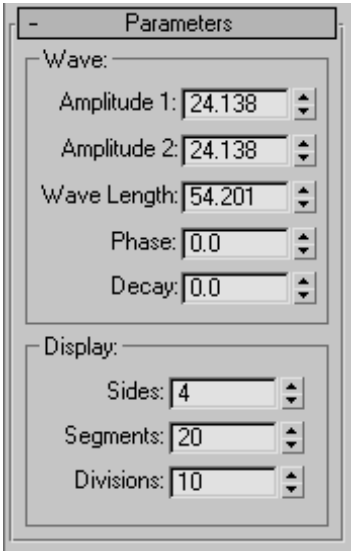
To create a Wave space warp:

1.  On the Create panel, click Space Warps. Choose Geometric/Deformable from the list, then on the Object Type rollout, click Wave.
2. Drag in a viewport to define the initial size of the wave object icon.

   The icon is displayed as a flat mesh wireframe.

3. Release the mouse button to set the icon size; then move the mouse to define the initial amplitude of the wave.
4. Click to set the wave amplitude.

Interface



The initial amplitude sets both Amplitude 1 and Amplitude 2. Set these parameters to unequal values to create a cross wave.

Wave group

These options control the wave effect.

Amplitude 1—Sets wave amplitude along the wave warp object's local X axis.

Amplitude 2—Sets wave amplitude along the wave warp object's local Y axis.

Amplitude is expressed in units. The wave is a sine wave along its Y axis and parabolic along its X axis. Another way to think of the difference between the amplitudes is that Amplitude 1 is at the center of the wave gizmo and Amplitude 2 is at the edge of the gizmo.

Wave Length—Sets the length of each wave along the wave's local Y axis, in active units.

Phase—Offsets the phase of the wave from its origin at the wave object's center. Whole values have no effect; only fractional values do. Animating this parameter makes the wave appear to travel through space.

Decay—When set to 0.0, the wave has the same amplitude or amplitudes throughout world space. Increasing the Decay value causes amplitude to diminish as distance increases from the position of the wave warp object. Default=0.0.

Display group

These options control the geometry of the Wave warp gizmo. In some cases, such as when the two Amplitude values differ, they change the effect of the wave.

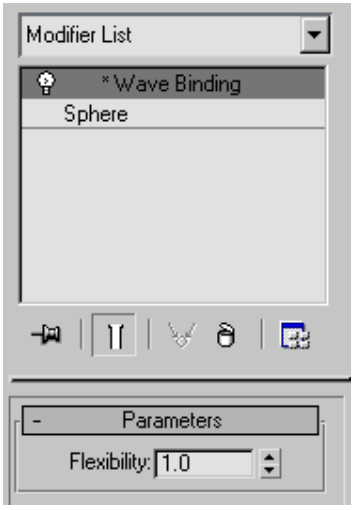
Sides—Sets the number of side segments along the wave object's local X dimension.

Segments—Sets the number of segments along the wave object's local Y dimension.

Divisions—Adjusts the size of the wave icon without altering the wave effect as scaling would.

Flexibility parameter (Modify panel)

The Wave space warp also has a Flexibility parameter that you can adjust individually in each bound object's stack, at the Wave Binding level. The parameter belongs to each binding; it doesn't appear with the Wave warp parameters.

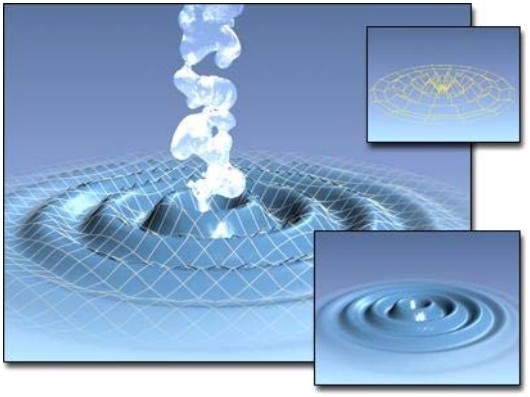


Flexibility—Makes the bound object more or less responsive to the wave by multiplying the amplitude by this value.

Ripple Space Warp

Create panel > Space Warps > Geometric/Deformable > Object Type rollout > Ripple

Create menu > Space Warps > Geometric/Deformable > Ripple



Using a ripple to deform a surface

The Ripple space warp creates a concentric ripple through world space. It affects geometry and behaves the same as the *Ripple modifier (page 1–783)*. Use the Ripple space warp when you want the ripple to affect a large number of objects, or to affect an object relative to its position in world space.

See also

Wave Space Warp (page 2–100)

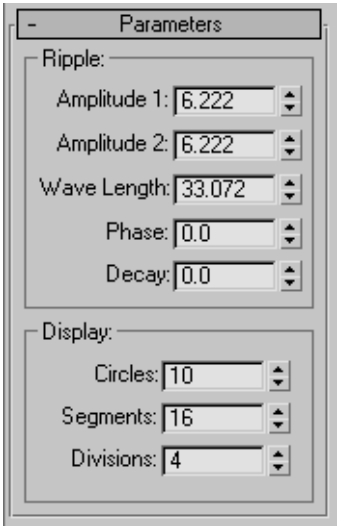
Procedure

To create a Ripple space warp:

1.  On the Create panel, click Space Warps. Choose Geometric/Deformable from the list, then on the Object Type rollout, click Ripple.
2. Drag in a viewport to define the initial size of the ripple object icon.

   The icon is displayed as a wireframe spider web.
3. Release the mouse button to set the icon size, and then move the mouse to define the initial amplitude of the ripple wave.
4. Click to set the wave amplitude.

Interface



The amplitude value set by dragging applies equally in all directions. The ripple's Amplitude 1 and Amplitude 2 parameters are initially equal. Set these parameters to unequal values to create a ripple whose amplitude varies relative to the local X and Y axes of the space warp.

Ripple group

Amplitude 1—Sets ripple amplitude along the ripple warp object's local X axis. Amplitude is expressed in active units.

Amplitude 2—Sets ripple amplitude along the ripple warp object's local Y axis. Amplitude is expressed in active units.

Wave Length—Sets the length of each wave, in active units.

Phase—Offsets the phase of the wave from its origin at the ripple object's center. Whole values have no effect; only fractional values do. Animating this parameter makes the ripple appear to travel through space.

Decay—When set to 0.0, the ripple has the same amplitude or amplitudes throughout world space. Increasing the Decay value causes amplitude to diminish as distance increases from the position of the ripple warp object. Default=0.0.

Display group

The options control display of the Ripple warp object's icon. They don't change the effect of the ripple.

Circles—Sets the number of circles in the ripple icon.

Segments—Sets the number of segments (pie slices) in the ripple icon.

Divisions—Adjusts the size of the ripple icon without altering the ripple effect as scaling would.

Flexibility parameter (Modify panel)

The Ripple space warp also has a Flexibility parameter that you can adjust individually in each bound object's stack, at the Ripple Binding level. The parameter belongs to each binding; it doesn't appear with the Ripple warp parameters.



Flexibility—Makes the bound object more or less responsive to the wave by multiplying the amplitude by this value.

Conform Space Warp

Create panel > Space Warps > Geometric/Deformable > Object Type rollout > Conform

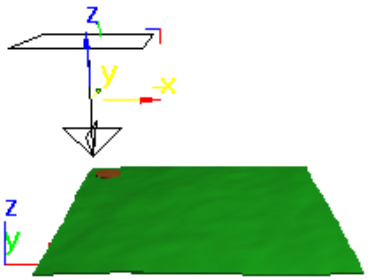
Create menu > Space Warps > Geometric/Deformable > Conform

The Conform space warp modifies its bound object by pushing its vertices in the direction indicated by the space warp icon, until they hit a specified target object, or until the vertices move a specified distance from their original position.

After creating a Conform space warp, you specify a target object in the Conform parameters, and then bind the Conform to the object you want to deform. You rotate the Conform icon to specify

the travel direction (toward the target object). The vertices of the deformed object move until they hit the target object.

There is also a *Conform compound object (page 1–324)* that provides additional methods of conforming one object to another.



Conform viewport icon (a surface is below it)

Procedures

Example: Using the Conform space warp:

Begin by making two objects.

1. Create a terrain by making a wide, flat box with plenty of Length and Width segments (or a quad patch). Apply a Noise modifier and adjust its parameters to result in a bumpy terrain (not mountainous, but low and irregular).
2. Create a short, wide cylinder whose radius is about one-eighth the area of the box (like a coin). You'll animate the cylinder to float diagonally over the surface of the terrain.
3. Set Cap Segments in the cylinder to **4**, and position the cylinder to float over the terrain.
4. Set the object color of the cylinder to contrast with the color of the terrain.
5. Move the cylinder inside one corner of the terrain as seen from the Top viewport. Turn on Animate, move to frame 100, and move the cylinder to the opposite corner of the terrain.

   The coin/disk moves from one corner of the terrain to the other.

The terrain will become the target object, and the cylinder the deformed object. The next step is to create the Conform space warp and bind it to the cylinder.

1. On the Create panel, choose Space Warps, and then, from the drop-down list, choose Geometric/Deformable. Click the Conform button.
2. In the Top viewport, in the center of the terrain, drag outward to create the Conform space warp.
3. Click the Pick Object button, and then click the terrain box.
4. In the Front viewport, drag the space warp up until it's above the cylinder.

   As seen in the Front viewport, the terrain is at the bottom, the cylinder is between the terrain and the space warp, and the space warp is at the top.
5. Bind the space warp to the cylinder.

   The cylinder becomes a disk that seems to be painted on the surface of the terrain.
6. Drag the time slider to see the cylinder/disk move across the box, following the terrain.

Because the vertices are pushed almost to the level of the terrain, the faces of the two object might intersect. In the following steps, you'll fix this by adjusting the standoff distance between the target surface and the pushed vertices. Then you'll go on to affect only selected vertices in the cylinder.

1. Select the Conform space warp, and open the Modify panel.
2. Set Standoff Distance to **3**.

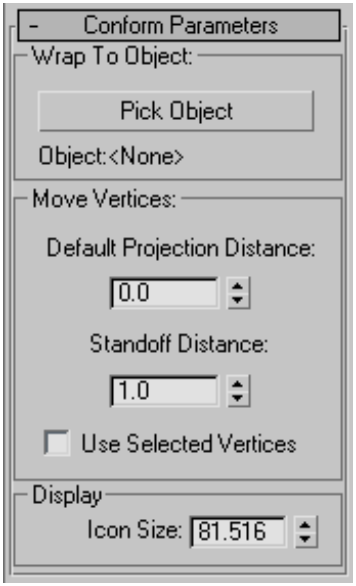
   You can now clearly see the surface of the disk above the terrain.
3. Set the Standoff Distance to **20**.

The disk floats 20 units above the terrain. Next, change the affected vertices.

4. Select the disk/cylinder.
5. In the modifier stack display, click the Cylinder item so it's highlighted in gray.
6. Apply a Mesh Select modifier.
7. At the Vertex sub-object level, in the Front viewport, region-select the bottom cap vertices of the cylinder.
8. Remain at the sub-object level, and in the stack display click the Conform Binding item.
9. In the viewport, select the Conform icon.
10. In the Modify panel, turn on Use Selected Vertices.

    Now that only the bottom cap vertices are selected, the rest of the cylinder is restored. If you adjust the viewing angle and play the animation, you'll see that the bottom face of the cylinder follows the terrain, while the rest of the cylinder retains its shape.

Interface



Wrap To Object group

These options provide controls to select the target object.

Pick Object—Click this, and then select an object in the scene. The object you select becomes the barrier against which the bound object's vertices will be pushed.

Object—Displays the name of the picked object.

Move Vertices group

These options affect how the vertices are moved.

Default Projection Distance—The distance a vertex in the bound object moves from its original location if it does not intersect the target object.

Standoff Distance—The distance maintained between the vertex and the surface of the target object. For example, if set to 5, the vertices can be pushed no closer than 5 units from the surface of the target object.

Use Selected Vertices—When on, only the sub-object selection of vertices on the Stack are pushed. When off, all vertices in the object are pushed regardless of the Stack selection.

Display group

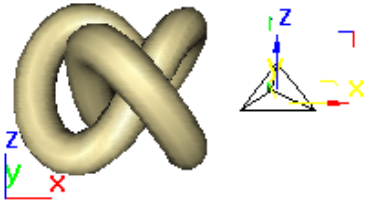
Icon Size—Specifies the size of the icon.

Bomb Space Warp

Create panel > Space Warps > Geometric/Deformable > Object Type rollout > Bomb

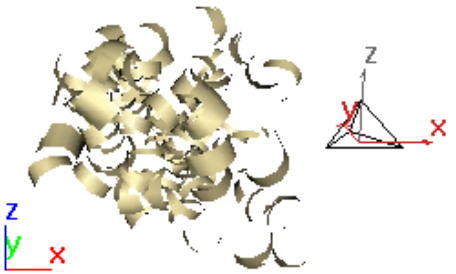
Create menu > Space Warps > Geometric/Deformable > Bomb

The Bomb space warp explodes objects into their individual faces.



Right: Bomb viewport icon

Left: Torus knot



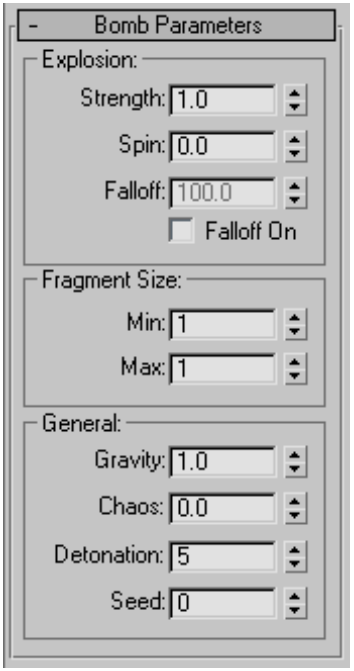
Effect of exploding the torus knot

Procedure

To create a Bomb space warp:

1.  On the Create panel, click Space Warps. Choose Geometric/Deformable from the list, and then on the Object Type rollout, click Bomb.
2. Create mesh objects to be exploded.
3.  On the toolbar, click the Bind to Space Warp button.
4. Drag the mouse between each object and the Bomb space warp.
5. Adjust Bomb parameters to achieve different effects.

Interface



Explosion group

Strength—Sets the power of the bomb. Larger values make the particles fly farther. The closer an object is to the bomb, the greater the effect of the bomb.

Spin—The rate at which fragments rotate, in revolutions per second. This is also affected by the Chaos parameter (which causes different fragments to rotate at different speeds), and by the Falloff parameter (which causes the force of the explosion to be weaker the farther the fragment is from the bomb).

Falloff—The distance from the bomb, in world units, of the effect of the bomb. Fragments past this distance are not affected by the Strength and Spin settings, but are affected by the Gravity setting.

For example, this is useful for blowing up the base of a building, and having the top of the building topple. To see the effect, place a bomb at the base

of a tall cylinder with many height segments, and adjust Falloff to be less than the height of the cylinder.

Falloff On—Turn on to use the Falloff setting. The falloff range appears as a yellow, tri-hooped sphere.

Fragment Size group

These two parameters define the number of faces per fragment. Any given fragment will have a number of faces, randomly determined, between the Min and Max values.

Min—Specifies the minimum number of faces per fragment to be randomly generated by the "explosion."

Max—Specifies the maximum number of faces per fragment to be randomly generated by the "explosion."

General group

Gravity—Specifies the acceleration due to gravity. Note that gravity is always in the direction of the world Z axis. You can have negative gravity.

Chaos—Adds random variation to the explosion to make it less uniform. A setting of 0.0 is totally uniform; 1.0 is a realistic setting. A value greater than 1.0 makes the explosion extra chaotic. Range=0.0 to 10.0.

Detonation—Specifies the frame at which the bomb goes off. Bound objects are unaffected before this time.

Seed—Change to alter randomly generated numbers in the bomb. You can achieve a different bomb effect by changing Seed while maintaining the other settings.

Although you can animate strength and gravity, the equations used for projectile motion assume they are constant. Therefore, the motion will not be physically correct, but it might look interesting. Also, if the bomb object is in motion during the blast, the result is not physically correct.

Modifier-Based

Create panel > Space Warps > Modifier-Based

Create menu > Space Warps > Modifier-Based



Modifier-based space warps duplicate the effects of standard object modifiers. Like other space warps, they must be bound to objects, and they work in world space. They are useful when you want to apply effects such as Twist or Bend to a widely scattered group of objects.

Creating Modifier-Based Space Warps

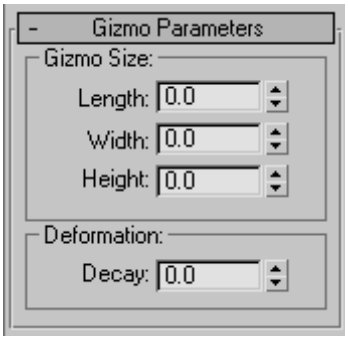
You create modifier-based space warps the same way that you create other space warps. Modifier-based space warps are listed on the Create panel as a separate category under Space Warps.

All the modifier-based space warps use a box-shaped (nonrenderable) object. When you create one, you use the mouse in the viewport as you do when creating a Box primitive.

Unlike their modifier versions, these space warps do not have sub-object levels.

Interface

Gizmo Parameters rollout



Gizmo Size group

Length/Width/Height—Let you adjust the warp object's dimensions.

Deformation group

Decay—When is set to 0, there is no decay, and the space warp affects its bound object regardless of its distance from the object. When you increase the decay, the effect on the bound object falls off exponentially. See the topics on the individual modifiers for more information.

Parameter rollout

The parameters for a modifier-based space warp are identical to those of the modifier on which the space warp is based:

Bend Modifier (page 1–560)

Taper Modifier (page 1–863)

Noise Modifier (page 1–743)

Twist Modifier (page 1–876)

Skew Modifier (page 1–790)

Stretch Modifier (page 1–836)

Particle Systems

Particle systems are useful for a variety of animation tasks. Primarily, they're employed when animating a large number of small objects using procedural methods; for instance, creating a snowstorm, a stream of water, or an explosion. 3ds Max provides two different types of particle systems: event-driven and non-event-driven. The event-driven particle system, also known as *Particle Flow (page 2–109)*, tests particle properties, and, based on the test results, sends them to different events. Each event assigns various attributes and behaviors to the particles while they're in the event. In the *non-event-driven systems (page 2–237)*, particles typically exhibit similar properties throughout the animation.

Important: Particle systems can involve a great many entities, each of which is potentially subject to any number of complex calculations. For this reason, when using them for advanced simulations, you should have a very fast computer and as much memory as possible. Also, a powerful graphics card helps speed the display of particle geometry in the viewports. Even so, it's still easy to overburden the system; if you encounter loss of responsiveness, wait for the particle system to finish its calculations, and then reduce the number of particles in the system, implement a cache, or use other methods to optimize performance.

Which Particle System to Use?

Having access to a wealth of particle systems in 3ds Max leads to the need to decide which system to use for a particular application. In general, for a simple animation, such as falling snow or a water fountain, setup is faster and easier with a non-event-driven particle system. With more complex animations, such as an explosion that generates different types of particles over time (for example: fragments, fire, and smoke), use Particle Flow for greatest flexibility and control.

Particle Flow

Particle Flow is a new, extremely versatile and powerful particle system for 3ds Max. It employs an event-driven model, using a special dialog called *Particle View (page 2–125)*. In Particle View, you combine individual *operators (page 3–985)* that describe particle properties such as shape, speed, direction, and rotation over a period of time into groups called *events (page 3–935)*. Each operator provides a set of parameters, many of which you can animate to change particle behavior during the event. As the event transpires, Particle Flow continually evaluates each operator in the list and updates the particle system accordingly.

To achieve more substantial changes in particle properties and behavior, you can create a *flow (page 3–942)*. The flow sends particles from event to event using *tests (page 3–1021)*, which let you *wire (page 3–1033)* events together in series. A test can check, for example, whether a particle has passed a certain age, how fast it's moving, or whether it has collided with a deflector. Particles that pass the test move on to the next event, while those that don't meet the test criteria remain in the current event, possibly to undergo other tests.

Note: Operators and tests are known collectively as *actions (page 3–903)*.

Following are the principal sections in the Particle Flow documentation:

Using Particle Flow (page 2–109)

Particle Flow User Interface (page 2–125)

Actions (page 2–141)

The Particle Flow Tutorials.

See also

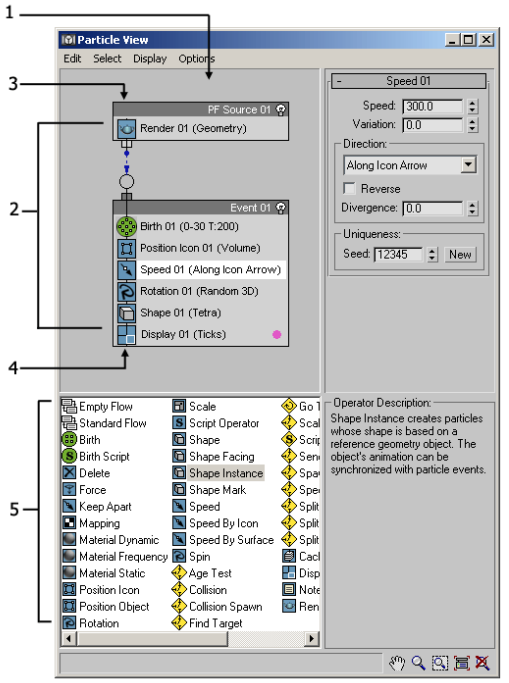
Non-Event-Driven Particle Systems (page 2–237)

Using Particle Flow

Particle View (page 2–125) is the primary interface for building and modifying Particle Flow systems. The first event in the system is always a *global event*, whose contents affect all particles in the system. It has the same name as the Particle Flow source icon. By default, the global event contains a single Render operator that specifies rendering properties for all particles in the system. You can add other operators here to have them act globally, such as Material, Display, and Speed. When you use an operator globally, be sure not to use the same operator locally (that is, in any other events in the system) to avoid potential conflicts.

The global event also serves as the Particle View representation of the particle system. You can create a new system by duplicating this event, or by adding an Empty Flow or Standard Flow. Conversely, if you clone the Particle Flow icon in a viewport, or add a new PF Source, the new system appears in Particle View as well.

The second event is called the *birth event*, because it must contain a Birth operator. The Birth operator should exist at the top of the birth event, and in no other place. The default birth event also contains a number of operators that act locally to specify properties of particles while in that event. The default particle system provides a basic global event and birth event that serve as a useful starting point for creating your own system. If you like, you can instead start with an empty system that lets you build a particle system from scratch.



1. Event display

2. Particle diagram

3. Global event

4. Birth event

5. Depot

To add an action to the *particle diagram*, you drag it to the *event display* from the *depot* (the area at the bottom of the Particle View dialog). If you drag an action to an event, you can add it to the event or replace an existing action, depending on where you drop it. If you drop it in an empty area, it creates a new event. Then, to customize the action, you click its event entry, and then edit its settings in the parameters panel at the side of Particle View.

To add complexity to the particle system, you can add a test to an event, and then wire the test to another event. You can adjust the test parameters to affect particle behavior, as well as determine whether specific conditions exist. When particles meet these conditions, they become eligible for redirection to the next event.

Particle Flow provides a number of tools for determining where in the system particles currently reside, including the ability to change particle color and shape on an event-by-event basis. You can also easily enable and disable actions and events, and determine the number of particles in each event. To speed up checking particle activity at different times during the animation, you can cache particle motion in memory. Using these tools, plus the ability to create custom actions with scripting, you can create particle systems of a level of sophistication previously unachievable.

The Life of a Particle

Another way of looking at Particle Flow is from the perspective of an individual particle. Each particle first comes into existence, or is *born*, via the *Birth operator (page 2–143)*, which lets you specify when to start and stop creating particles, and how many to create.

The particles first appear at an object called an emitter. By default, the emitter is the Particle Flow source icon using the *Position Icon operator (page 2–147)*, but you can alternatively use the *Position Object operator (page 2–148)* to specify that particles should be born on the surface of or within any mesh object in the scene.

After being born, particles can remain stationary at the emission point, or start moving in two different ways. First, they can move, physically, within the scene at a speed and in a direction specified by various actions. These are typically Speed operators, but other actions can also affect particle motion, including *Spin (page 2–154)* and *Find Target (page 2–218)*. In addition, you can use the *Force operator (page 2–204)* to affect their motion with outside forces.



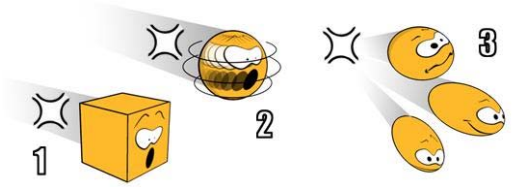
1. Particle immediately after creation, with no speed.

2. The Speed operator sets the particle in motion.

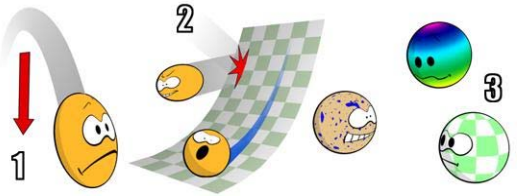
3. The particle continues moving until acted upon by another action.

The second way that particles move is logically, from event to event through the particle diagram, as constructed in *Particle View (page 2–125)*. Each event can contain any number of operators that can affect, in addition to motion, a particle's surface appearance, its shape and size, and others.

The particles start out in the birth event, which typically is the first event after the global event. During a particle's residence in an event, Particle Flow evaluates each of the event's actions from top to bottom, once per integration step, and makes any applicable changes to the particle. If the event contains a test, Particle Flow determines whether the particle tests True for the test's parameters, such as whether it has collided with an object in the scene. If it does, and if the test is wired to another event, Particle Flow sends the particle to the next event. If it doesn't, the particle remains in the current event, and may be further acted upon by its operators and tests. Thus, each particle exists in only one event at a time.

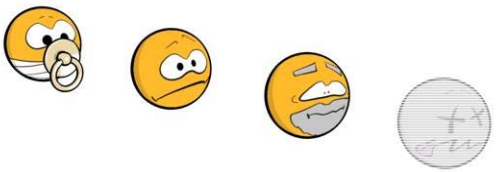


An action in an event can change the particle shape (1), or the particle spin (2), or spawn new particles (3).



Actions can also apply forces to particles (1), specify collision effects (2), and alter surface properties (3).

In this way, the particle continues to travel through the system. Due to the flexible nature of schematic construction in Particle Flow, a particle may be redirected to the same event several times. But at some point, you might want the particle's life to end. For this purpose, you'd use the *Delete operator (page 2–146)* or the *Collision Spawn test (page 2–215)* or *Spawn test (page 2–230)*. Otherwise, the particle lives throughout the entire animation.



Particle age can be used to kill a particle.

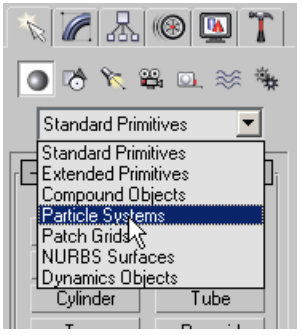
As a particle moves through the system, it's accompanied by a number of channels. For example, each particle has a speed channel that defines how fast it moves, and a material ID channel that lets Particle Flow know which sub-material to apply to it. However, the material itself is not defined by a channel, but by a Material operator that acts locally or globally. Properties that are defined by channels persist, unless altered by an action. For example, the *Material Dynamic operator (page 2–191)* can change a particle's material ID. In effect, by setting up a particle diagram and modifying how particles look and act during the animation, you're deciding how channel values change based on events and animation keyframes.

Introductory Tutorial

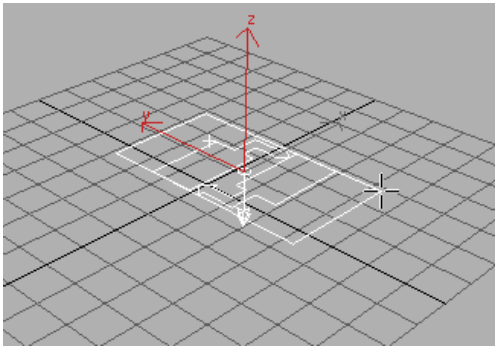
The best way to understand Particle Flow is by using it. This tutorial will acquaint you with some of the basic methods of working with Particle Flow.

Create the Particle Flow system:

1. Start 3ds Max or reset the program.
2. On the Create panel > Geometry category, click the drop-down list and choose Particle Systems.

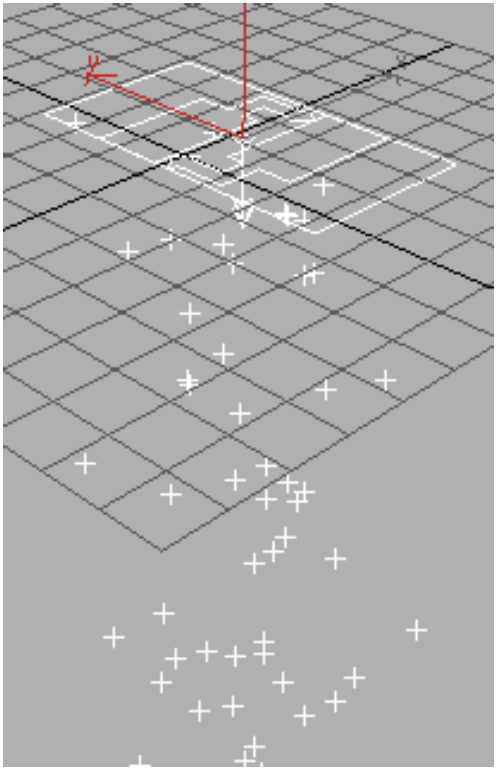


3. On the Object Type rollout, click PF Source.
4. In the Perspective viewport, drag out a rectangle.



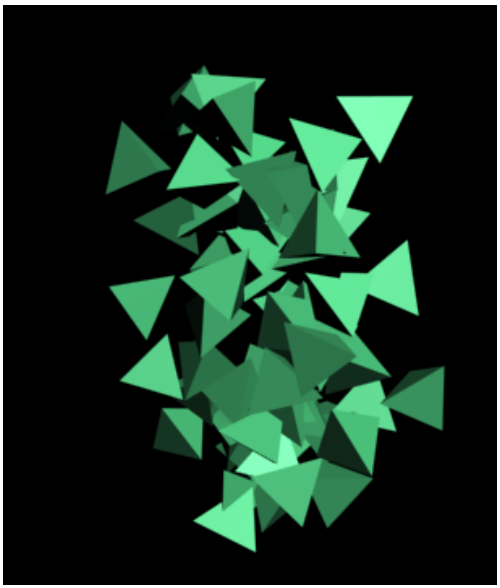
This is the Particle Flow icon, or source, named PF Source 01. By default, it acts as an emitter, but you can also use any other object as an emitter.

5. Drag the time slider.



By default, the icon emits particles downward from its entire surface. In the viewport, the particles appear as ticks.

6. Go to frame 10, and press Shift+Q to render the Perspective viewport.

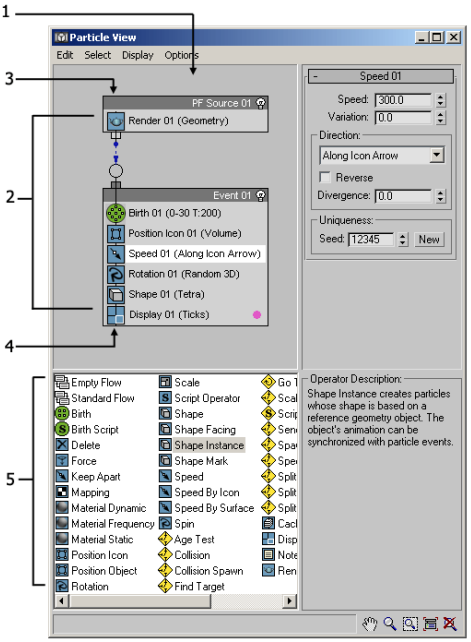


The rendered particles appear in a window. The default particle shape is a tetrahedron, a four-sided triangular solid. Its geometry is very simple, so the system can handle many particles quickly and efficiently, but it gives a good idea of how particles are behaving. Other basic shapes are a low-poly sphere and a cube; Particle Flow also lets you use any scene object as particle geometry.

Modify the particle system in Particle View:

1. Press the 6 key to open Particle View. The particle source icon need not be selected.

   Tip: You can also open Particle View from the command panel when a Particle Flow source icon is selected.



1. Event display

2. Particle diagram

3. Global event

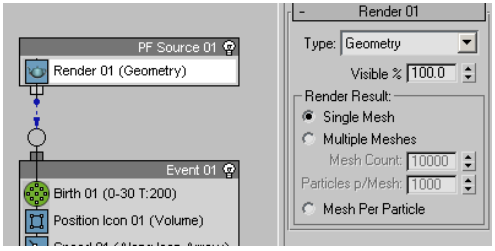
4. Birth event

5. Depot

Particle View is where you manage and edit the particle system. The main window, known as the *event display*, contains the *particle diagram*, which describes the particle system. The default system starts with a *global event*, with the same name as the Particle Flow icon, followed by a *birth event*, containing a Birth operator and several others that define the system's initial properties. Each operator's name is followed by its most important setting or settings in parentheses. Above the event display is a menu bar, and below is the *depot*, containing all operators available for use in the particle system. For more information about this dialog, see Particle View.

2. In the global event, PF Source 01, click the name of the Render 01 (Geometry) operator to highlight it and access its parameters. Be sure to click the text area; clicking the icon simply toggles the operator's active state.

   Because this operator is in the global event, it affects the entire particle system. Any operator you place here does likewise. For example, you can define a global material here, or you can define different materials locally in each event. For best results, don't combine global and local operators of the same type in a single system; use one or the other.



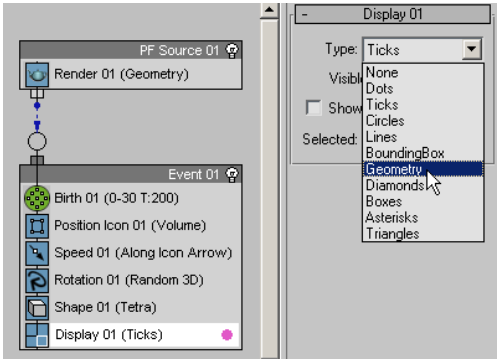
   The settings for the Render operator appear in a rollout on the *parameters panel*, on the right side of the Particle View dialog. These include a drop-down list for how particles are to render, the percentage of particles that will render, and how to split up the particles into individual meshes.

   Next, you'll change the particle display type.

3. In the birth event, Event 01, click the Display 01 (Ticks) operator at the bottom of the list.

   The display type is set to Ticks, as noted in the operator name in the event.

4. Next to the Type label, click Ticks, and from the drop-down list, choose Geometry.

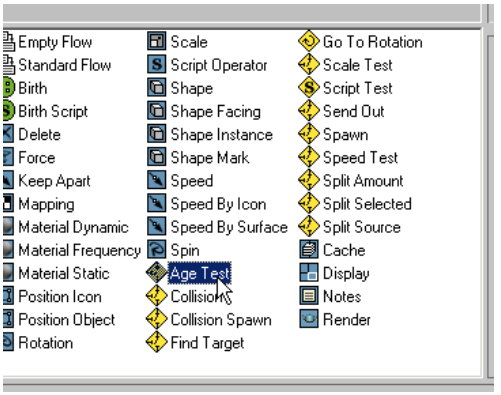


   Now the particles appear as tetrahedrons in the viewports.

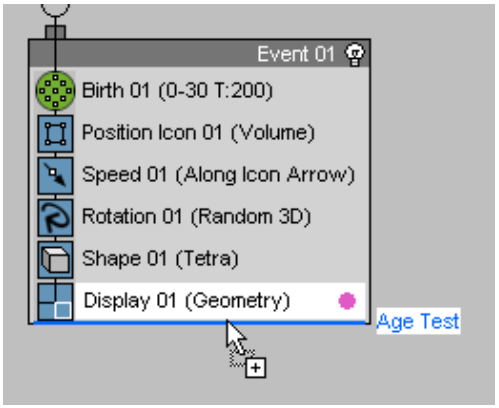
Wire two events together:

Next, you'll add a test and use it to wire the birth event to a new event.

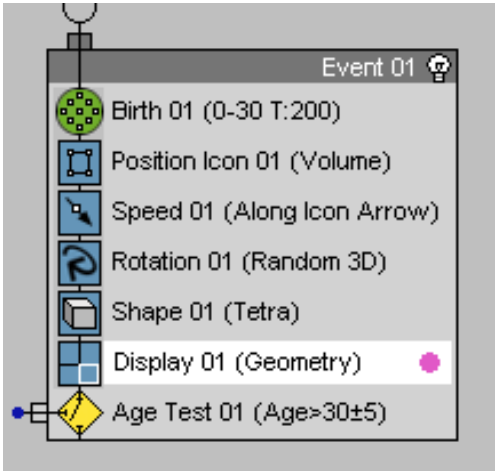
1. In the depot at the bottom of the Particle View dialog, find Age Test. It's the first item that uses a yellow, diamond-shaped icon.



2. Drag an Age Test from the depot into the Event 01 list, at the bottom of the list.

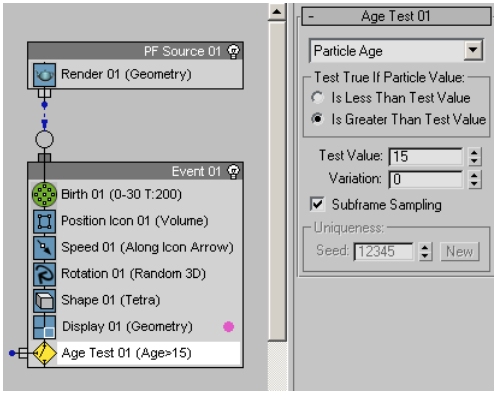


Before you release the mouse button, make sure you see a solid blue line in Event 01 under the Display operator. If the line is red and goes through an existing operator, the Age Test will replace that operator. If you drop the Age Test outside of Event 01 it will create a new event.



Age Test appears in the list, with its *test output* sticking out to the left. This is the part you use to connect the test to the next event.

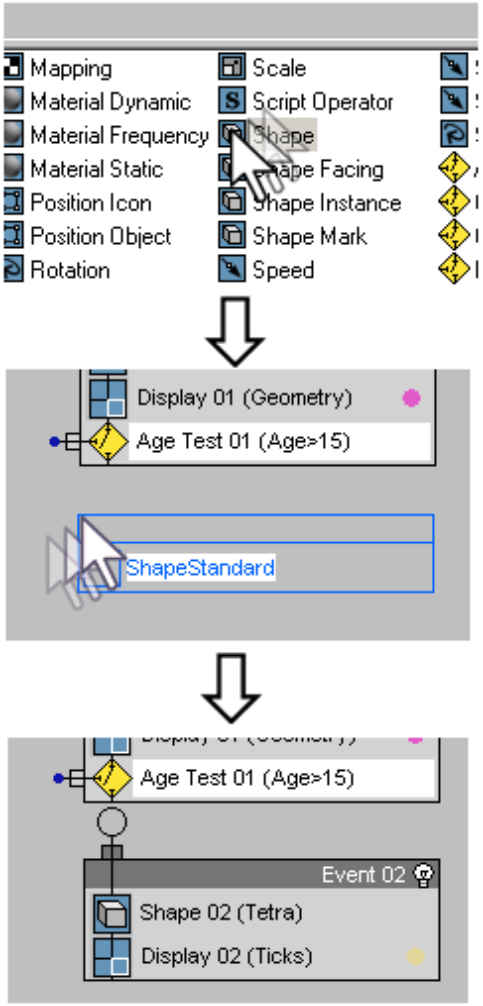
3. Click the Age Test item in the list, and then in the Age Test 01 rollout on the right side of Particle View, set Test Value=**15** and Variation=**0**.



The test type is Particle Age, so this means that all particles that have existed for more than 15 frames will test True, and be passed on to the next event.

Next, you'll create a new event and wire it to the test.

4. From the depot, drag the Shape operator ("Shape") to an empty part of the event display, below Event 01.



The Shape operator appears in a new event, named Event 02. Like Event 01, the event has a circular *event input* sticking out from the top. Also, Particle Flow automatically adds a local Display operator to the event so its particles will be visible in the viewports. You can disable the automatic creation of local Display operators by choosing Options menu > Default Display > Global.

The actual location of an event in the event display doesn't matter; the recommended placement is for the sake of convenience when wiring the events. It also helps to make sense of complex schematics if the events are arranged logically.

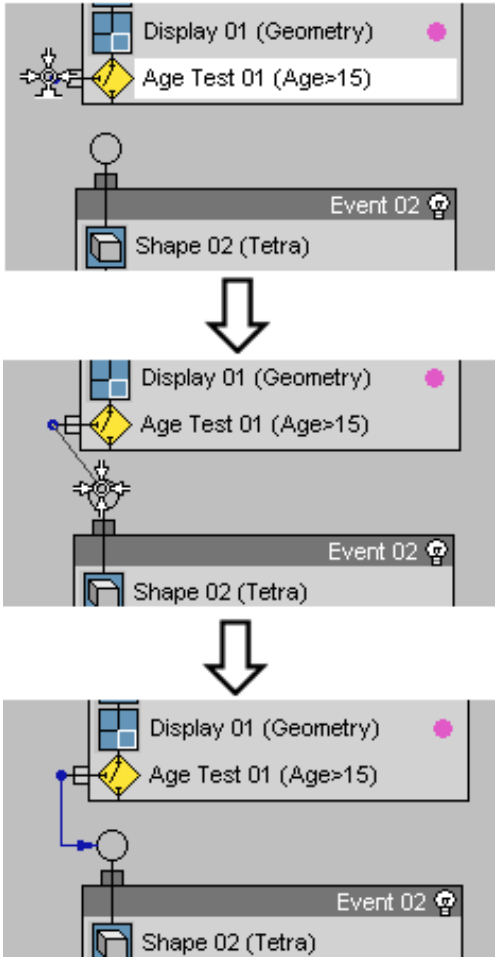
You can move an event by dragging its title bar.

Next, you'll perform the actual wiring of the Age Test to the new event.

5. Position the mouse cursor over the blue dot at the left end of the Age Test's test output.

   The cursor image changes to an icon depicting three arrows pointing inward toward a circular connector.

6. Drag from the event output on the Age Test in Event 01 to the Event 02 input, and then release the mouse button.



As you drag, the cursor image changes from a three-arrow icon to a circular connector by itself when it's over an empty space in the event display area, and then to a four-arrow icon when it's over the Event 02 input.

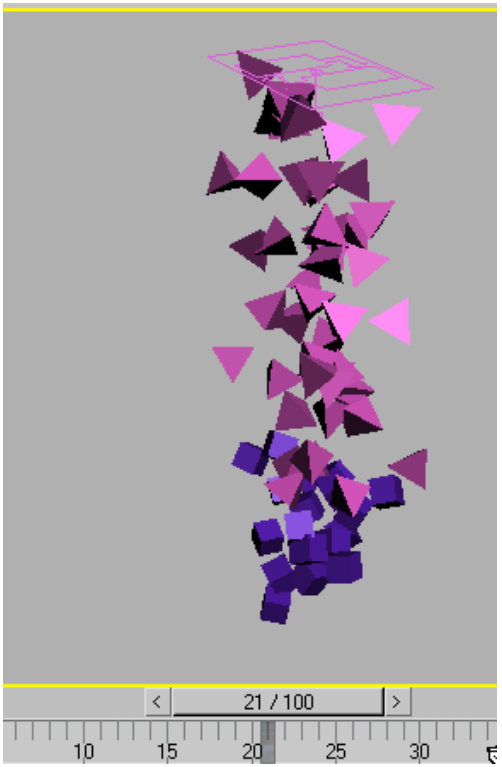
When you release the mouse button, a blue "wire" appears connecting the two events. This wire indicates that particles that meet the Age Test conditions will pass "through" this wire to Event 02, to be affected by its actions.

Tip: You can delete a wire (and anything else in the particle diagram) by selecting it and pressing the **Delete** key. Feel free to try this now, but undo any changes when you're finished.

7. Click the Shape 02 operator and set Shape=Cube. Also, click the Display 02 operator and set Type=Geometry.

View the wiring results:

1. Play the animation. Adjust the viewports so you can see the entire particle stream, if necessary.



Starting at frame 16, the particles at the head of the stream change into cubes, indicating that they've entered Event 02. As time goes on, more and more particles pass the age of 15 and become eligible to go to the next event.

2. While the animation is playing, try modifying the different operator settings to see what happens. For instance, click Speed 01 and

change the Speed and Direction settings. When you change a setting, the change is reflected in the viewports in real time, even during playback.

Try right-clicking actions and events and using Rename to give them custom names. Even with a custom name, an action's tooltip reveals its type.

Another facility of the right-click menu is to add comments to actions and events. Once you do so, a small red triangle indicates the comment's presence.

Congratulations! You now have a working knowledge of Particle Flow. If you'd like to try some more tutorials, please go to this page and download the 3ds Max 8 Tutorials – Online Help (or PDF versions) and the Scene and Support Files.

Particle Flow FAQ

This topic offers answers to a number of questions users commonly ask when first learning to use Particle Flow. The first section contains links to all the questions, and the subsequent sections contain the questions and answers organized by category.

Question List

How does Particle Flow handle time? (page 2–119)

When I go to a different frame, the software sometimes seems to freeze for a while. What's going on? (page 2–119)

Is there a way to pre-calculate portions of a particle simulation, as in "baking" the animation? (page 2–120)

What else can I do to optimize performance? (page 2–120)

How do I use Particle Flow to make an object explode? (page 2–120)

How can I control the accuracy of the Particle Flow simulation? (page 2–120)

Can I use MAXScript to affect particles? (page 2–120)

Can I use more than one object as instanced geometry? (page 2–121)

How can I exclude particles from certain light sources? (page 2–121)

How do I open Particle View without clicking Particle View button on the command panel? (page 2–121)

Do I need multiple Particle Views to manage more than one particle system? (page 2–121)

Can I selectively disable or "turn off" portions of a particle system? (page 2–121)

I sometimes see events named "Action Recovery" in Particle View. Where did they come from? (page 2–121)

How do I animate action parameters? (page 2–121)

How can I synchronize an animated bitmap texture to particle age? (page 2–121)

Do I still need to bind my Particle Flow particles to space warps? (page 2–121)

How do I make particles follow a path? (page 2–122)

Can I use Snapshot or dynamics with Particle Flow? (page 2–122)

How do I use motion blur with Particle Flow? (page 2–122)

How do I use the Particle Age and Particle MBlur maps with Particle Flow? (page 2–122)

I'm trying to apply Particle Flow to an animation created with a dynamics system. Why am I getting strange results such as unexpected spawning of particles? (page 2–122)

How can I make all particles appear in the first frame while giving them different ages? (page 2–123)

How can I specify the time frame in which animated parameters are applied to particles? (page 2–123)

How can I apply bubble motion to particles? (page 2–123)

How can I branch an event to several other events? (page 2–123)

Does it matter which order I put actions in an event? (page 2–123)

How can I tell which particles are in a certain event? (page 2–124)

Why do my particles lose their material when they move to another event? (page 2–124)

Can I have an event receive input from multiple events? (page 2–124)

I changed an operator setting, but it doesn't seem to have any effect on the particle system. (page 2–124)

Can an event be isolated in Particle View and not connected to anything? (page 2–124)

General

How does Particle Flow handle time?

Parameters that measure time in Particle Flow, such as *Birth (page 2–143)* > Emit Start/Stop and *Age Test (page 2–211)* > Test Value, are specified in frames. However, Particle Flow is always aware of the current system frame rate (fps), and if you change this rate, it adjusts all time-related parameters to keep the same timing. For instance, if you set Test Value to 60 when you're using the NTSC frame rate (30 fps), and then switch to PAL (25 fps), Particle Flow automatically changes the value to 50, so the age that's tested for is still two seconds.

On the other hand, rate parameters, such as Speed, are measured in units per second, so they don't change when you go to a different frame rate.

When I go to a different frame, the software sometimes seems to freeze for a while. What's going on?

Most of the animation in Particle Flow is history dependent; that is, to be able to draw the particles in a particular frame, the program needs to know what happened in all previous frames. Normally, when you change a parameter value, the program needs to recalculate all frames between the start and the current frames. Or, if you go to a different frame, Particle Flow must recalculate one or more animation frames. If you go forward, it must calculate the frames between the current frame and the one you go to. So, for example, if you just go to the next frame, relatively little calculation is needed. But if you go backward, even only by one frame, it must calculate all frames from the start of the animation to the frame you go to.

If a lot of calculation is needed, there is a delay. Meanwhile, the software displays a message like "PF Source 01 Update xx% (Press Esc to cancel)" in the status bar, so you can get an idea of how the recalculation is progressing. If, when you see this message, you press the Esc key, Particle Flow displays an alert with the message "Click OK to turn off PF Source 01." If you click OK, the recalculation stops, giving you the opportunity to optimize the animation. For example, you could reduce the number of particles for testing purposes. You must turn the source back on to continue. If you click Cancel, the calculation continues.

Tip: An easy way to speed up recalculation and rendering of particles is by adjusting the total number of particles. To do this, select the Particle Flow source icon, and then go to the Modify panel > Emission rollout > Quantity Multiplier group. Here you can increase or decrease by percentage

the total number of particles, separately for viewport display and rendering.

Tip: To speed up animation playback when you're not adjusting parameters, use the Cache operator (see the following item).

Is there a way to pre-calculate portions of a particle simulation, as in "baking" the animation?

Yes. Particle Flow's *Cache operator (page 2–197)* lets you store all or part of a particle animation in memory, and then play back the animation from memory rather than having to recalculate particle motion. This makes it much faster to jump between different parts of the animation. You can even save the cached animation to disk as part of the scene file.

What else can I do to optimize performance?

Particle Flow can place heavy processing and resource demands on your computer. For optimal performance, the most important thing you can do is to use the fastest available CPU. Also, when using particle systems with many particles, install as much memory as possible in your computer, especially if you're using caching.

Other ways to improve performance include reducing the percentage of viewport particles with the Quantity Multiplier setting, and temporarily disabling flows and actions that you're not currently working with. When making parameter changes, return to the first frame and play forward, or set Particle View > Options menu > Update Type to "Forward". That way, if you change a setting, the particle system is not forced to recalculate its state from the very beginning. The change will affect only animation from the current frame forward. On the other hand, the result could be misleading, because you won't be able to see the difference right away.

Also, be cautious when using spawning; it can quickly create very large numbers of particles, especially when you use the By Travel Distance option in the *Spawn test (page 2–230)*.

How do I use Particle Flow to make an object explode?

Particle Flow doesn't have a fragmentation operator, but by utilizing PArray in a Birth Script operator, you can implement fragmentation in Particle Flow. You can find an example scene, with a commented script operator, in the file *scriptfragmentationwithcache.max*, on the install disc in the directory *Samples\Scenes\Version6Features\Scenes\Particle Flow\Fragments*.

How can I control the accuracy of the Particle Flow simulation?

On the command panel > *System Management rollout (page 2–139)* of the PF Source icon, you can adjust the integration step independently for viewport playback and rendering. The smaller the integration step, the more times Particle Flow calculates particle motion per frame, resulting in greater accuracy at the cost of calculation time.

Can I use MAXScript to affect particles?

Yes. Particle Flow includes a *Script operator (page 2–208)* and *test (page 2–229)*, as well as a *Birth Script operator (page 2–145)*, that let you fully customize the particle system. Each scriptable action includes a sample script, which also lists all relevant scriptable functions. You can also control parameters of the *Force operator (page 2–204)* and *Keep Apart operator (page 2–172)* with script wiring, described in the respective topics.

In addition, you can execute a script at each integration step, and another just before each frame is viewed or rendered; see *Script rollout (page 2–139)*.

Can I use more than one object as instanced geometry?

Yes. With the *Shape Instance operator (page 2–178)*, you can use groups, hierarchies, and objects consisting of multiple elements, with each members of the combined object constituting a separate particle. These objects can be emitted in a specific order, or in a random order. For example, you can use a text object, with the letters emitted in the order in which they appear in the text.

How can I exclude particles from certain light sources?

To prevent particles from being affected by a light source, use the light's Exclude function, found on the General Parameters rollout (Modify panel), to specify any events containing particles to be excluded. Specifying the PF Source XX object (default name) has no effect; you must specify all objects listed as PF Source XX->Event XX.

Particle View

How do I open Particle View without clicking Particle View button on the command panel?

 Press the 6 key. It might be necessary to first turn on the Keyboard Shortcut Override Toggle.

Do I need multiple Particle Views to manage more than one particle system?

No: All particle systems appear in Particle View. You can scroll and zoom the event display to see the different systems.

Can I selectively disable or "turn off" portions of a particle system?

Yes. Particle View gives you a number of different ways of doing this. You can click an action's icon or an event's light-bulb icon to turn it off, or use the right-click menu, or use Edit menu's Turn On and Turn Off commands. Also, if you press Esc while Particle Flow is calculating, the software gives you the opportunity of turning off the entire particle system, thus immediately returning control of the software to you. You can then analyze the system to determine the area of slowdown, optimize or simplify the particle flow, and then recalculate the animation.

I sometimes see events named "Action Recovery" in Particle View. Where did they come from?

When you merge a scene containing Particle Flow data, you can import entire events and isolated actions. If you merge an action without its events, the software places it in an event named "Action Recovery."

Animation

How do I animate action parameters?

Use the same *Auto Key (page 2–278)* method as with animating any other parameter in 3ds Max.

Note: Avoid animating Particle Flow parameters with Set Key mode or with Shift+right-click. Also, to see animation keys for actions in the track bar, make sure the Particle Flow source icon is selected.

How can I synchronize an animated bitmap texture to particle age?

Use the *Material Dynamic operator (page 2–191)* in conjunction with the *Bitmap 2D map (page 2–1631)*. You'll find a procedure that describes the method in detail *here (page 2–191)*.

Do I still need to bind my Particle Flow particles to space warps?

No; you use the *Collision test (page 2–212)*, *Collision Spawn test (page 2–215)*, or *Force operator (page 2–204)* to affect particle motion with 3ds Max space warps. The ability to do this on a

global and local (per-event) basis gives you much greater control over how space warps affect the particles than with previous systems.

How do I make particles follow a path?

You can use the *Speed By Icon operator (page 2–162)* or the *Find Target test (page 2–218)*, assigning the icon or target to a spline path. You'll find illustrative procedures in both topics.

Can I use Snapshot or dynamics with Particle Flow?

No. The Snapshot tool is not currently compatible with Particle Flow, nor is **reactor** or the legacy dynamics system.

How do I use motion blur with Particle Flow?

You can apply motion blur on a per-event basis by editing the event's *object properties (page 1–117)* with the right-click menu in Particle View, or on a global basis by editing the global event's object properties. In the Particle View dialog, highlight the event to edit and then right-click and choose Properties. On the Object Properties dialog, edit the Motion Blur group settings.

Alternatively, you can apply motion blur to an entire Particle Flow system without the need to set properties by using the *multi-pass motion blur effect (page 2–1386)* with a camera.

Note: You can also use Object Properties to set other properties for events and the particle system, as with any object in 3ds Max. Do this *only within Particle View*, not using the source icon in the viewports (right-click menu > Properties or Edit menu > Object Properties).

Note: When using Object Motion Blur, if an event contains a Material Dynamic operator that uses a material with a Particle Age, Particle MBlur, or Bitmap map, the event should not also contain a Delete operator, or a Spawn or Collision Spawn test. Also, the event should not contain any tests that are wired to another event. The only exception to this is the Age Test operator when set to Absolute Time without any variation; that is, all particles leave the event at the same time. This applies to the use of Object Motion Blur only; there are no restrictions with Image Motion Blur.

How do I use the Particle Age and Particle MBlur maps with Particle Flow?

To use the Particle Age map with Particle Flow, you must use the Delete operator to give the particles a finite life span. For further information, see *Delete Operator (page 2–146)*.

Also note the following:

- Particle Age, Particle MBlur and a Bitmap map containing an animated bitmap can be used only with *Material Dynamic operator (page 2–191)*, with Same As Particle ID turned on. You'll find procedures for using Bitmap and Particle Age at this topic.
- The particle material cannot be a submaterial; it must be the main material.
- In the non-event-driven particle systems in 3ds Max, Particle MBlur is used with a control named Direction Of Travel\MBlur and an accompanying Stretch parameter. In Particle Flow, you can replicate the stretching effect by using a Scale operator, turning off Scale Factor > Constrain Properties, and scaling the particle along one axis.

I'm trying to apply Particle Flow to an animation created with a dynamics system. Why am I getting strange results such as unexpected spawning of particles?

The dynamics system probably generates rotation keys using the Euler XYZ controller. To avoid interpolation discontinuities, change the rotation controller for such objects to TCB Rotation.

How can I make all particles appear in the first frame while giving them different ages?

Use a negative frame range in the *Birth operator (page 2–143)*. For example, to get a particle-age spread of 30 frames, set Emit Start to **-29** and Emit Stop to **0**.

How can I specify the time frame in which animated parameters are applied to particles?

You can animate many of the Particle Flow parameter values with keyframing. In most actions, you can choose the time frame by which to apply this animation to the particles from a drop-down list labeled Sync By. You can apply this animation to particles in the time frame of the entire animation, or at a specific time of each particle's life (particle age), or based on the length of time the particle has been in the current event. See the individual operator and test topics for details.

How can I apply bubble motion to particles?

Although Particle Flow doesn't have the *bubble motion option (page 2–270)* found in PArray, you can simulate the effect by following this procedure:

1. Add an object to serve as particle geometry and a *dummy helper object (page 2–16)*.
2. Position the dummy away from the center of the particle geometry object and *link (page 2–421)* your particle geometry to the dummy so that the dummy is the parent object (that is, drag from the geometry object to the dummy). The distance between the *pivots (page 2–487)* of the two objects determines the radius of the bubble motion.
3. In Particle View, add a *Shape Instance operator (page 2–178)* and a *Spin operator (page 2–154)* to the event in which the bubble motion is to occur.
4. Click the Shape Instance operator and specify the dummy object as the particle geometry object.
5. Click the Spin operator and set appropriate values for Spin Rate and Variation. For Spin Axis, keep the default setting of Random 3D.

The result is particle motion in which the particles bob about while moving along their paths, similar to bubble motion.

Events

How can I branch an event to several other events?

Particle Flow includes several Split tests, which let you send some particles to another event based on quantity, selection, or source. You can use any number of these in a single event to send parts of the particle stream to different events, and then use a *Send Out test (page 2–230)* to redirect the remaining particles to another event. Keep in mind that any tests subsequent to the first can work only with particles remaining in the event, not necessarily all particles that begin in the event.

Does it matter which order I put actions in an event?

It depends. If an event contains two or more operators of the same type, such as Shape, the last one overrides the rest. If an event contains two or more tests, they are evaluated and particles redirected in the order in which the tests appear.

However, if two actions in an event control the same properties, the interaction is more complex. An action that works on a continuous basis will typically prevail over one that affects particles only when they first enter the event. For example, both the *Find Target test (page 2–218)* and *Speed operator (page 2–159)* can control particle speed and direction, but the test works continuously, while the operator takes effect only once. If an event contains both, particle speed and direction

will be primarily controlled by the test, in general, even if the operator comes after the test. However, the operator's settings will still have some influence over particle behavior, particularly if its Speed value is significantly higher than that of the test. For a list of actions' effective time frames, see *Action Time Frames (page 2–141)*.

Also, if you're testing for a specific condition that can be affected by other actions in an event, be sure to place the test *after* the actions. For example, in an event with a *Force operator (page 2–204)* and a *Collision test (page 2–212)*, place the Collision test after the Force operator. This avoids the possibility of the force pushing particles past the deflector before the software can test for the collision, which would allow the particles to penetrate the deflector. In general, place tests at the end of the event.

How can I tell which particles are in a certain event?

One way is to set the Type option for the *Display operator (page 2–202)* in each event to a different choice. For example, the first event could use Ticks, the second Circles, and the third Lines. This way the particles change appearance in the viewports as they move from event to event. You can also use the Display operator to change particle colors, to further distinguish them.

Another way is to select all particles in a certain event. Select the Particle Flow source icon, and then go to the Modify panel > Selection rollout and click the Event icon. You can then click an event in the Select By Event list to highlight all of its particles in the viewports.

Why do my particles lose their material when they move to another event?

A material is a static property of an event. It does not travel along with the particles from event to event. A particle's material ID does, but its material does not. If you want particles always to use the same material, define the material in the *global event (page 3–949)* with a *Material operator (page 2–186)* or a *Shape Instance operator (page 2–178)*. Otherwise, you need to define it in each local event.

Can I have an event receive input from multiple events?

Particle Flow lets you wire any number of tests to a single event.

I changed an operator setting, but it doesn't seem to have any effect on the particle system.

A similar operator in the global event might be overriding your local operator. By default, Particle Flow evaluates local operators first, and then global operators. If a global operator affects the same property, such as speed, as a local one, the particle system will use the value set by the global operator. You can set local operators to override global ones by choosing Particle View > Options menu > Action Order > Globals First.

Can an event be isolated in Particle View and not connected to anything?

Yes, but it won't affect the particle system at all.

Particle Flow User Interface

Particle View

Select a Particle Flow source icon > Modify panel > Setup rollout > Click Particle View (or press 6).

Create panel > Geometry > Particle Systems > Object Type rollout > Click PF Source. > Setup rollout > Click Particle View (or press 6).

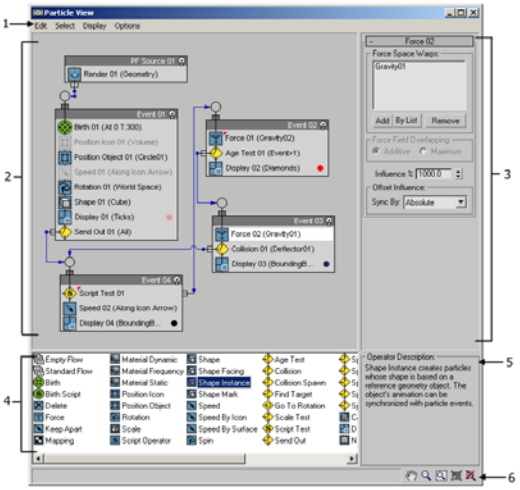
Particle View provides the main user interface for creating and modifying particle systems in Particle Flow. The main window, known as the *event display*, contains the *particle diagram*, which describes the particle system. A particle system consists of one or more *events* wired together, each of which contains a list of one or more *operators (page 2–142)* and *tests (page 2–210)*. Operators and tests are known collectively as *actions*.

The first event is called the *global event*, because any operators it contains can affect the entire particle system. The global event always has the same name as the Particle Flow icon; by default, this is PF Source ## (starting with 01 and counting upward). Following this is the *birth event*, which must contain a *Birth operator (page 2–143)* if the system is to generate particles. By default, the birth event contains this operator as well as several others that define the system's initial properties. You can add any number of subsequent events to a particle system; collectively, the birth event and additional events are called *local events*. They're called this because a local event's actions typically affect only particles currently in the event.

You use tests to determine when particles are eligible to leave the current event and enter a different one. To indicate where they should go next, you *wire* the test to another event. This wiring defines the schematic, or flow, of the particle system.

By default, the name of each operator and test in an event is followed by its most important setting or settings in parentheses. Above the event display is a menu bar, and below is the *depot*, containing all actions available for use in the particle system, as well as a selection of default particle systems.

Tip: The easiest way to open Particle View is by pressing the 6 key. It's not necessary to first select a Particle Flow icon.



1. Menu bar
2. Event display
3. Parameters panel
4. Depot
5. Description panel
6. Display tools

Particle View comprises the following elements:

- The *menu bar (page 2–126)* provides functions for editing, selection, adjusting the view, and analyzing the particle system.
- The *event display (page 2–131)* contains the particle diagram, and provides functions for modifying the particle system.

- The *parameters panel* contains rollouts for viewing and editing parameters of any selected actions. Basic functionality is identical to that of rollouts on the 3ds Max command panels, including usage of the right-click menu.

  To toggle display of the parameters panel, choose Display menu > Parameters.

- The *depot* contains all Particle Flow actions, as well as several default particle systems. To see an item's description, click its entry in the depot. To use an item, drag it into the *event display (page 2–131)*.

  The contents of the depot fall into three categories: *operators (page 2–142)*, *tests (page 2–210)*, and *flows (page 2–208)*.

  To toggle display of the depot, choose Display menu > Depot.

- The *Description panel* displays a brief description of the highlighted depot item.

  To toggle display of the description panel, choose Display menu > Description.

- The icon-based *display tools* in the bottom-right corner of the Particle View dialog let you pan and zoom the event display window. For descriptions, see *Display menu (page 2–129)*.

Particle View Menu Bar

Select a Particle Flow source icon. > Modify panel > Setup rollout > Click Particle View (or press 6).

Create panel > Geometry > Particle Systems > Object Type rollout > Click PF Source. > Setup rollout > Click Particle View (or press 6).

The menu bar provides access to a number of important Particle View functions.

Procedure

To render only particles in specific events:

1. In Particle View, highlight the events containing the particles you want to render.
2. Choose Select menu > Assign Selection To Viewport.
3. Render using one of the *Render Type (page 3–13)* > Selected options.

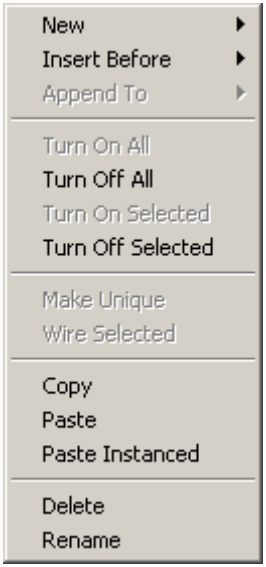
To bring a Particle Flow setup into a different scene:

1. In Particle View, highlight the events containing the particles you want to merge with a different scene. For example, to designate only a particular flow, you might highlight its global event, and then choose Select menu > Select Downstreams.
2. Choose Select menu > Save Selected, and use the file dialog to specify a file to save.

   Note: Any scene objects on which Particle Flow depends are also saved. For instance, if a *Shape Instance operator (page 2–178)* in one of the highlighted events refers to a certain object, that object is saved along with the designated events.
3. Open or create the scene with which the saved Particle Flow setup will be combined, and then use *Merge (page 3–463)* to merge the file from step 2.

Interface

Edit menu



Each of the first three commands on this menu presents a submenu containing all actions. Choose the command, and then choose an action from the submenu.

Note: The Undo and Redo commands are available from the main 3ds Max Edit menu, and their default keyboard shortcuts are the same: Ctrl+Z and Ctrl+Y, respectively.

New—Adds a new event containing the chosen action to the event display.

Insert Before—Inserts the chosen item above each highlighted action. Available only when one or more actions are highlighted.

Append To—Inserts the chosen item at then end of each highlighted event. Available only when one or more events are highlighted.

Turn On All—Turns on all actions and events.

Turn Off All—Turns off all actions and events.

Turn On Selected—Turns on any highlighted, turned-off actions or events. Available only when one or more highlighted items are turned off.

Turn Off Selected—Turns off any highlighted, turned-on actions or events. Available only when one or more highlighted items are turned on.

Make Unique—Converts an instanced action to a copy that's unique to its event. Available only when one or more instanced actions are highlighted.

Wire Selected—Wires one or more highlighted tests to a highlighted event, or one or more highlighted global events to a highlighted birth event. Available only when one or more tests and a single event are highlighted, or when one or more global events and a single birth event are highlighted.

Copy—Copies any highlighted events, actions, and wires to the paste buffer. Keyboard equivalent: Ctrl+C.

Paste—Pastes the contents of the paste buffer to the event display. Keyboard equivalent: Ctrl+V.

The result of copying and then pasting multiple items depends on what you copied:

- If you copied multiple actions, pasting adds them all to a single new event, even if they originally came from different events.
- Copying multiple events or wires pastes them exactly as copied.
- If you copied a combination of actions and events, pasting creates a new event for each group of actions from a single event. For example, if you highlight event A, some actions from event B, and some actions from event C, copying and then pasting would add three events: a copy of event A, a second event with the actions from event B, and a third event with the actions from event C.

Paste Instanced—Pastes the contents of the paste buffer to the event display, making instances of any pasted actions and their originals.

For the results of pasting multiple copied items, see Paste, above.

Delete—Deletes any highlighted items. Keyboard equivalent: Delete.

Deleting an event also deletes all of its actions.

Rename—Lets you enter a new name for any single highlighted item in the event display. Available only when one action or test is highlighted

Select menu



By default, the Select tool is active when Particle View is open, as indicated by the arrow-shaped mouse cursor. You can use this tool to highlight, move, and copy events, actions, tools, and wires. You can also use the commands on this menu to highlight all elements, no elements, or elements by category.

Note: When you render using any of the Selected options, the software renders only events that are selected in the viewports. You can use the last three commands on this menu to manage and coordinate selection of events and flows between Particle View and the viewports. This also applies to use of the Views menu > Shade Selected command.

Select Tool—Activates the Select tool. Choose this to return to the Select tool after using interactive tools for panning and zooming the event display. Alternatively, simply right-click anywhere in the event display to activate the Select tool.

Select All—Highlights all items in the event display.

Select None—Deselects all items in the event display. Alternatively, click an empty area in the event display.

Select Actions—Highlights all operators and tests in the event display.

Select Operators—Highlights all operators in the event display.

Select Tests—Highlights all tests in the event display.

Select Events—Highlights all events in the event display.

Select Wires—Highlights all wires in the event display.

Select Downstreams—Highlights all events after currently highlighted events. Available only when one or more events are highlighted.

Save Selected—Saves only highlighted elements in the event display to a *.max* file. You can then open this file, or combine it with an existing scene with File menu > Merge.

Note: The Particle View > Save Selected command saves all items selected in Particle View *and* in the scene. So if a PF Source icon is selected in the scene, then this command will save its global event as well as any birth events associated with it. if you want to save only parts of a Particle Flow setup, first unselect everything in the scene.

Also, if an action references an object (for example, a *Shape Instance operator (page 2–178)* uses a geometry object), and the action or its parent event is highlighted when you use Save Selected, then the reference object is saved as well.

Get Selection From Viewport—Highlights global events whose source icons are selected in the viewports.

Assign Selection To Viewport—Transfers an event selection to the viewports. Use this to render only particles in specific events. First, in Particle View, highlight the events to render. Next, choose this command, and finally, render using one of the *Render Type (page 3–13)* > Selected options.

Sync Source/Events Selection In Viewport—Selects all events of any source icons selected in the viewports. You can then propagate this selection to Particle View with the Get Selection From Viewport command.

Use this function to render a specific particle flow using one of the Render > Selected options.

Display menu



The first five commands on this menu are also available as icons, in the Display tools section on the right side of the bottom border of the Particle View dialog. Each command's Display tools icon is shown below.


Pan Tool—Drag in the event display to move the view. The mouse cursor changes to a hand icon. You can also pan the view by dragging with the middle mouse button or wheel button held down.

To exit this mode, right-click in the event view or choose the command again.


Zoom Tool—Drag in the event display to zoom the view. The mouse cursor changes to a magnifying-glass icon. Drag upward to zoom in, or downward to zoom out.

To exit this mode, right-click in the event view or choose the command again.


Zoom Region Tool—Drag in the event display to define a zoom rectangle. The mouse cursor changes to the image of a magnifying glass inside a zoom region. When you release the mouse button, the display zooms to show only the area defined by the region.

To exit this mode, right-click in the event view or choose the command again.


Zoom Extents—Sets the zoom to show the entire particle diagram in the event display.


No Zoom—Sets the zoom to the default level. This is the level displayed when you first open Particle View in a given session.

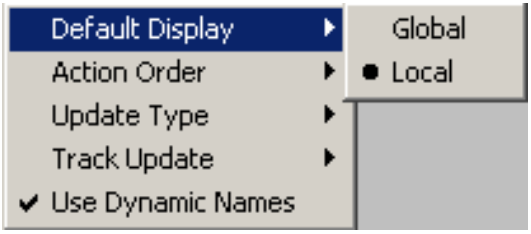
Parameters—Toggles display of the parameters panel, on the right side of the Particle View dialog. Default=on.

Depot—Toggles display of the depot, below the Particle View dialog. Default=on.

Description—Toggles the Description panel, to the right of the depot. Default=on.

The Description panel displays a brief description of any action highlighted in the depot.

Options menu



This menu contains a number of options useful in analyzing and debugging particle systems.

Default Display—Determines whether the Display operator is applied locally or globally to new particle systems and events. Default=Local.

- **Global—**When you create a new particle system, Particle Flow adds a single Display operator to the global event. It does not add a Display operator to each new event.
- **Local—**Particle Flow adds a Display operator to each new event. This lets you easily distinguish particles in different events in the viewports.

Action Order—For predictable results, don't use similar actions in the global and local events. However, if comparable actions do exist in both the global event and other events, Particle Flow applies them to the system in the order specified here. Typically, the effect that's applied *last* is the one that's visible in the particle system. Default=Locals First.

For example, consider a situation in which the global event contains a Shape operator set to Cube, and the other events in the same system contain Shape operators set to Tetra or Sphere. By default, Particle Flow applies the global Shape operator last at each integration step, overwriting any local shapes, so all particles would be cubes. But if you set Action Order to Globals First, the local shapes are applied last, and particles appear as tetrahedrons or spheres, depending on which event they're in.

- **Globals First—**At each integration step, Particle Flow first applies actions in the global event, and then actions in the other (local) events. Typically, the result is that actions in the local events override comparable actions in the global event.
- **Locals First—**At each integration step, Particle Flow first applies actions in the local events, and then actions in the global event. Typically, the result is that actions in the global event override comparable actions in the local events.

Update Type—This setting determines how Particle Flow updates the system when you change a parameter during playback. Because the state of the particle system at any given frame can depend on events in previous frames, using the Complete option can provide a more accurate depiction of the result of the change, at a cost in speed; it takes more time to recalculate the entire system starting at the first frame. Default=Complete.

- **Complete—**When you change a setting during playback, Particle Flow updates the entire system, starting at the first frame.
- **Forward—**When you change a setting during playback, Particle Flow updates the system starting at the current frame.

Track Update—Provides options for visualizing the particle system status in Particle View.

- **Particle Count—**Adds a tab above each event that shows the number of particles in the event. The global event's count shows the total number of particles in the particle system.

  Use this, in addition to the ability to display particles differently in each event, to trace particle progress through the system.
- **Update Progress—**Highlights each action in color whenever Particle Flow evaluates it.

  The highlighting is very fast, but this option can add significant computational overhead to the particle system, with the result that

real-time playback might skip a greater number of frames. To see all frames, turn off Time Configuration > Real Time.

Use Dynamic Names—When on, action names in events are followed by their most important setting or settings, in parentheses. When off, only the names appear. Default=on.

Particle View Event Display

Select a Particle Flow source icon. > Modify panel > Setup rollout > Click Particle View (or press 6).

Create panel > Geometry > Particle Systems > Object Type rollout > Click PF Source. > Setup rollout > Click Particle View (or press 6).

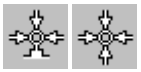
The *event display (page 3–936)* in Particle View contains the particle diagram, and provides a complete range of intuitive functions for editing the particle system.

Note: A number of these functions are also available from the *event display right-click menu (page 2–133)*, as well as the *menu bar (page 2–126)*.

Editing the Particle Diagram and Particle System

Interactive functionality in event display includes the following:

General

- To add an action to the particle system, drag it from the depot to the event display. If you drop it in an empty area of the event display, it creates a new event. If you drag it to an existing event, the result depends on whether a red line or a blue line appears when you drop it. If a red line, the new action replaces the underlying action. If a blue line, the action is inserted in the list.
- To edit an action's parameters, click its name in an event. The parameters appear on the right side of Particle View. If not, it means the parameters panel is hidden; to display it, choose Display menu > Parameters.
-  To wire a test to an event, drag from its *test output*, the blue dot that by default sticks out to the left of the test, to the event's *event input*, which sticks out from the top, or vice-versa. Similarly, you can wire a global event to a birth event by dragging between the *source output* on the bottom of the global event and the event input. The mouse cursor changes to the first depicted image when you can begin this operation, and to the second image when you can complete it.
- To pan or zoom the event display, use the controls available from the *Display menu (page 2–129)*, the *right-click menu (page 2–133)*, or as icons at the bottom-right of the dialog.

Modifying Actions and Events in the Event Display

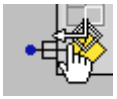
- To change the color used by a Display operator, click its color swatch to the right of its name.
-  To toggle an operator (that is, turn it off and on), click its icon in the event. The mouse cursor changes to the depicted image when this action is possible.

   A turned-off operator does not affect the particle system, and appears as gray in the event list; however, its parameters are still available for editing.

  Toggling an instanced operator doesn't affect its other instances.
- 

To toggle an event (that is, turn it off and on), click the light-bulb icon next to its name in the title bar. Turning off an event is equivalent to turning off all of its

operators and tests, although you cannot turn the operators and tests back on individually. The event's items do not affect the particle system, and appear as gray in the event list. All particles entering a turned-off event are immediately passed on to the next event, if any.

-  To make a test result always True or False, click its icon's left or right side, respectively. The mouse cursor changes to the first image at the left (with the left-pointing arrow) when you can click for always True, and to the second image at the left (without the left-pointing arrow) when you can click for always False.

   The test icon changes to a green or red light bulb to indicate current functionality: green for always True, and red for always False.

   To return to the original functionality, click the icon again. The mouse cursor changes to the depicted image when this action is possible.

Selecting, Moving, and Copying

- To copy an action or event, first press and hold the Shift key and position the mouse cursor over the item to copy. When the plus sign (+) appears next to the mouse cursor arrow, drag the item to a new location. When you release the mouse button, the Clone Options dialog appears; choose Copy or Instance and click OK. In the Clone Options dialog, the Event And Particle System group is unavailable, as a reminder that you cannot instance these elements. Reference is not an option when cloning Particle System elements. Cloning an event or particle system simply creates a new event or particle system containing clones of the original's actions.

  Any instanced actions have identical properties. Changing one instance's values changes them all. On the other hand, copied actions can have different values.

  To clone multiple items, highlight them all before Shift+dragging. If you attempt to clone a combination of actions and events, or actions from different events, the result depends on where you start dragging from. If you drag a highlighted action, you'll clone only actions from that event. If you drag an event, only highlighted events are cloned.

  Note: When you highlight an instanced action, Particle View indicates all of its other instances by changing their background color to a lighter shade of gray.

- To move an event, drag its title bar, or the icon of any of its actions. If you drag the event to an edge of the display, and scrolling is possible, the display window will automatically scroll in that direction.
- To resize an event, drag its right edge.
- To highlight multiple items, drag a rectangle around them, or Ctrl+click them individually.
- To delete an event, action, or wire, highlight it and then press Delete. Be sure the Particle Dialog is active, or you might inadvertently delete a selection in the scene instead.
- To move an action, drag its *name* (not its icon) to the new location. If you drag the event to an edge of the display, and scrolling is possible, the display window will automatically scroll in that direction.

  If you drag an action to an event, you can insert the action or replace an existing action, depending on where you position the mouse cursor before releasing the button. A blue line (before releasing) indicates that it will be inserted at that location, and a red line indicates that it will replace the underlying action. If you

drag the action to an empty space in the event display, it will form a new event, and, by default, the system will add a new Display operator.

Note: Particle Flow doesn't let you perform illegal operations. For instance, a Birth operator must always be the first item in a *birth event (page 3–916)*; the system doesn't let you position it elsewhere. You can place a Birth operator in an isolated event, but you can't then wire that event into a particle system, except in parallel with another birth event. For more information, see the *Birth operator topic (page 2–143)*.

-  To move a test output or switch it to the opposite side of the event, drag from just inside the test output to the desired new location. The mouse cursor changes to the depicted image when this operation is possible.

Getting Information

- To see a renamed action's original type, position the mouse cursor over its event list entry until the tooltip appears.
- To see a brief description of an action, click it in the depot. The description appears on the Description panel, to the right of the depot.

Event Display Right-Click Menu

Select a Particle Flow source icon. > Modify panel > Setup rollout > Click Particle View (or press 6). > Right-click in event display.

Create panel > Geometry > Particle Systems > Object Type rollout > Click PF Source. > Setup rollout > Click Particle View (or press 6). > Right-click in event display.

The event display right-click menu provides handy access to a number of contextual commands. The contents of the menu, as well as the results of its commands, depend in some cases on where you click, in other cases on what's highlighted, and in still other cases, on both.

In the following descriptions, "item" means an action or event, or, in some cases, a wire.

Interface

Insert—Displays a submenu of all actions, and when you choose a new action, inserts it where you originally right-clicked. Available only when you right-click over an event or action.

If you right-click over an action, and then insert an action, the inserted action replaces the old one. If you right-click above or below an action (no tooltip appears), the new action is added to the event at that location. If you right-click over an event, the action is inserted at the first available position.

Append—Displays a submenu of all actions, and when you choose a new action, inserts it at the end of the event. Available only when you right-click over an event or action.

Turn On—Turns on the event or action under the mouse cursor. Available only when the item under the mouse cursor is turned on.

To affect an event, the mouse cursor must be over its title bar.

Turning on a highlighted item affects all highlighted items.

Turn Off—Turns off the event or action under the mouse cursor. Available only when the item under the mouse cursor is turned off.

To affect an event, the mouse cursor must be over its title bar.

Turning off a highlighted item affects all highlighted items.

Make Unique—Converts an instanced action to a copy that's unique to its event. Available only when the mouse cursor is over an instanced action.

If you right-click over a highlighted, instanced action, Make Unique affects all such actions.

Wire—Wires one or more highlighted tests to a highlighted event, or one or more highlighted global events to a highlighted birth event. Available only when one or more tests and a single event are highlighted, or when one or more global events and a single birth event are highlighted, and when you right-click over an eligible, highlighted item.

Copy—Copies any highlighted events, actions, and wires to the paste buffer. Available only when the mouse cursor is over an highlighted action, but affects all highlighted items.

Paste—Pastes the contents of the paste buffer to the event display at the mouse cursor.

If the cursor is over a list, and you copied actions, the actions are added to the event. If the cursor is over an empty area of the event display, the software creates a new event containing the pasted actions. A copied event can be pasted only in an empty area of the event display.

The result of copying and then pasting multiple items depend on what you copied and the position of the mouse cursor:

- If you copied multiple actions, pasting adds them all to the event under the mouse cursor, or, if the cursor is over an empty area, to a new event, even if they originally came from different events.
- Copying multiple events and/or wires pastes them exactly as copied, but in this case Paste is available only when the mouse cursor is over an empty area.
- If you copied a combination of actions and events, you can paste them only when the mouse cursor is over an empty area. Pasting creates a new event for each group of actions from a single event. For example, if you highlight event A, some actions from event B, and some actions from event C, copying and then pasting would add three events: a copy of event A, a second event with the actions from event B, and a third event with the actions from event C.

Paste Instanced—Pastes the contents of the paste buffer to the event display, making instances of any pasted actions and their originals.

For the results of pasting multiple copied items, and the limitations based on mouse position, see Paste, above.

New—Add a new item to the event display from the submenu. The submenu contents are identical to that of the depot. Available only when the mouse cursor is over an empty area.

Delete—Deletes the item under the mouse cursor. Available only when the mouse cursor is over an item.

If the mouse cursor is over a highlighted item, choosing Delete deletes all highlighted items.

Rename—Lets you rename the item under the mouse cursor. Enter a new name from the keyboard.

Properties—Opens the Object Properties dialog for the event under the mouse cursor. Available only when the mouse cursor is over the title bar of a highlighted event.

Particle Flow lets you set object properties on a per-event basis. This lets you set attributes such as Hide, Renderable, and Motion Blur separately for each event. Alternatively, you can set properties for the entire particle system from the global event.

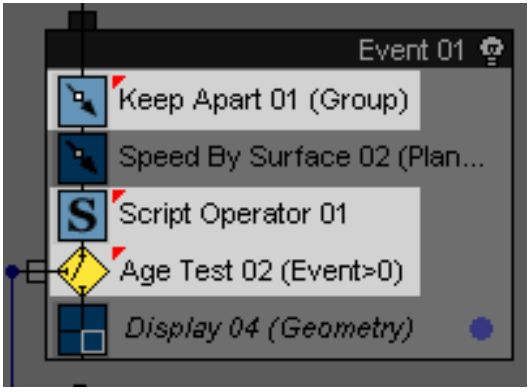
For more information, see *Object Properties (page 1–117)*.

Note: Set properties for events and the particle system *only within Particle View*, not via the source icon in the viewports (using the right-click menu > Properties or Edit menu > Object Properties).

Use Script Wiring—Lets you use a script to specify certain parameters in the *Force operator (page 2–204)* and *Keep Apart operator (page 2–172)*. This command appears only when you right-click either of the specified, highlighted operators in an event.

Comments—Lets you add comments to each action and event, and view existing comments. Available only when the mouse cursor is over an action or the title bar of an event.

Choosing Comments opens a window for entering text from the keyboard. Begin typing, and click OK to finish. A commented event or action has a red triangle near its name to indicate the presence of the comment. You can click this triangle to open the comment.



Commented actions; click the triangle to view a comment.

Note: A comment is specific to the action or event to which it's applied. If you copy a commented event or action, the comment is not included in the copy, because it's probably not applicable to the copy.

Pan—Switches to the Pan tool. Drag in the event display to change the view; right-click to exit. Available only when you click over an empty area of the event display.

Zoom—Switches to the Zoom tool. Drag upward in the event display to zoom in, and downward to zoom out; right-click to exit. Available only when you click over an empty area of the event display.

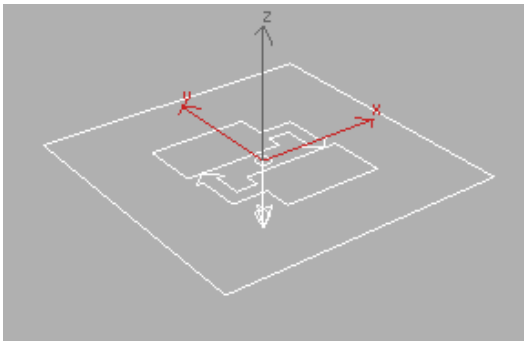
Zoom Region—Switches to the Zoom Region tool. Drag a rectangle in the event display to specify the area to zoom into; right-click to exit. Available only when you click over an empty area of the event display.

Particle Flow Source

Add a Particle Flow system to the scene. > Create panel

Select a PF Source icon. > Modify panel

Particle View (page 2–125) > Click a global event (title bar). > Parameters panel



The Particle Flow source is the viewport icon for each *flow (page 3–942)*, and also serves as the default emitter. By default, it appears as a rectangle with a central logo (see above illustration), but you can change its shape and appearance using the controls described below.

When you select a source icon in the viewport, the Particle Flow emitter-level rollouts appear on the Modify panel. Alternatively, click the title bar of a global event in Particle View to highlight it, and to access the emitter-level rollouts from the parameters panel on the right side of the Particle View dialog. Use these controls for setting global attributes, such as icon properties and the maximum amount of particles in the flow.

The particle source icon is roughly equivalent to the corresponding *global event (page 3–949)* in Particle View. They have the same name, but selecting one does not select the other. If you delete a particle source icon from the scene, the software converts the global event to an isolated local event in Particle View, retaining its operators with their settings intact. Any other events in the system remain in Particle View, along with their wiring. However, if you delete a global event, the software also removes any local events used exclusively by that system, as well as the corresponding particle source icon. To retain the local events, first delete the wire from the global event, and then delete the global event.

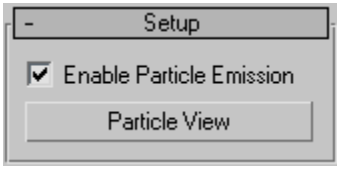
If you clone a particle source in a viewport with Shift+transform or Edit menu > Clone, an equal number of copies of the global event appear in Particle View, each wired to the original *birth event (page 3–916)*. The Clone Options dialog offers only the Copy option. However, if you clone a global event in Particle View, the Clone Options dialog also lets you create instances of the cloned operators and tests. It's not possible to create instances of global and local events, so these options are unavailable in the Clone Options dialog, as a reminder. Also, global events cloned in Particle View are not automatically wired to the original birth event.

Interface

Modifier stack

In the modifier stack, expanding the hierarchy of a Particle Flow source object provides access to two sub-object levels: Particle and Event. For further information, see *Selection rollout (page 2–138)*.

Setup rollout



Use these controls to turn the particle system on and off, and to open Particle View.

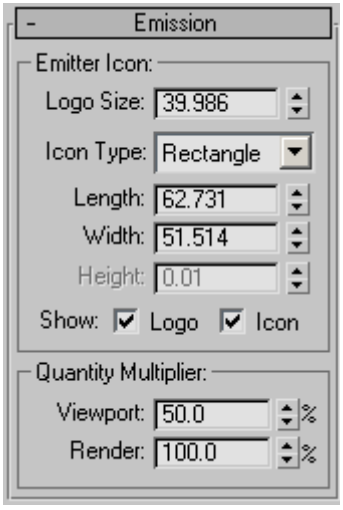
Note: This rollout appears only on the Create and Modify panels, not on the Particle View dialog > parameters panel.

Enable Particle Emission—Turns the particle system on and off. Default=on.

You can also turn off all particle flows in Particle View with Edit menu > Turn Off All, or a specific particle flow by right-clicking its global event's title bar and choosing Turn Off.

Particle View—Click to open the *Particle View dialog (page 2–125)*.

Emission rollout



Sets the physical characteristics of the emitter (particle source) icon, and the percentage of

particles produced in the viewports and when rendering.

Emitter Icon group

Logo Size—Sets the size of the Particle Flow logo, which appears at the center of the source icon, as well as the arrow that indicates the default direction of particle motion.

By default, the logo size is proportional to that of the source icon; with this control, you can make it larger or smaller.

This setting affects only the viewport display of the logo; changing it has no effect on the particle system.

Icon Type—Choose the basic geometry of the source icon: Rectangle, Box, Circle, or Sphere. Default=Rectangle.

This choice matters only if you use the source icon as the particle emitter. The available size settings depend on which icon type you choose, and, again, are important only if you use the source icon as an emitter.

The default icon type is Rectangle. If you add a particle system, and then change the icon type to Box, the icon continues to resemble a rectangle. To make it look like a box, increase the Height setting.

Length/Diameter—Sets the length of the Rectangle and Box icon types, and the diameter of the Circle and Sphere icon types.

Width—Sets the width of the Rectangle and Box icon types. Unavailable with the Circle and Sphere icon types.

Height—Sets the height of the Box icon type. Available only with the Box icon type.

Show Logo/Icon—Turns display of the logo (with arrow) and icon on and off, respectively.

These settings affect only the viewport display of these items; they have no effect on the particle system.

Quantity Multiplier group

These settings determine the percentage of the total number of particles in each *flow (page 3–942)* that are actually produced in the viewports and at render time. They don't affect the percentage of particles that are visible; those are determined by the *Display operator (page 2–202)* and *Render operator (page 2–206)*. You can use them to quickly decrease or increase the number of particles consistently throughout all events in the particle system. The maximum setting, 10,000%, lets you multiply the number of particles generated by the flow by 100.

The total number of particles is determined by the combined effects of the following operators and tests: *Birth (page 2–143)*, *Birth Script (page 2–145)*, *Delete (page 2–146)*, *Collision Spawn (page 2–215)*, and *Spawn (page 2–230)*. Scripted operators and tests can also affect this number.

Viewport—Sets the percentage of the total number of particles in the system produced in the viewports. Default=50.0. Range=0.0 to 10000.0.

Render—Sets the percentage of the total number of particles in the system produced at render time. Default=100.0. Range=0.0 to 10000.0.

Selection rollout



Use these controls for selecting particles on a per-particle or event basis. Selection of particles at the Event level is for debugging and tracking purposes. Particles selected at the *Particle level* can be acted upon by the *Delete operator (page 2–146)* and the *Split Selected test (page 2–235)*. You cannot directly manipulate selected particles with standard 3ds Max tools such as Move and Rotate.

Note: This rollout appears only on the Modify panel, not on the Create panel or Particle View dialog > parameters panel.

Particle—Lets you select particles by clicking them or dragging a region.

Event—Lets you select particles by event. At this level, you can select all particles in one or more events by highlighting the event(s) in the *Select By Event list*, or in the viewports with standard selection methods. To convert a selection from the Event level to the Particle level for use with the Delete operator or Split Selected test, use *Get From Event Level*.

Selected particles appear in the viewports in red (if not geometry), in the form designated by the *Display operator > Selected setting*.

Select by Particle ID group

Each particle has a unique ID number, starting with 1 for the first particle and counting up. Use these controls to select and deselect particles by their ID numbers. Available only at the Particle selection level.

Tip: You can display particle IDs in the viewports by turning on Display operator > Show Particle IDs.

ID—Use this to set the ID number of the particle you want to select. You can set only one number at a time.

Add—After setting the ID number of a particle to select, click Add to add it to the selection. By default, selecting a particle doesn't deselect any others.

Remove—After setting the ID number of a particle to deselect, click Remove to remove it from the selection.

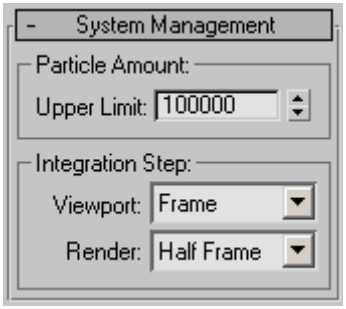
Clear Selection —When on, clicking Add to select a particle deselects all other particles.

Get From Event Level—Click to convert an Event-level selection to the Particle level. Available only at the Particle level.

Select By Event—This list shows all events in Particle Flow, and highlights selected events. To select all of an event's particles, click its list entry, or use standard viewport-selection methods.

 # Particles Selected—Shows the number of selected particles.

System Management rollout



Use these settings to limit the number of particles in the system, and to specify the frequency of updating the system.

Particle Amount group

Upper Limit—The maximum number of particles the system may have. Default=100000. Range=1 to 10000000.

Tip: You can have more than 10,000,000 particles in a single system by using multiple particle sources and wiring them to the same birth event. Note, however, that Particle Flow is limited to sending a maximum of 5,000,000 particles per event to the renderer.

Integration Step group

At each integration step, Particle Flow updates the particle system, applying each active action to particles in its event. A smaller integration step can improve accuracy, at the cost of calculation time. These settings let you apply different integration steps to the particle animation in the viewports and at render time.

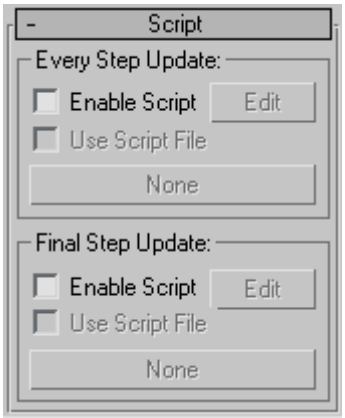
In most cases, the default Integration Step settings work fine. One instance in which increasing the integration-step frequency might help is when fast-moving particles that should collide with a deflector penetrate it instead.

Viewport—Set the integration step for animation playback in the viewports. Default=Frame (once per animation frame). Range=1/8 Frame to Frame.

Render—Set the integration step at render time. Default=Half Frame (twice per animation frame). Range=1 Tick to Frame.

There are 4,800 ticks in a second; thus, at the NTSC video rate of 30 frames per second, there are 160 ticks per frame.

Script rollout



This rollout lets you apply a script to the particle system at each integration step, as well as after the last integration step of each frame you view. Use an Every Step Update script to set up history-dependent properties, and a Final Step Update script to set up history-independent properties.

Every Step Update group

The Every Step Update script is evaluated at the end of each integration step, after all actions in the particle system are evaluated, and all particles are finally in their respective events. When, for example, you are setting up Material ID according to a particle index, it is important to be sure that particles are not about to jump to another event.

When you set up history-dependent properties, such as speed, it is important to do that at every

integration step, because otherwise the final position would be quite different.

Enable Script—Turn on to cause a script in memory to be executed at each integration step. You can modify this script by clicking the Edit button, or load and use a script file with the remaining controls in this group.

The default script modifies particle speed and direction, causing particles to follow a wavy path.

Edit—Click this button to open a text-editor window with the current script.

When Use Script File is off, this is the default Every Step Update script (*3dsmax\scripts\particleflow\example-everystepupdate.ms*). When Use Script File is on, this is the loaded script, if you've loaded one. If you haven't, clicking Edit displays the Open dialog.

Use Script File—When on, you can load a script file by clicking the button below.

[button]—Click this button to display an Open dialog that lets you specify a script file to load from disk. After you load a script, the name of script file appears on the button.

Final Step Update group

The Final Step Update script is executed after the last integration step has been completed for each frame that you view (or render). For example, if you play the animation in the viewport with Real Time turned off, Particle Flow runs this script at each frame, immediately before the particle system is rendered to the viewport. However, if you simply jump to a different frame, the script is run once only, so if the script assumes a certain history, you might get unexpected results.

For this reason, it's best to use the Final Step Update script to modify history-independent properties. For example, if no operators in the system depend on the material indices, you could use it to modify the material index. In this case there's no need to set those indices in every intermediate integration step. Also, you can set up a position channel in the Final Step Update script if you know the analytical expression for the position.

Enable Script—Turn on to cause a script in memory to be executed after the final integration step. You can modify this script by clicking the Edit button, or load and use a script file with the remaining controls in this group.

The default script modifies particle speed and direction, causing particles to follow a bulb-shaped path.

Edit—Click this button to open a text-editor window with the current script.

When Use Script File is off, this is the default Final Step Update script (*3dsmax\scripts\particleflow\example-finalstepupdate.ms*). When Use Script File is on, this is the loaded script, if you've loaded one. If you haven't, clicking Edit displays the Open dialog.

Use Script File—When on, you can load a script file by clicking the button below.

[button]—Click this button to display an Open dialog that lets you specify a script file to load from disk. After you load a script, the name of script file appears on the button.

Particle Flow Keyboard Shortcuts

To use Particle Flow keyboard shortcuts, the Keyboard Shortcut Override Toggle on the main toolbar must be turned on.

To view and customize these shortcuts, open the Customize User Interface dialog from the Customize menu, and then, from the Group drop-down menu, choose Particle Flow.

Particle Flow Function	Keyboard Shortcut	Description
Copy Selected In Particle View*	Ctrl+C	
Particle Emission Toggle	;	Toggles active status of all particle systems.
Particle View Toggle	6	Toggles Particle View dialog.
Paste In Particle View*	Ctrl+V	
Select All In Particle View*	Ctrl+A	Selects all events, actions, and wires.
Selected Particle Emission Toggle	Shift+;	Toggles active status only of particle systems whose icons are selected in the viewports.

* These functions are specific to Particle View. It is recommended that you use Customize User Interface to change their keyboard shortcuts only, and not assign them to toolbars, quads, or menus.

Actions

The Particle Flow components for creating particle systems are known collectively as *actions*. These are subdivided into three main categories:

- *Operators (page 2–142)*
- *Flows (page 2–208)*
- *Tests (page 2–210)*

Action Time Frames

Most actions in Particle Flow operate on particles in either of two time frames: once, when the particle first enters the event, or on a continuous basis, potentially changing particle behavior at each integration step. Some actions can work only in one time frame, while others can work on an instantaneous or continuous basis, depending on their settings. The tables below list each action with its time frame.

Operators

Operator	Time Frame
Birth	n/a
Birth Script	n/a
Delete	n/a
Force	Continuous
Keep Apart	Continuous
Mapping	Continuous
Material Dynamic	Continuous
Material Frequency	Once
Material Static	Once
Position Icon	Once, except continuous with Lock On Emitter
Position Object	Once, except continuous with Lock On Emitter
Rotation	Once, except continuous with Speed Space Follow
Scale	Once, except continuous with Absolute and Relative options
Script Operator	Depends on script
Shape	Once
Shape Facing	Continuous
Shape Instance	Once, except continuous with Animated Shape
Shape Mark	Once, except continuous with Align To Surface Animation
Speed	Once
Speed By Icon	Continuous
Speed By Surface	Depends on setting
Spin	Once, except continuous with Speed Space Follow

n/a=not applicable

Tests

Most tests in Particle Flow function only as tests. At each integration step, they check each particle for the specified conditions, and then return the test result: True or False. So, as tests, they work on a continuous basis. For example, Age Test checks each particle's age at every integration step, because particles might not reach the specified age until remaining in the event for a while, and also because another action in the event might change or reset particle age.

The principal exceptions to this are the Split tests, which test each particle only when it first enters the event. That is, a Split test splits the particle stream only once for each particle that enters the event. Any particles that remain in the event are not subject again to being split off from the stream by the same test. Also, the Send Out test performs no test, but simply moves particles along to the next event.

Some tests also function as operators, in that they directly affect particle behavior. These are the ones listed here, and the specified time frame is related to the operator functionality, not the test.

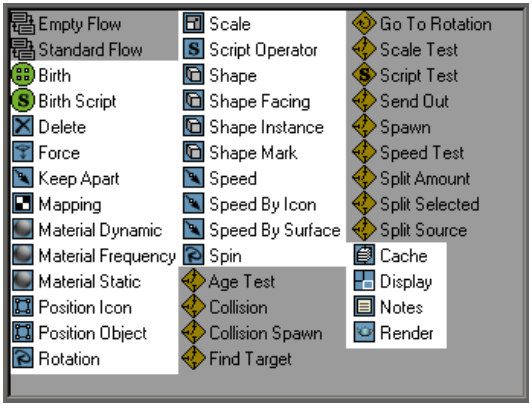
Test	Time Frame
Collision Spawn	Continuous
Find Target	Continuous
Go To Rotation	Continuous
Script Test	Depends on script
Spawn	Continuous, except once with Once

Operators

The operator is the basic element of the particle system; you combine operators into events to specify the particles' characteristics over a given period of time. Operators let you describe particle speed and direction, shape, appearance, and more.

The operators reside in two groups in the Particle View depot, and within each group appear in alphabetical order. Each operator's icon has a blue background, except for the Birth operators, which have a green background. The first group contains operators that directly affect particle behavior, such as transformation.

The second group, found at the end of the depot listing, contains four operators that serve more of a utility function: Cache, for optimizing particle-system playback; Display, for determining how particles appear in the viewports; Notes, for adding comments; and Render, for specifying render-time characteristics.



The Particle Flow operators in the Particle View depot

The primary operators in Particle Flow are:

Birth Operator (page 2–143)

Birth Script Operator (page 2–145)

Delete Operator (page 2–146)

Force Operator (page 2–204)

Keep Apart Operator (page 2–172)

Mapping Operator (page 2–195)

Material Dynamic Operator (page 2–191)

Material Frequency Operator (page 2–189)

Material Static Operator (page 2–187)

Position Icon Operator (page 2–147)

Position Object Operator (page 2–148)

Rotation Operator (page 2–153)

Scale Operator (page 2–156)

Spin Operator (page 2–154)

Select and Uniform Scale (page 1–441)

Script Operator (page 2–208)

Shape Operator (page 2–176)

Shape Facing Operator (page 2–176)

Shape Instance Operator (page 2–178)

Shape Mark Operator (page 2–183)

Speed Operator (page 2–159)

Speed By Icon Operator (page 2–162)

Speed By Surface Operator (page 2–167)

The utility operators are:

Cache Operator (page 2–197)

Display Operator (page 2–202)

Notes Operator (page 2–206)

Render Operator (page 2–206)

See also

Flows (page 2–208)

Tests (page 2–210)

Birth and Death

Birth Operator

Particle View (page 2–125) > Click Birth in an event or add a Birth operator to the particle system and then select it.

The Birth operator enables creation of particles within the Particle Flow system using a set of simple parameters. In general, use Birth as the first operator in any event connected directly to a *global event (page 3–949)*; this is called the *birth event (page 3–916)*.

You can specify a total number of particles, or a rate of particles born per second. You can also tell the system when to begin emitting particles, and when to stop.

Note: The Birth operator must always come at the beginning of a particle stream; the system doesn't let you position it elsewhere. You can place a Birth operator in an isolated event, but you can't then wire that event in series with a stream that already uses a Birth operator. However, you can wire multiple Birth operators, each in its own event, into a particle stream with an existing Birth operator, *in parallel*. The following procedure illustrates this. If you need to create particles midstream, use the *Spawn Test (page 2–230)* or *Collision Spawn Test (page 2–215)* test.

See also

Birth Script Operator (page 2–145)

Procedure

To use the Birth operator:

This procedure demonstrates the impossibility of using multiple Birth operators in series, and shows how to use multiple Birth operators in parallel.

1. Start or reset the software, and add a new PF Source object to the scene.
2. Press 6 to open Particle View.

   The default particle system contains a Birth operator as the first action in the birth event.
3. Try to drag the Birth operator elsewhere in Event 01.

   You can't. As you drag over the different actions in Event 01, a red line appears at the top of the event, showing that the Birth operator will be placed here, no matter where you release the mouse button.
4. Try to drag a new Birth operator from the depot to Event 01.

   As in the previous step, the only place you can drop the Birth operator is at the top of the event, replacing the existing Birth operator.
5. From the depot, drag the Birth operator to an empty area in the event display.

   Particle Flow creates a new birth event, Event 02, containing the Birth operator and a Display operator.
6. Wire the output of the global event, PF Source 01, to the event input of Event 02.

   Each birth event must be associated with a global event to be able to generate particles.
7. From the depot, add a Send Out test at the end of both Event 01 and Event 02.
8. Try to wire the test output of Event 02 to the event input of Event 01.

   Particle Flow doesn't let you, because this would result in two Birth operators in series.
9. Try to wire the test output of Event 01 to the event input of Event 02.

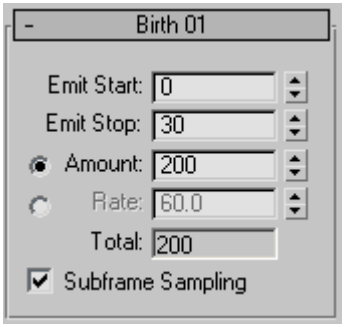
   Again, Particle Flow doesn't let you, because this would result in two Birth operators in series.
10. Drag a Speed operator to an empty area of the event display.

    Particle Flow creates a new event, Event 03.
11. Wire the test output of Event 01 to the event input of Event 03.
12. Wire the test output of Event 02 to the event input of Event 03.

    There's no problem wiring the two birth events to a single, third event. The birth events exist in the particle stream in parallel, each generating particles independently and then feeding its particle stream into a common event, where the two streams are combined.

    If the second birth event had its own global event, you could, at any point further downstream, separate the streams back out according to their origin using the *Split Source test (page 2–236)*. To do this, delete the wire from PF Source 01 to Event 02, add an *Empty Flow (page 2–209)* to the system, and then wire it to Event 02.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Emit Start—The frame number at which the operator begins emitting particles.

Emit Stop—The frame number at which the operator stops emitting particles.

Note: The Emit Start and Emit Stop values are tied to the system frame rate. If you change the frame rate, Particle Flow automatically adjusts the Emit values accordingly. For example, if you set Emit Start to 120 and Emit Stop to 300 using the default NTSC frame rate (30 fps), and then switch to PAL (25 fps) using the Time Configuration dialog, you've reduced the frame rate to 5/6 of the original value. Thus, Particle Flow automatically uses the same ratio to adjust the Emit settings, resulting in an Emit Start value of 100 and an Emit Stop value of 250. This allows the particle system to retain the timing you specify, no matter which frame rate you use.

Note: With Subframe Sampling off, such adjustments are rounded off to the nearest integer frame number. With Subframe Sampling on, fractional frame values that result from such adjustments will be used, but not displayed.

Amount—To specify the total number of particles emitted by the operator, choose Amount, and then set the quantity of particles.

Using the Amount option, the first particle is always emitted at the Emit Start frame, and the last particle is always emitted at the Emit Stop frame. Particles emitted between these endpoints appear at equal intervals within the emission period. For example, if you set Amount to 3, the second particle is emitted halfway through the emission period.

To determine the number of particles emitted per frame when using Amount, divide the Amount value by the number of emission frames (Emit Stop-Emit Start+1).

Rate—To specify the number of particles emitted *per second*, choose Rate, and then set the value. The operator emits this number of particles per second starting at the Emit Start frame and ending at the Emit Stop frame.

If you specify a birth rate value that isn't an integer multiple of the system frames-per-second value (set in the Time Configuration dialog), Particle Flow uses interpolation to determine when to emit particles. For example, if you use the system default rate of 30 frames per second, and set the birth rate to 4, the system would emit each particle at intervals of seven or eight frames if Subframe Sampling is off, or at intervals of 7.5 frames if Subframe Sampling is on.

Total—The calculated total number of particles emitted by the operator.

Subframe Sampling—Turning this on helps avoid particle "puffing" by emitting particles at a much higher subframe resolution (that is, throughout each frame), rather than using the relatively coarse frame resolution. Default=on.

"Puffing" is the effect of emitting separate "puffs" or clusters of particles, rather than a continuous stream. This effect is especially noticeable when the emitter is animated.

Turn off Subframe Sampling to cause particles to be emitted exactly at frame times. This makes it easier to sort particles by their age later.

Tip: If, when using the *Collision (page 2–212)* or *Collision Spawn test (page 2–215)*, you experience an irregular particle stream, try turning off Subframe Sampling.

Birth Script Operator

Particle View (page 2–125) > Click Birth Script in an event or add a Birth Script operator to the particle system and then select it.

The Birth Script operator enables creation of particles within the Particle Flow system using a MAXScript script. The script can use any program functionality available to MAXScript.

The default script (*3dsmax\scripts\particleflow\example-scriptbirth.ms*) emits particles for 100 frames in a wavy, circular path. To see this, turn off or delete any Speed and Position operators in the event.

See also

Birth Operator (page 2–143)

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Edit Script—Click this button to open the current script in a MAXScript Editor window.

For detailed information about the MAXScript utility, open the MAXScript Reference, available from Help menu > MAXScript Reference.

Emit Start—The frame number at which the operator begins emitting particles.

Uniqueness group

The Uniqueness setting provides a randomization seed that the script can use or ignore.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

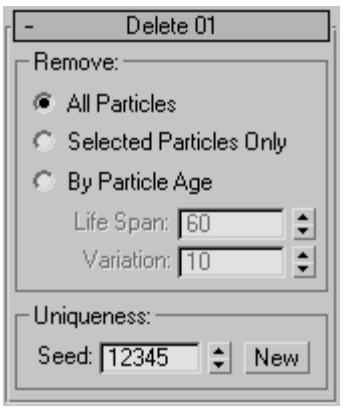
Delete Operator

Particle View (page 2–125) > Click Delete in an event or add a Delete operator to the particle system and then select it.

Use the Delete operator to remove particles from the particle system. By default, particles live "forever," that is, for the duration of the animation. The Delete operator lets you give them a finite life span. This is useful for eliminating particles once they've served their purpose in the animation.

The Delete operator is also necessary when using the Particle Age map, which, when incorporated into the Material Dynamic operator, applies different materials to particles depending on their age. Because the map works on the basis of a percentage of the particle's life span, you need to use a Delete operator to define the particles' maximum age. For an example of usage, see *Material Dynamic operator (page 2–191)*.

Interface



The user interface appears in the Parameters panel, on the right side of the Particle View dialog.

Remove group

Choose whether to delete all particles, selected particles, or particles past a specific age.

The Uniqueness setting enables randomization of maximum particle age using the By Particle Age > Variation setting.

All Particles—Deletes all particles in the event immediately.

Selected Particles Only—Deletes particles selected at the Particle sub-object level in the event immediately. See *Selection rollout (page 2–138)*.

By Particle Age—Deletes particles in the event after they've existed for a specific length of time, with an optional random variation. Choosing this makes the Life Span and Variation settings available.

Life Span—The number of frames of life allowed to the particles. After this period, they're deleted. Default=60.

Variation—The maximum amount by which Life Span may vary. To get each particle's actual life span, this value is multiplied by a random number between -1.0 and 1.0, and then added to Life Span. Default=10.

Uniqueness group

The Uniqueness setting enables randomization of maximum particle age using the By Particle Age > Variation setting.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Transforms

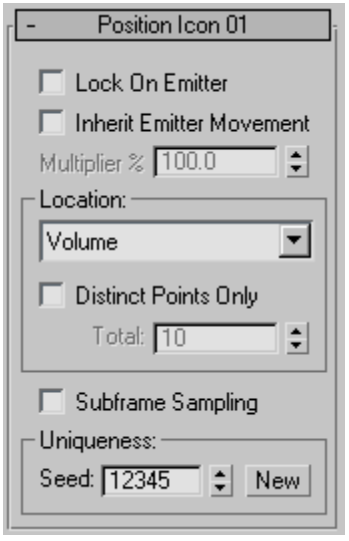
Position Icon Operator

Particle View (page 2–125) > Click Position in an event or add a Position operator to the particle system and then select it.

By default, particles are born, or *emitted*, from the Particle Flow icon. This topic uses the term *emitter* to refer to this icon. Use the Position Icon operator to control the initial placement of particles on the emitter. You can set the emitter to emit particles from its surface, volume, edges, vertices, or center. And, if you animate the emitter, you can cause its motion to be imparted to the particles it emits.

Alternatively, you can use *Position Object (page 2–148)* to emit particles from any other object.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Lock On Emitter—When on, causes all particles to remain at their initial positions on the emitter. In

effect, they're "glued" to the emitter. When off, each particle's birth position is determined by the emitter's current location. Default=off.

Inherit Emitter Movement—When on, the software sets each particle's rate and direction of motion to that of the emitter at the time of birth. When off, the emitter movement doesn't affect that of the particles. Available only when Lock On Emitter is off. Default=off.

Note: If a Speed operator exists later in the event, it overrides any motion imparted by Inherit Emitter Movement.

Multiplier—Determines, as a percentage, the extent to which particles inherit the emitter's motion. Available only when Inherit Emitter Movement is on. Default=100.

To have the particles move half as fast as the emitter, set Multiplier to 50. To have them move the same speed, use the default setting of 100. To make the particles move twice as fast as the emitter, set Multiplier to 200.

Location group

The Location drop-down list lets you specify where on the emitter the particles appear. You can create a variety of emission behaviors by varying the *icon type*, Location setting, and direction of emission. Default=Volume.

- **Pivot—**Emits particles from the center of the icon.
- **Vertices—**For the Box and Rectangle icon types, emits particles from the corners of the icon. For the Sphere icon type, emits particles from the six intersections of the three circles that define the sphere. For the Circle icon type, emits particles from the center.
- **Edges—**Emits particles from random points along the edges of the icon.
- **Surface—**Emits particles from random points on the surface of the icon.
- **Volume—**Emits particles within the volume of the icon.

Distinct Points Only—Limits emission to a specific number of points (see Total, below) on the specified Location type.

Emission points are still randomized; change the Uniqueness setting to alter these.

Total—Sets the number of emission points. Available only when Distinct Points Only is on. Default=10.

Subframe Sampling—When on, the operator acquires animation of the emitter icon on a tick basis (every 1/4800th of a second) rather than a frame basis. This provides greater precision in allowing the particle positions to follow animation of the emitter icon. Default=off.

Uniqueness group

The Uniqueness setting lets you alter the randomization of emission points.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Position Object Operator

Particle View (page 2–125) > Click Position Object in an event or add a Position Object operator to the particle system and then select it.

By default, particles are born, or *emitted*, from the Particle Flow icon. The Position Object operator lets you emit particles from any other object or objects in the scene instead. This topic uses the term *emitter* to refer to an object or objects used to emit particles with this operator. For example, use it to create a meteor's fiery trail.

Use the Position Object settings to control the initial placement of particles on the emitter. You can set the emitter to emit particles from its surface, volume, edges, vertices, or pivot, or from a sub-object selection. You can also control particle emission with a material applied to the object.

Alternatively, you can use the *Position Icon operator (page 2–147)* to emit particles from the Particle Flow icon.

Note: Typically, you would use Position Object in a birth event, so the particles appear initially at the emitter object(s). If you use it in a non-birth event, the particles jump to the designated emitter(s) upon entering the event. This could be useful for a teleportation effect.

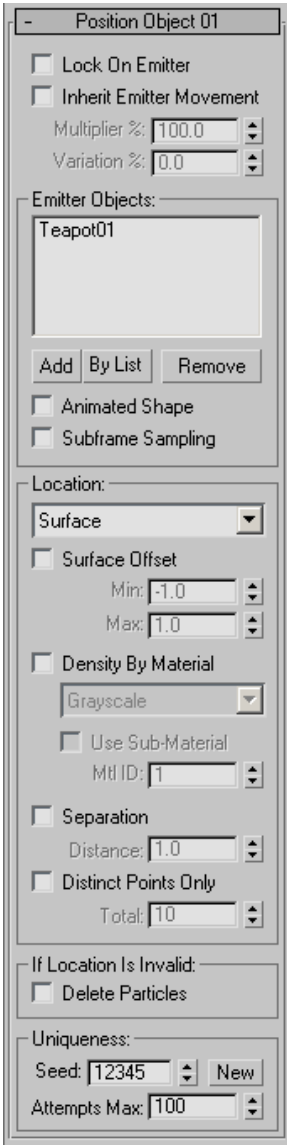
Tip: The default speed operator, *Speed (page 2–159)*, typically uses the Particle Flow icon to control the direction of the particle flow, so if the emitter object(s) and the icon aren't coincident and aligned, you might get unexpected results. For better control over the direction of particle movement when using Position Object, use the *Speed By Surface operator (page 2–167)* instead and designate as Surface Geometry the same object(s) as with Position Object.

Procedure

To use Position Object:

1. Create a Particle Flow system and one or more mesh objects to use as emitters. Animate the emitter objects and set particle system parameters as necessary.
2. In Particle View, add a Position Object operator to the birth event, replacing the Position Icon operator if one exists. Click the operator to display its parameters.
3. In the Emitter Objects group, click Add, and then select an object to serve as the emitter. Alternatively, click By List and then use the Select Emitter Objects dialog to designate one or more objects as emitters.
4. Use the Location drop-down list to choose where the particles should appear on the emitter: surface, vertices, and so forth.
5. Set other options as necessary. For example, if you want the particles to use the same motion as the emitter at the time of emission, turn on Inherit Emitter Movement. Or if you're using an emitter object whose shape is animated, turn on Animated Shape.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Lock On Emitter—When on, causes all particles to remain at their initial positions on the emitter. In effect, they're "glued" to the emitter. When off, each particle's birth position is determined by the emitter's current location. Default=off.

Inherit Emitter Movement—When on, the software sets each particle's speed and direction of motion to those of the emitter at the time of birth. When off, the emitter movement doesn't affect that of the particles. Available only when Lock On Emitter is off. Default=off.

Note: If a Speed operator exists later in the event, it overrides any motion imparted by Inherit Emitter Movement.

Multiplier—Determines, as a percentage, the extent to which particles inherit the emitter's motion. Available only when Inherit Emitter Movement is on. Default=100.

To have the particles move half as fast as the emitter, set Multiplier to 50. To have them move the same speed, use the default setting of 100. To make the particles move twice as fast as the emitter, set Multiplier to 200.

Variation—Determines, as a percentage, the extent to which inherited emitter motion can vary per particle. To determine the final multiplier for each particle, the Variation value is multiplied by a random number between -1.0 and 1.0, and then added to the Multiplier value. Default=0.

Emitter Objects group

Use these controls to assign objects to be used as particle emitters. The list in this group shows the objects, or reference geometry, that the operator uses as emitters.

If you don't assign any objects, Particle Flow uses the world origin (0,0,0) as the emitter source.

With more than one emitter object, the division of particles among the objects depends on which option you choose for Location. If you set Location to Pivot, then each object has one emission point and each object emits the same number of particles

(total number of particles/number of objects). But if you set Location to any other option, each object emits a number of particles proportionate to the number of available emission points per object. For instance, if you set Location to Volume, then a larger object will emit more particles than a smaller one. Similarly, with Location set to Vertices, an object with 100 vertices will emit twice as many particles as an object with 50.

If you delete from the scene an object designated as an emitter, its entry in the list is replaced with "<deleted>".

Use the Add and Remove buttons to edit this list.

Add—Adds an object to the list. Click Add, and then click an object in the viewport.

By List—Adds multiple objects to the list. Click By List to open the Select Emitter Objects dialog. This works just like the *Select Objects dialog (page 1–78)*; highlight the objects to use as emitters, and then click the Select button.

Remove—Removes an object from the list. Highlight the object in the list, and then click Remove.

Animated Shape—Turn on to allow particles to emit from the surface of an object whose form is animated by morphing or with modifiers.

Subframe Sampling—When on, the operator acquires animation of the emitter shape on a tick basis (every 1/4800th of a second) rather than a frame basis. This provides greater precision in allowing the particle positions to follow animation of the emitter object's form.

Location group

The Location drop-down list lets you specify where on each emitter the particles appear. This choice applies to all emitter objects. Default=Surface.

The Selected Vertices/Edges/Faces options in this list let you emit particles from a specific part of the emitter object by using an existing sub-object selection. To create this selection, first convert the object to an editable mesh or poly, or apply a Mesh Select or Poly Select modifier (or equivalent), select the sub-objects that are to emit particles, and then choose the corresponding Location option. If the software doesn't find a sub-object selection, it uses all sub-objects of the indicated type.

Pivot—Emits particles from the *original position* of the emitter object's pivot. Moving the pivot does not affect this option.

Vertices—Emits particles from randomly selected vertices of the emitter objects.

Edges—Emits particles from random points along the edge sub-objects of the emitter objects.

Surface—Emits particles from random points on the surface of the icon.

Volume—Emits particles from random points within the volume of the icon.

Selected Vertices—Emits particles from the current vertex sub-object selection.

Selected Edges—Emits particles from the current edge sub-object selection.

Selected Faces—Emits particles from the current face or polygon sub-object selection.

You can modify the Location choice with any combination of the following:

Surface Offset—Lets you specify a range of distances from the object surface for particle placement.

Surface Offset is unavailable when Location is set to Pivot or Volume. Use the Min and Max parameters to set the range.

Note: Negative values offset the particles below the surface.

Min/Max—When Surface Offset is on, these let you specify the range of offset distances from the

object surface. The Min value cannot exceed the Max value, and vice-versa. If you try to raise Min above Max, or lower Max below Min, both values change equally.

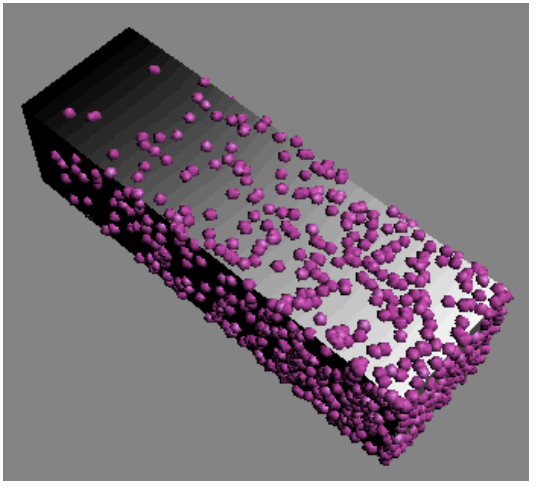
Density by Material—Varies emission over the emitter surface based on properties of the material applied to the emitter. For example, if the object is assigned a black-and-white checkered diffuse map and you choose the Grayscale option, particles are emitted only from the white checks.

Important: For material-influenced emission to appear properly in the viewports, two conditions are required:

- At least one viewport must be set to a shaded display mode.
- The material or map must have Show Map In Viewport turned on in the Material Editor.

The options are as follows:

- **Grayscale—**The software internally converts the material-based coloring (diffuse) to grayscale, and then emits more particles in lighter areas and fewer in darker areas.



The box emitter is mapped with a gradient. With Position Object set to Density By Material > Grayscale, the particles appear with greater frequency in the lighter areas of the box.

- **Opacity—**Particles are more likely to appear on opaque areas than on transparent areas.
- **Grayscale & Opacity—**Combines the two: More particles appear on light, opaque areas than on dark, transparent areas.
- **Red/Green/Blue—**Considers only the specified color channel. The higher its value at any given pixel, the more likely particles are to appear there.

Use Sub-Material—When on, uses a sub-material from the Multi/Sub-Object material assigned to the emitter to define density values.

This option allows usage of "invisible" materials for particle placement. If the emitter uses a Multi/Sub-Object material but its geometry doesn't use the ID that corresponds to one of the sub-materials, the sub-material doesn't appear. However, the operator can use it to calculate the density of particle placement. The software assumes the material to be applied to the entire object surface.

Mtl ID—Specifies the material ID of the sub-material to be used for particle emission.

Separation—When on, the software attempts to keep the particles apart by the amount specified in Distance. The likelihood of successfully separating particles depends on the number of particles, the specified distance, and the Attempts Max value.

Distance—Specifies the distance, in system units, by which the software should try to keep the particles separate. Default=1.0.

Distinct Points Only—Limits emission to a specific number of points (see Total, below) on the specified Location type.

Emission points are still randomized; change the Uniqueness setting to alter these.

Total—Sets the number of emission points. Available only when Distinct Points Only is on. Default=10.

If Location Is Invalid group

Delete Particles—When on, if the software cannot place a particle according to the current option, it deletes the particle. When off, particle placement is undefined; that is, it depends on other variables. Default=off.

If the number of particles is more important than the position on the object, leave this off. However, if exact placement is of higher priority than the number of particles, turn Delete Particles on.

Uniqueness group

The Uniqueness setting enables randomization of placement on the emitters.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Attempts Max—When using the Separation option, this specifies the maximum number of times Particle Flow will try to set particle positions that maintain the requested distance. If unsuccessful within this number of attempts, particles might be closer together than the requested distance. Also affects placement with Density By Material using either of the Grayscale options.

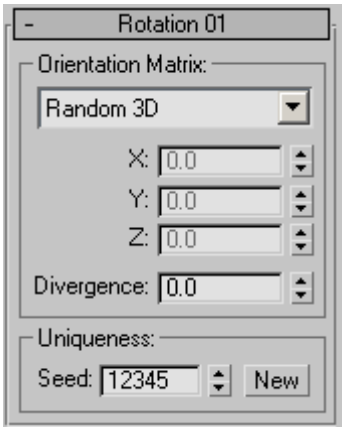
Rotation Operator

Particle View (page 2–125) > Click a Rotation operator in an event or add a Rotation operator to the particle system and then select it.

The Rotation operator lets you set and animate particle orientation during an event, with optional random variation. You can apply orientation in any of five different matrices: two random and three explicit. For some options you can set a degree of random variation or divergence from the specified orientation.

To cause particles to spin, use the *Spin operator (page 2–154)*.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Orientation Matrix group

The Orientation Matrix setting includes options that give particles random orientations, or let you specify orientation in any of three different ways. Default=Random 3D.

Random 3D—Gives each particle an arbitrary, random 3D orientation.

Random Horizontal—Gives each particle an arbitrary, random 3D orientation about the Z axis only. Rotation about the X and Y axes doesn't change (unless Divergence is used), so the particles remain vertical.

World Space—Orientation is specified in the world coordinate space.

Use the X/Y/Z settings to specify the orientation for all particles.

Speed Space—The coordinate space for particle orientation is determined by the particles' direction upon entering the event.

By default, using the *Speed operator (page 2–159)* > Along Icon Arrow option, particles are aimed straight down when born.

Use the X/Y/Z settings to specify the orientation for all particles.

Speed Space Follow—The coordinate space for particle orientation is determined continually while the particles are in the event. Thus, by default, particles constantly reorient themselves to "aim" in the direction they're traveling.

For example, if you add a *Find Target test (page 2–218)* to the end of the default Event 01, move the target sideways, and set Rotation to Speed Space Follow, the particles start out pointing straight down, as oriented by the Speed operator, and then gradually reorient themselves to point in the direction they're traveling: toward the target.

Use the X/Y/Z settings to specify the orientation for all particles. The Divergence setting is unavailable when using Speed Space Follow.

X/Y/Z—Set the basic orientation about the particles' local axes. Unavailable with the Random 3D and Random Horizontal options.

Divergence—Defines the range of variation (in degrees) for particle orientation. The actual deviation is calculated at random within this range. Unavailable with the Random 3D or Speed Space Follow option. Default=0.0.

Restrict Diverg[ence] To Axis—When on, lets you use the Divergence Axis controls to set the axis to which divergence will be applied. When off, the software uses a random axis for each particle. Unavailable with the Random 3D or Speed Space Follow option. Default=off.

Divergence Axis—Use the X/Y/Z settings to set the axis to which divergence will be applied. Default=1,0,0. Range=-1.0 to 1.0.

To specify one of the world axes, set the corresponding parameter to any non-zero value, and the others to 0. A negative value flips the axis. The numeric values come into play when you want to use an axis that's not aligned with the X, Y, or Z axis. In that case, you specify multiple non-zero values whose effect is relative to one another. For example, if you want the axis to be oriented halfway between the positive X and Y axes, you would set X and Y to the same positive amount. The actual value doesn't matter. Similarly, to set the axis to 30 degrees (1/3 of the angle) from the X axis to the Y axis, you'd set the Y value to twice that of the X value. For example, X=0.2 and Y=0.4, or X=0.5 and Y=1.0.

Uniqueness group

The Uniqueness setting affects the randomization of orientation with the Random 3D and Random Horizontal options, and also Divergence.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

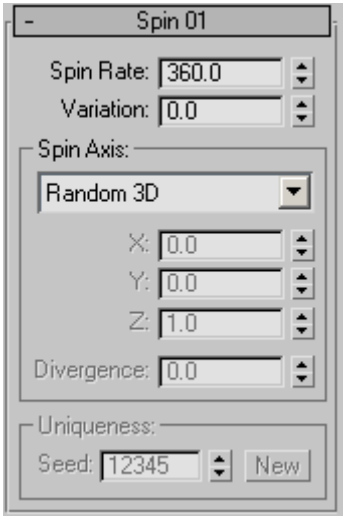
Spin Operator

Particle View (page 2–125) > Click a Spin operator in an event or add a Spin operator to the particle system and then select it.

The Spin operator gives an angular velocity to particles in an event, with optional random variation. Spin is applied once per event per particle, except when using the Speed Space Follow option; however, the settings can be animated.

To simply specify particle orientation, use the *Rotation operator (page 2–153)*.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Spin Rate—The rate of revolution in degrees per second.

Variation—The maximum amount, in degrees per second, by which the spin rate can vary. The actual variation is calculated once, at random, for each particle.

Spin Axis group

The Spin Axis setting includes options that let you apply the spin on a random or specific axis, with optional random variation of the spin axis. Default=Random 3D.

Random 3D—Spins each particle about an arbitrary, random 3D axis.

World Space—The spin axis is specified in the world coordinate space.

Use the X/Y/Z settings to specify the spin axis.

Particle Space—The spin axis is specified in the local coordinate space of each particle.

Use the X/Y/Z settings to specify the spin axis.

Speed Space—The coordinate space for particle orientation is determined by the particles' direction upon entering the event.

The X axis is aligned with the particle direction, while the Z axis is perpendicular to the X axis and is pointed upward with respect to world space as much as possible.

Use the X/Y/Z settings to specify the orientation for all particles.

Speed Space Follow—The coordinate space for particle orientation is determined by the particles' direction throughout the event. Use this option to allow particles that change direction during the event to keep the spin axis aligned with their direction.

The X axis is aligned with the particle direction, while the Z axis is perpendicular to the X axis and is pointed upward with respect to world space as much as possible.

Use the X/Y/Z settings to specify the orientation for all particles.

X/Y/Z—Use these to set the spin axis. Unavailable for the Random 3D option. Default=0,0,1. Range=-1.0 to 1.0.

To specify a single axis, set the corresponding parameter to any non-zero value, and the others to 0. A negative value flips the axis, and reverses the direction of rotation. The numeric values come into play when you give more than one axis a non-zero value, in which case their effect is relative to one another. For example, if you want the spin axis to be oriented halfway between the positive X and Y axes, you would set X and Y to the same positive amount. The actual value doesn't matter. Similarly, to set the spin axis to 30 degrees (1/3 of the angle) from the X axis to the Y axis, you'd set the Y value to twice that of the X value. For example, X=0.2 and Y=0.4, or X=0.5 and Y=1.0.

Divergence—Defines the range of variation (in degrees) for spin-axis orientation. The actual deviation is calculated at random within this range. Unavailable for the Random 3D option. Default=0. Range=0 to 180.

Uniqueness group

The Uniqueness setting affects the randomization of spin rate variation, spin axis with the Random 3D option, and Divergence with the other options.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Scale Operator

Particle View (page 2–125) > Click a Scale operator in an event or add a Scale operator to the particle system and then select it.

The Scale operator lets you set and animate particle size during an event, with optional random variation. Options for how scaling and animation are applied give this operator a great deal of flexibility.

Procedure

Example: To animate particle scale:

This procedure gives an overview of animating particles to grow for 30 frames, remain at full size for the next 30 frames, and then shrink for 30 frames. It also gives an option for having the particles pulsate in size during the middle period. To follow this procedure, you should be familiar with the basic usage of Particle Flow and Track View.

1. Set up a default Particle Flow system. Set the Shape operator to Sphere, and set the Display operator to Geometry.
2. Add a Send Out test to Event 01.
3. Add a new Scale operator to the event display, creating a new event. Set the new event's Display operator to Geometry, and make sure it's a different color than the one in Event 01.
4. Wire the Send Out test in Event 01 to Event 02.
5. In the Event 02 > Scale operator, set Type to Absolute, and animate the Scale Factor values from **10** to **100** over frames 0 to 30. In the Animation Offset Keying group, set Sync By to Event Duration.
6. Add an Age Test to Event 02. Set it to Event Age, set Test Value to **30**, and Variation to **0**.

   In this event, particles grow from 10 percent of their original size to full size over the first 30 frames of their existence, which is the same as their duration in the event. At that point, they'll be eligible to move to the next event.
7. Use a new Scale operator to create a new event, Event 03. Set the new event's Display operator to Geometry, and make sure it's a different color than the others.
8. Wire the Age Test in Event 02 to Event 03.
9. In the Event 03 > Scale operator, Set Type to Absolute, and in the Animation Offset Keying group, set Sync By to Event Duration.
10. Copy and paste the Age Test as an instance from Event 02 to Event 03.

    Particles in Event 03 remain at full scale for 30 frames, and then move on to the next event.
11. Use a new Scale operator to create a new event: Event 04. Set the new event's Display operator to Geometry, and make sure it's a different color than the others.
12. Wire the Age Test in Event 03 to Event 04.
13. In the Event 04 > Scale operator, set Type to Absolute, and animate the Scale Factor values from **100** to **10** over frames 0 to 30. In the Animation Offset Keying group, set Sync By to Event Duration.

14. Copy and paste the Age Test as an instance from Event 03 to Event 04.

    In this event, particles shrink from full size to 10 percent of their original size over 30 frames.

15. Play the animation.

    Next, you'll use a Noise controller to cause the particles in Event 03 to pulsate in size.

16. Add a second Scale operator to Event 03, below the existing Scale operator. Set it to Relative Successive, and set Sync By to Event Duration.

    The name of this operator should be Scale 04.

17. In Particle view, right-click the Scale Factor X % parameter field on the Scale 04 rollout, and choose Show In Track View.

    The Curve Editor opens with the X Scale Factor parameter highlighted.

18. Right-click the X Scale Factor item and choose Assign Controller from the menu.

19. In the Assign Float Controller dialog, double-click Noise Float.

    The Noise Controller dialog opens.

20. Set Strength to **200**, and to the right of Strength, turn on the >0 check box.

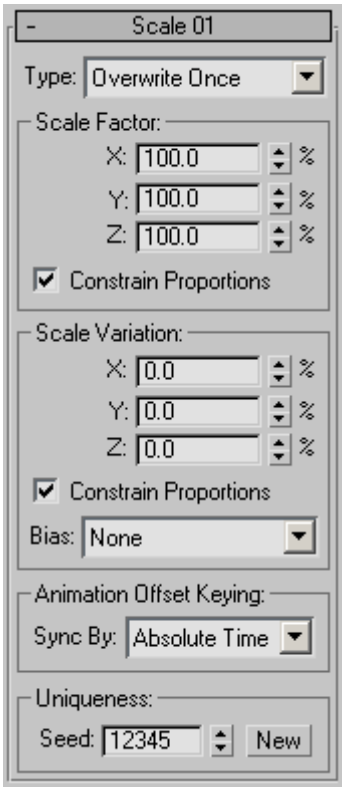
21. Copy this controller, and then paste it as instances to Y Scale Factor and Z Scale Factor.

    Although the Scale operator defaults to constraining all scale factors to the same value, thus scaling particles uniformly, changing controllers in Track View works on a per-parameter basis.

22. Play the animation.

    This time, the particles grow in size for 30 frames, then pulsate in size for the next 30 frames, and then shrink from the size at the end of Event 03 to 10 percent of that over the next 30 frames.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Type

The scaling Type setting includes options that let you scale particles once in an event or repeatedly, and apply scaling as an absolute or relative factor. Default=Overwrite Once.

Overwrite Once—Sets the scale *one time only* as an absolute percentage value, disregarding any previous scaling.

To produce a range of different birth sizes, and specify the scaling explicitly, use Overwrite Once with animated scale factors.

Inherit Once—Sets the scale *one time only* as a percentage of existing scaling. The existing scaling

should be specified with a previous Scale operator, the Scale setting in the *Shape operator (page 2–176)*, or a *Shape Instance operator (page 2–178)*.

For instance, if you previously used Scale with Overwrite Once to scale particles' birth size within a range, and then want to scale them to half their previous size, use Inherit Once and set Scale Factor to 50%.

Absolute—Sets the scale *continuously*, while the particle is in the event, as an absolute percentage value, disregarding any previous scaling.

Use Absolute if you want to animate the particles' scale explicitly.

Relative First—Sets the scale *continuously*, while the particle is in the event, as a percentage of existing scaling. The existing scaling should be specified with a previous Scale operator, or a Shape Instance.

Use Relative First when you want to scale the particles relative to the scaling at which they enter the event, and optionally animate the scaling, or before a Relative Successive operator.

Relative Successive—Scales the particles *continuously*, relative to scaling set earlier in the same event.

Note: Always precede Relative Successive *in the same event* with an operator that modifies the scale channel all the time, such as a Scale operator set to Absolute or Relative First, or a Shape Instance operator with the Scale check box turned on (it's on by default).

Scale Factor group

The operator performs scaling on each particle's local axes. It can scale particles on a single axis or on any combination of axes.

X/Y/Z—Sets the scaling as a percentage of the particle's current size. To scale uniformly, turn on Constrain Proportions, and then change any axis setting. Range=0 to 10000000. Default=100.

Constrain Proportions—When on, retains the current ratio of scale factor settings, so that changing any axis setting changes all of them. Default=on.

Scale Variation group

X/Y/Z—Sets the scaling variation as a percentage of the particle's previous size. To scale uniformly, turn on Constrain Proportions, and then change any axis setting. Range=0 to 100. Default=0.

Constrain Proportions—When on, retains the current ratio of scale variation settings, so that changing any axis setting changes all of them. Default=on.

Bias—Lets you choose how to distribute the scaling variation within the specified range or ranges. Default=None.

- **None—**No bias; scaling variation is distributed equally through the range.
- **Centered—**Scaling variation is concentrated near the middle of the range; that is, at 0.0%.

  With this choice, scaling will occur more frequently with low percentage values than with ones near the values you set. In other words, most scaling variants will be close to the specified scaling value.
- **Towards Minimum—**Scaling variation is concentrated near the lower end of the range (base-variation); that is, most scaling variants will be smaller than the scaling value.
- **Towards Maximum—**Scaling variation is concentrated near the upper end of the range (base+variation); that is, most scaling variants will be larger than the scaling value.

Animation Offset Keying group

If you animate the Scale Factor or the Scale Variation settings or both, the software can begin applying this animation to all particles as of the start frame of the animation or the first frame of

the current event, or to each particle based on its age. For instance, if you set Sync By to Particle Age, and set Scale Factor keys at frames 0 and 30, then the software will animate the scaling factor for each particle between its birth and its 30th frame of existence, if the particle is in the scaling event or a prior event. Following this example, any particle that has already reached its 30th frame of existence before reaching the event will enter the event fully scaled. That is, the software will apply the Scale Factor value that you set at frame 30 to each particle aged 30 or above at the moment it enters the scaling event. However, any particle that leaves the scaling event before reaching age 30 will stop scaling as of its exit frame. In other words, with respect to particles in other events, animation of action parameters is retroactive, but not post-active.

Alternatively, if you set Sync By to Absolute Time, the scaling is animated from frame 0 to frame 30 of the animation, regardless of particle age, and even if no particles are in the event at that time. Or, if you set Sync By to Event Duration, the scaling animation is applied to each particle as of the time that it enters the event.

Note: With the Overwrite Once and Inherit Once scaling types, scaling always occurs with respect to the entire animation; that is, in Absolute Time mode. Thus, when either of those scaling types is in effect, the Sync By setting is unavailable. Also, if you animate Scale Factor or Scale Variation when using Overwrite Once or Inherit Once, it doesn't cause scaling animation in the particles, but rather applies one-time scaling to particles born during that period. For instance, if you animate Scale Factor on all three axes from 100% to 200% over frames 0 to 30, particles born at frame 0 are normal size, particles born at frame 15 are one-and-one-half times normal size, and particles born at frame 30 (and thereafter) are twice normal size.

Sync By—Choose the time frame for applying animated parameters:

- **Absolute Time—**Any keys set for parameters are applied at the actual frames for which they're set.
- **Particle Age—**Any keys set for parameters are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for parameters are applied to each particle starting when it first enters the event.

Uniqueness group

The Uniqueness setting affects the randomization of scale variation.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Speed

Speed Operator

Particle View (page 2–125) > Click Speed in an event or add a Speed operator to the particle system and then select it.

This is the default Speed operator, which appears in the first event when you create a new Particle Flow icon. It provides basic controls over particle speed and direction.

Directional controls provided by the Speed operator are based on the position and orientation of the Particle Flow icon. For best results when using other objects as emitters, use the *Speed By Surface Operator (page 2–167)* operator instead.

See also

Speed By Icon Operator (page 2–162)

Keep Apart Operator (page 2–172)

Procedure

Example: To change particles' speed:

The Speed operator works on an instantaneous basis: It sets each particle's speed once only, when it enters the event. Even if you animate the Speed value, each particle moves at a constant rate of speed, defined by the value at the time it enters the event. This procedure demonstrates a trick you can use to change particle speed with an animated Speed value, thanks to Particle Flow's looping ability.

1. Create a default particle system, and position it at the top of the Perspective viewport.
2. Play the animation.

   The particles fall downward at the default rate: 300 units per second.
3. Add a Send Out test at the bottom of Event 01.
4. Add a Speed operator to an empty area of the event display.

   This creates a new event.
5. 

Turn on Auto Key, and move the time slider to frame 30.
6. In Particle View, click the new Speed 02 operator, and then in the Particle View parameters panel, set Speed to **0**.

   This animates the Speed value from 300 at frame 0 to 0 at frame 30.
7. 

Turn off Auto Key.
8. Wire the Send Out test in Event 01 to Event 02.
9. Play the animation.

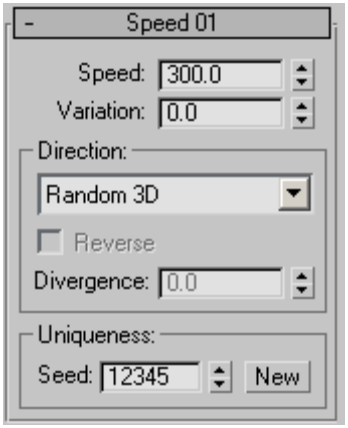
   The particles born later move slower, but all still move at a constant rate of speed.
10. Add a Send Out test at the end of Event 02.
11. Create a new event using an Age Test. Click the Age Test to display its parameters, and then set the following:
    - Event Age
    - Test Value=**1**
    - Variation=**0**
12. Wire the Send Out test in Event 02 to Event 03.
13. Wire the Age Test in Event 03 to Event 02.
14. Play the animation.

    All the particles slow down simultaneously and eventually come to a stop.

    Here's how it works: As each particle enters Event 02, its speed is set to the current Speed value in the Speed operator. Particle Flow then sends the particle immediately to Event 03, where it sits for one frame. Event 03 then returns the particle to Event 02, whose Speed value is now lower. Particle Flow perceives the returned particle as newly entering the event, so it changes its speed to the current Speed value. Thus, the particles continually return to Event 02 one frame later than before, and are assigned a progressively lower speed.

    If you wanted the particles to do something else after they stop, you could add a Speed Test to Event 02, above the Send Out test, set Test True If Particle Value to Is Less Than Test Value, and set Test Value to a very low value, such as 0.01. Then wire the Speed Test to a different event.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Speed—The particle speed in system units per second. Default=300.

A positive Speed value moves the particles in the direction determined by the Direction setting; a negative Speed value moves the particles in the opposite direction.

Note: Speed sets each particle's speed once only: when the particle enters the event (or is born, in the case of a birth event). If you animate the Speed value, particle speed does not vary: rather, each particle is given a constant speed equal to the current Speed value when it enters the event.

Variation—The amount by which particle speed can vary, in system units per second. Default=0.0.

To obtain each particle's speed, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the Speed value. For example, if Speed=300 and Variation=100, then each particle's speed would be between 200 and 400.

Direction group

The Direction drop-down list lets you specify which way the particles go after they're born. Default=Along Icon Arrow. In most cases, the actual direction also depends on the icon orientation. The primary exception is when Position > Location is set to Pivot.

Particle movement is always in a straight line unless influenced by other factors.

Along Icon Arrow—Particles move parallel to the icon arrow. Rotate the icon to change the direction.

Icon Center Out—Each particle moves along an imaginary line drawn between the particle's location and the icon center.

With the flat icon types (Rectangle and Circle), this results in all the particles moving in one plane, unless you increase Divergence above 0.0. With icons that have height (Box and Sphere), the particles move outward in three dimensions.

Particles at the center, as is the case when the Position operator's Location parameter is set to Pivot, arbitrarily move along the world X axis.

Icon Arrow Out—Each particle moves along an imaginary line drawn between the particle's location and the icon arrow. The line is perpendicular to the icon arrow, which is considered an infinite line for this purpose.

With the flat icon types (Rectangle and Circle), this results in all the particles moving in one plane, unless you increase Divergence above 0.0. With icons that have height (Box and Sphere), the particles move outward in a cylindrical formation.

Random 3D—Particles move in all directions. This option is affected by the Uniqueness setting.

Random Horizontal—Each particle moves in a random horizontal direction; that is, parallel to the world XY plane. This option is affected by the Uniqueness setting.

Inherit Previous—Uses the current direction of motion.

If you choose Inherit Previous but no direction was previously specified, the speed and direction are undefined; the particles don't move.

Reverse—When on, the direction is reversed. Default=off.

Using Reverse is the equivalent of multiplying the Speed value by -1. This option is unavailable if Random 3D or Random Horizontal is chosen.

Divergence—When on, spreads out the particle stream. Use the numeric setting to define the extent of the divergence, in degrees. Range=0 to 180. Default=0. The value can be animated.

This option is unavailable if Random 3D is chosen.

Tip: For a fountain-like spray, set Position > Location to Pivot, set Direction to Along Icon Arrow, set Divergence to the desired angle, and rotate the icon so its arrow points upward.

Uniqueness group

The Uniqueness setting enables randomization of speed variation, and randomization of direction with the Random 3D and Random Horizontal options.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

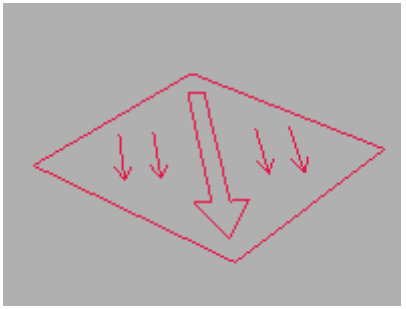
Speed By Icon Operator

Particle View (page 2–125) > Click Speed By Icon in an event or add a Speed By Icon operator to the particle system and then select it.

Click a Speed By Icon operator icon in the viewport > Modify panel

Create panel > Helpers > Particle Flow > SpeedByIcon

The Speed By Icon operator lets you use a special, non-rendering icon to control particle speed and direction.



Speed By Icon operator icon

When you add a Speed By Icon operator to the particle system in *Particle View (page 2–125)*, the Speed By Icon *operator icon* icon appears in the scene at the world origin (0,0,0). When animated, the operator icon's motion is imparted to the particles. If you delete the icon, the software also deletes the operator.

Note: If you add Speed By Icon from the Create panel, Particle Flow creates a separate event for the operator in the particle diagram.

Warning: If you delete a Speed By Icon operator icon in the viewport, you also delete the corresponding operator in the Particle Flow system.

See also

Speed Operator (page 2–159)

Speed By Surface Operator (page 2–167)

Keep Apart Operator (page 2–172)

Procedures

To use the Speed By Icon operator:

1. Add a Speed By Icon operator to an event. Be sure to add it *after* any existing Speed operators in the event. Alternatively, delete or turn off any existing Speed operators.

   The Speed By Icon operator icon appears at the world origin (0,0,0).

2. Animate the icon as you would any object. You can do this manually, or with a controller such

as Path Constraint. See the following procedure for an example of this.

The particles' motion is now under the influence of the animated icon.

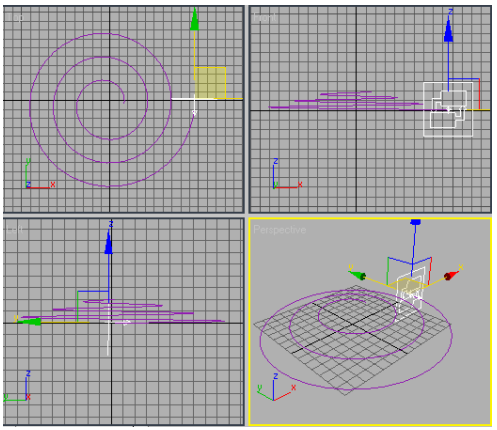
3. Set the Speed By Icon parameters. With the icon selected, you can do this in the Modify panel.

   In general, the default Icon Animation > Sync By setting, Event Duration, works best. It causes the icon animation to be applied to the particles within the time frame of the event containing the Speed By Icon operator.

Example: To send particles along a path:

1. Reset the software, and then add a Particle Flow system in the Perspective viewport.
2. On the Create panel, choose Shapes > Splines > Helix, and then, in the Perspective viewport, at the center of the grid, create a helix shape. Then, on the Parameters rollout, set the following:
   - Radius 1=**100**
   - Radius 2=**20**
   - Height=**20**
   - Turns=**3**
   - Bias=**0**

   This will serve as the particle path.
3. Move and rotate the Particle Flow source icon so that it's at the start of the helical path, and aimed along the path, as shown. Get the base of the arrow as close as possible to the start of the path.



4. Open Particle View (press 6).
5. In Event 01, click the Position Icon operator and set Location to Pivot.

   This causes the particles to be emitted in a thin stream.
6. Insert a Speed By Icon operator at the end of Event 01.

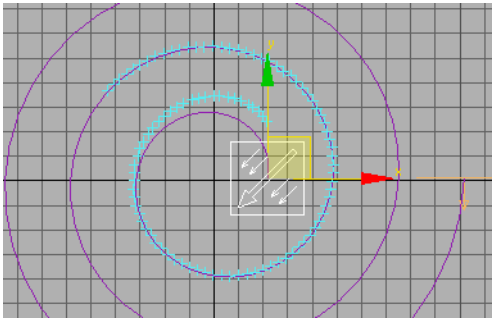
   The operator icon appears at the world origin.
7. Select the operator icon, and then, from the Animation menu, choose Position Controllers > *Path Constraint (page 2–398)*.

   When you move the mouse cursor into the viewport, a rubber-band line joins the cursor to the icon.
8. Select the Helix object.

   The command panel switches to the Motion tab, and the controls indicate that the software has applied a Position List controller to the icon, with the listed controllers being a Position XYZ and a Path Constraint. You can delete the former if you like, but it doesn't make any difference for the purposes of this procedure.
9. Close Particle View, and then drag the time slider back and forth between its extents a few times.

Tip: The best way to view the results is from the Top viewport.

The particles follow the path fairly closely, but diverge noticeably at the end of the path, where the turns are tightest. Correcting divergence in tight turns typically requires an increase in the acceleration limit.



The particle stream diverges from the path at its endpoint.

10. Stop at frame 100.

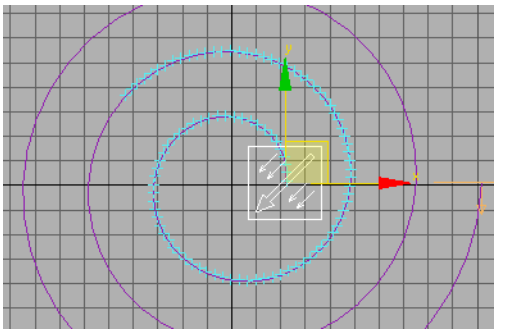
This will let you see the results of changing the Speed By Icon parameters as you make the changes.

11. Open the Modify panel, and, if necessary, select the Speed By Icon operator icon.

The operator parameters appear on the Modify panel. This is the case with any action that uses a unique icon, and lets you adjust the parameters without using Particle View.

12. On the Parameters rollout, use the Accel Limit spinner to slowly increase the value as you observe the changes to the particle path.

As you approach a value of 150, the end of the particle path comes closer and closer to the Helix. You shouldn't see much change above 150.



The endpoints now coincide.

13. Drag the time slider again.

The particle path remains similar to that of the Helix for the entire duration of the animation. With other setups, you might need to use different values for Influence %, and for paths with tight twists and turns, you might need to increase Accel Limit as well.

You might be wondering why you didn't simply replace the default Speed operator with the Speed By Icon operator. To see why not, try this:

14. Go to frame 100, open Particle View, and click the Speed 01 operator's icon to turn it off. Play the animation.

The entire particle path is offset from the Helix, so it's difficult to tell how closely it follows the latter's curves.

15. Try adjusting the Speed By Icon's Accel Limit and Influence % setting to get the particles to follow the helical path more closely.

It's not easy. You get more reliable results using the Speed operator to set the initial speed, and then Speed By Icon to set the particle path.

To use the Use Icon Orientation option:

The Use Icon Orientation option applies arc-like motion to the particles based on rotation animation of the icon. To best understand how it works, it's necessary to isolate its influence by eliminating any potentially conflicting factors.

1. Reset the software, and then add a Particle Flow system in the Perspective viewport. Set its position to the world origin: (0,0,0).
2. Open Particle View and click the Birth operator. Set Emit Stop=**0**.

   This causes all particles to appear at frame 0.
3. Click the Speed operator and press the Delete key to delete it.
4. Drag a Speed By Icon operator from the Depot to the end of Event 01.

   The Speed By Icon operator icon appears at the world origin.
5. Auto Key Select the Speed By Icon operator icon, and then right-click the icon and choose Rotate from the menu. Go to frame 20, turn on Auto Key, and rotate the icon 180 degrees about the X axis.
6. Auto Key Turn off Auto Key.
7. Drag the time slider.

   The particles aren't affected by the icon animation.

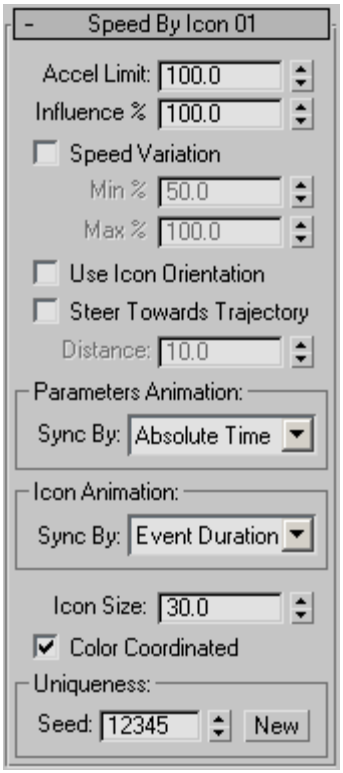
   Because the Speed By Icon operator icon is selected, the operator's parameters appear on the Modify panel.
8. On the Modify panel, turn on Use Icon Orientation. Drag the time slider again.

   This time, the particles move in unison with the icon rotation, as if they were glued to an infinite plane coincident with the icon.
9. Go to frame 10, select the Particle Flow source icon (not the Speed By Icon operator icon), and then right-click the icon and choose Move from the menu. Move the icon straight up, on the Z axis, about 100 units.

   As you drag upward, the particles move away from the icon.
10. Drag the time slider again. This time the particles move in an arc around the icon.
11. Try moving and rotating the two icons, playing the animation each time you make a change. Also set Emit Stop back to 30.

    As you can see, the possibilities with just this simple setup are myriad. In combination with the many other variables and options in Particle Flow, they're endless.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog. When the icon is selected, the parameters also appear on the Modify panel.

Accel(eration) Limit—The maximum amount, in system units per second per second, by which the

particles' speed can change in order to match the speed of the operator icon. Default=100.0.

If the operator icon changes speed or direction rapidly, increase this value to let the particles follow it more closely.

Tip: Use a lower Accel Limit value for smooth motion, and a higher value when greater accuracy is needed, such as when the particles should hit a small target. You can animate this setting (use Sync By > Event Duration) to specify different appropriate values, depending on the required results.

Influence %—Determines the mix of the previous speed with the speed of the operator icon. Default=100.0. Range=0.0 to 100.0.

At the default value of 100, the speed is controlled only by that of the icon. At 0, the icon speed doesn't affect particle speed at all. At in-between values, the mix proportion is determined by the Influence % value.

Speed Variation—Enables random variation of the particles' speed, so they don't all move at exactly the same rate as the icon. Turn this on, and then set the minimum and maximum percentages of variation. Default=off.

Using Speed Variation can help keep the particles from bunching up as they travel along the animation path. Note, however, that each individual particle's speed doesn't vary; it travels at a constant speed.

Min %/Max %—Set the minimum and maximum percentages, respectively, of particle speed variation. Default=50.0 (Min %), 100.0 (Max %).

For example, if you set Min % to 50 and Max % to 200, the particles will vary in speed between half and two times the speed of the icon.

Use Icon Orientation—Applies animation of the icon's orientation to the particles.

By default, Speed By Icon controls particle motion only by animation of the icon's position. If you turn on Use Icon Orientation, the software also applies rotation of the icon to the particle motion. In essence, this motion is circular, and its extent depends on the distance between each particle and the icon. Default=off.

The best way to understand how this works is by using it. To try out Use Icon Orientation, see *this procedure (page 2–164)*.

Steer Towards Trajectory—Particles farther away than the Distance value from the icon are moved directly toward the icon. Default=off.

Distance—Sets the distance between the particles and the operator icon beyond which Steer Towards Trajectory takes effect. Default=10.0.

Parameters Animation group

If you animate the operator settings, the software can begin applying this animation to all particles as of the start frame of the animation or the first frame of the current event, or to each particle based on its age. For instance, if you set Sync By to Particle Age, and set Speed Variation keys at frames 0 and 30, then the software will animate the speed variation for each particle born while the event is active between its birth and its 30th frame of existence. Alternatively, if you set Sync By to Absolute Time, the speed variation is animated from frame 0 to frame 30 of the animation, even if no particles are in the event at that time. Or, if you set Sync By to Event Duration, the speed variation animation is applied to each particle as of the time that it enters the event. For further information, see *Animation Offset Keying group (page 2–158)*.

Sync By—Choose the time frame for applying animated parameters:

- **Absolute Time—**Any keys set for parameters are applied at the actual frames for which they're set.

- **Particle Age—**Any keys set for parameters are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for parameters are applied to each particle relative to the frame at which it first enters the event.

Icon Animation group

The software can begin applying animation of the Speed By Icon operator icon to all particles as of the start frame of the animation or the first frame of the current event, or to each particle based on its age. For an explanation, see *Animation Offset Keying group (page 2–158)*.

Sync By—Choose the time frame for applying animation of the icon to the particles:

- **Absolute Time—**Any keys set for icon motion are applied at the actual frames for which they're set.
- **Particle Age—**Any keys set for icon motion are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for icon motion are applied to each particle relative to the frame at which it first enters the event.

Icon Size—Sets the size of the operator icon. This setting is for visibility only; it doesn't affect particle behavior.

Uniqueness group

The Uniqueness setting enables randomization of the speed variation range.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Speed By Surface Operator

Particle View (page 2–125) > Click Speed By Surface in an event or add a Speed By Surface operator to the particle system and then select it.

The Speed By Surface operator lets you control particle speed and direction with any object or objects in the scene. This is in contrast to the default *Speed operator (page 2–159)*, which uses the Particle Flow icon to control particle speed and direction. Speed By Surface also provides options for controlling speed by materials in the scene, and continuous speed control.

Typically, you use Speed By Surface to assign speed and direction characteristics based on the objects used as emitters with the *Position Object operator (page 2–148)*. That is, you assign the same objects as emitters with Position Object and as Surface Geometry objects with Speed By Surface.

Tip: For greater control over particle motion on the surface, apply a *Speed operator (page 2–159)* before the Speed By Surface operator in the same event. Use the Speed operator to specify particle direction on the surface, and the Speed By Surface operator to adjust the vertical component of the speed vector to follow the contour of the surface.

See also

Speed By Icon Operator (page 2–162)

Keep Apart Operator (page 2–172)

Procedure

Example: To use Speed By Surface:

The Speed By Surface operator offers a wealth of possibilities for usage. This procedure will lead you through several examples, but you're encouraged to explore the operator further on your own to learn more about it.

1. Start or reset the software, and add a Particle Flow system. Position the source icon at the

world origin (X/Y/Z=**0**), and set the icon's Length and Width values both to **30**.

2. Add a Cylinder primitive to the scene. Position it above the source icon: X/Y=**0** and Z=**30**. Set Radius=**20** and Height=**40**.
3. Play the animation, and then stop.

   The particles move downward, using the default starting setup.
4. Open Particle View and add a Speed By Surface operator to the end of Event 01. Click the operator in the event to display its parameters panel in Particle View.
5. In the Surface Geometry group, click Add, and then select the cylinder.
6. Play the animation again, and then stop.

   The particles still move downward.
7. Go to frame 15, so you can see the particles, and then slowly move the cylinder downward along the Z axis, while watching the particles in the Perspective and Front viewports. Stop when the emitter is above the cylinder.

   As the bottom of the cylinder passes below the emitter, more and more particles start moving on the XY plane rather than perpendicular to it. That's because they eventually become closer to one of the vertical sides of the cylinder, rather than the bottom, at which point they move perpendicular to the vertical sides. When the top becomes the closest side, the particles again move vertically, but upward instead of downward.
8. Move the cylinder back up until the emitter is at its vertical center. Then, on the Modify panel, click the lower part of the Sides spinner to decrease the number of sides, one at a time.

   Each time you click, the particle streams traveling outward change, to move perpendicular to the vertical sides as they change position. You might also try rotating the cylinder, and note that the particle streams also rotate, like the spokes of a wheel.
9. In the Speed By Surface parameters, change Direction to Out Of Surface, and then move the cylinder up and down.

   The overall behavior doesn't change much. Out Of Surface becomes more useful when you use the Control Speed Continuously option, as you'll observe shortly.
10. In the Speed By Surface parameters, change Direction to Parallel To Surface, and then move the cylinder up and down.

    Now the particle behavior is effectively opposite of that with the other two options. When the particles are closer to the top or bottom, they move along the XY plane, and when they're closer to the vertical sides, they move downward. If you want them to move upward instead, set Speed to a negative value.

    Next, you'll discover how the Control Speed Continuously option alters particle behavior.
11. In the Speed By Surface parameters, choose Control Speed Continuously, and then drag the time slider.

    The particles spread out from the emitter, and then start orbiting the cylinder in a roughly cylindrical overall formation. The software continually checks to see which side of the cylinder a particle is closest to, and, if necessary, changes its direction to make it travel parallel to that side. You can affect the size of the particles' orbits by changing their speed.

    Lastly, you'll see how to contain particle motion within the cylinder.
12. Make the cylinder considerably larger: Radius=**60** and Height=**90**. Activate the Perspective viewport, if necessary, and then press **F3** to set it to Wireframe view.
13. In Particle View, click the Speed 01 operator and set Direction to Random 3D.

14. Click the Speed By Surface operator. Set Speed to 300 if necessary. In the Direction group, choose Out Of Surface, and then drag the time slider.

    After exiting the emitter, the particles always move away from the nearest surface, with the result that they end up milling about in the center of the cylinder.

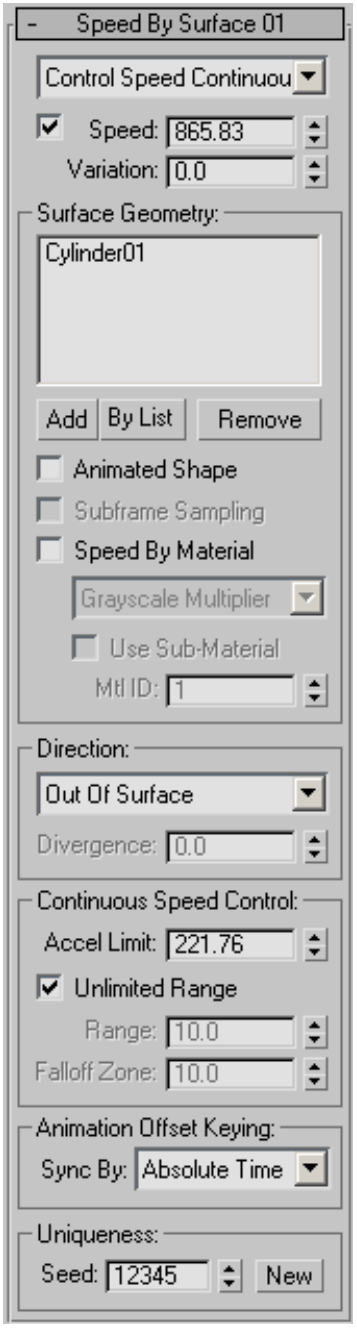
15. Gradually increase the Speed value, up to **1,000** or so.

    The particles start to spread out vertically, and eventually start moving so fast that they escape the cylinder. You can make them stay inside the cylinder by letting them accelerate more quickly, so they can turn before moving past the nearest surface.

16. Increase the Accel Limit setting until the particles no longer exit the cylinder.

17. Try giving the particles more room to move about in by increasing the size of the emitter. Also try changing other settings throughout the particle system to see their effects. The possibilities don't end here, and the more you experiment, the more you'll learn about how this powerful operator works.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

The first setting lets you choose whether the operator controls speed once or continuously.

Set Speed Once—The operator sets the speed for each particle once: when the particle enters the event.

Control Speed Continuously—The operator sets particle speed throughout the event. When you choose this option, the *Continuous Speed Control group (page 2–172)* becomes available.

Speed—The particle speed in system units per second. Default=300.

With the Control Speed Continuously option, you can turn off Speed. This causes Particle Flow to use the current particle speed.

Using a negative Speed value causes particles to move in the opposite direction of that effected by a positive speed.

Variation—The amount by which particle speed can vary, in system units per second. Default=0.0.

To obtain each particle's speed, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the Speed value. For example, if Speed=300 and Variation=100, then each particle's speed would be between 200 and 400.

Surface Geometry group

Use these controls to assign objects to affect particle speed and direction. The list in this group shows the objects, or reference geometry, that the operator uses to control speed and direction.

If you don't assign any objects, the current particle speed and direction are not affected.

You can assign multiple Surface Geometry objects, but if they're not all emitters as well, the results can be difficult to interpret. In general, assign the same objects as Surface Geometry that your system uses as emitters.

Tip: For optimal performance, when moving particles with respect to high-polygon objects, create a low-poly proxy version to use as surface geometry, link it as a child of the high-poly object if necessary, and then hide the low-poly proxy.

Use the Add and Remove buttons to edit this list.

Add—Adds an object to the list. Click Add, and then click an object in the viewport.

By List—Adds multiple objects to the list. Click By List to open the Select Surface Objects dialog. This works just like the *Select Objects dialog (page 1–78)*; highlight the objects to use to control speed and direction, and then click the Select button.

Remove—Removes an object from the list. Highlight the object in the list, and then click Remove.

Animated Shape—Turn on to allow particles to follow the surface of an object whose form is animated by morphing or with modifiers.

Subframe Sampling—When on, the operator acquires animation of the Surface Geometry shape on a tick basis (every 1/4,800th of a second) rather than a frame basis. This provides greater precision in allowing the particle positions to follow animation of the Surface Geometry object's form.

Speed by Material—Varies particles' existing speed and direction based on properties of the material applied to each Surface Geometry object. For example, if an object is assigned a black-and-white checkered diffuse map and you choose the Grayscale Multiplier option, particles near the white-checked areas move faster than those from the black-checked areas.

Note: For material-influenced speed to appear properly in the viewports, two conditions are required: at least one viewport must be set to a shaded display mode, and the material or map

must have Show Map In Viewport turned on in the Material Editor.

The options are as follows:

- **Grayscale Multiplier—**Lets material luminance control speed, with darker areas producing slower particles and lighter areas producing faster ones. The software multiplies the luminance of the material near each particle, converted to a percentage, by the particle's current speed. A luminance of 0 converts to 0%, of 128 converts to 50%, and of 255 converts to 100%.

  For example, if the speed of a particle traveling at 50 units per second is influenced by a pixel whose luminance is 90, the resulting speed is 90/255*50, or about 17.6 units per second.

- **Signed Grayscale—**Works like Grayscale Multiplier, but the multiplier can be negative as well, causing reversal of motion. Signed Grayscale uses a material luminance value of 128 as the midpoint, and assigns it a multiplier of 0%. Luminance values from 0 to 127 result in multipliers of -100% to about -1%, respectively, and values of 129 to 255 result in multipliers of about 1% to 100%, respectively.

- **RGB as World XYZ Mult.—**Works like Grayscale Multiplier, but uses the intensity of the material's red, green, and blue channels to affect particle speed on the world X, Y, and Z axes, respectively. So, for example, if the material pixel is pure red, that is, its RGB value is (255,0,0), then the particle will retain its current speed on the world X axis, but its speed on the Y and Z axes will be reduced to 0. Similarly, a medium-yellow pixel (128,128,0) will cause speed on the world X and Y axes to be reduced by half, and will cut speed on the Z axis to 0.

- **RGB as Local XYZ Mult.—**Works like RGB as World XYZ Mult., but uses the object's local coordinates rather than world coordinates.

Use Sub-Material—When on, uses a sub-material from the Multi/Sub-Object material assigned to the Surface Geometry object to define speed.

This option allows usage of "invisible" materials for controlling particle speed. If the emitter uses a Multi/Sub-Object material but its geometry doesn't use the ID that corresponds one of the sub-materials, the sub-material doesn't appear. However, the operator can use it to calculate the density of particle placement. The software assumes the material to be applied to the entire object surface.

Mtl ID—Specifies the material ID of the sub-material to be used for particle speed control.

Direction group

The Direction drop-down list lets you specify which way the particles go after they're born. Default=Surface Normals. In most cases, the actual direction also depends on the icon orientation. The primary exception is when Position > Location is set to Pivot.

Particle movement is always in a straight line unless influenced by other factors.

Surface Normals—Each particle moves along a line perpendicular to the nearest face. The direction the surface faces doesn't matter.

Out Of Surface—Particles move away from the closest face.

Tip: You can use this option to confine particles to the interior of an object. Position the emitter inside the object, designate the object as the Surface Geometry, choose Control Speed Continuously, and choose Out Of Surface. Each time a particle comes close to a surface, it turns to travel directly away from the surface. Control the particles' travel range with the Speed setting, but keep in mind that particles moving very fast might "escape" their container. If this happens, increase the Accel Limit setting; this lets the particles turn more quickly.

Parallel To Surface—Each particle travels parallel to the nearest face.

To make particles orbit an object, use this with Control Speed Continuously, and set appropriate Continuous Speed Control values (see the following section).

Divergence—When on, spreads out the particle stream. Use the numeric setting to define the extent of the divergence. Range=0 to 180. Default=0. The value can be animated.

This option is unavailable if Control Speed Continuously is chosen.

Tip: For a fountain-like spray, set Position > Location to Pivot, set Direction to Along Icon Arrow, set Divergence to the desired angle, and rotate the icon so its arrow points upward.

Continuous Speed Control group

When you choose the Control Speed Continuously option, these controls become available. Basically, these controls let the particles move about within the region of the Surface Geometry object rather than in a straight line away from the emitter.

Accel Limit—Sets the maximum acceleration. The higher this value, the more quickly particles can turn and change speed.

Tip: Use a lower Accel Limit value for smooth motion, and a higher value when greater accuracy is needed, such as when the particles should hit a small target. You can animate this setting (use Sync By > Event Duration) to specify different appropriate values, depending on the required results.

Unlimited Range—When on, the surface controls speed and direction of particles at any distance. When off, particles must be within a specified range. Default=on.

Range—The maximum distance, in system units, between the particles and controlling surface. The Surface Geometry will not control particles beyond this distance.

Falloff Zone—The distance, in system units, beyond the Range value within which the Surface Geometry exerts partial control over particles. The degree of control diminishes from 100% at the Range+Falloff Zone distance.

Animation Offset Keying group

Choose the time frame for applying animated parameters. For an explanation, see *Animation Offset Keying group (page 2–158)*.

Absolute Time—Any keys set for parameters are applied at the actual frames for which they're set.

Particle Age—Any keys set for parameters are applied at the corresponding frames of each particle's existence.

Event Duration—Any keys set for parameters are applied to each particle starting when it first enters the event.

Uniqueness group

The Uniqueness setting enables randomization of speed variation, and randomization of direction with the Random 3D and Random Horizontal options.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Keep Apart Operator

Particle View (page 2–125) > Click Keep Apart in an event or add a Keep Apart operator to the particle system and then select it.

The Keep Apart operator, a member of the Speed operator family, lets you apply forces to particles

to cause them to separate, in order to prevent or minimize collisions among particles. Alternatively, you can use a negative force to keep particles from separating too much. The operator works by controlling particle speed and acceleration.

Note: Keep Apart doesn't use particle geometry; rather, it creates a spherical force field centered on the pivot of each particle. You can adjust the size of the force field by changing the particle size.

Tip: In certain cases, the default settings might not be sufficient to keep particles from interpenetrating. For increased separation, use higher values for Force and Accel Limit, choose Relative To Particle Size, and increase the Core % value.

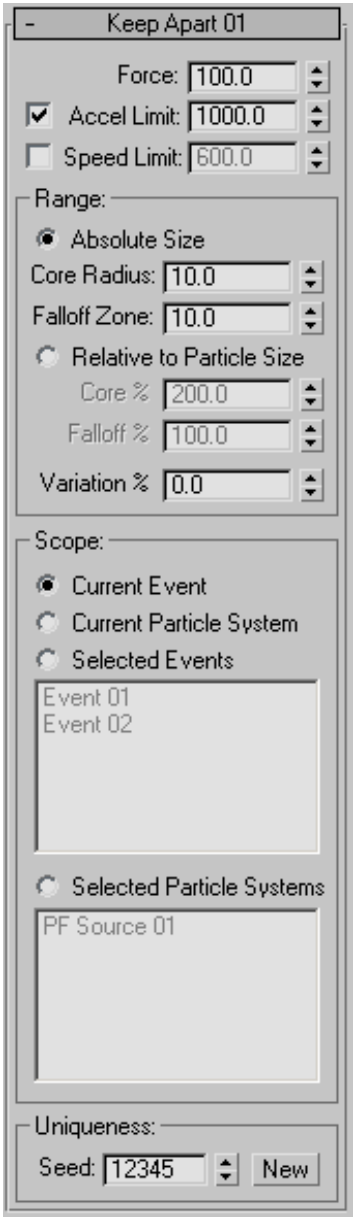
See also

Speed Operator (page 2–159)

Speed By Surface Operator (page 2–167)

Speed By Icon Operator (page 2–162)

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Force—The amount of force applied to the particles. Use positive values to separate particles,

or negative values to move them closer together. Default=100.0

Accel(eration) Limit—When on, lets you set a maximum acceleration value that can be applied to particles' motion. When off, the software uses any necessary acceleration. Default=on, 1000.0.

Tip: Use a lower Accel Limit value for smooth motion, and a higher value when greater accuracy is needed, such as when the particles should hit a small target. You can animate this setting (use Sync By > Event Duration) to specify different appropriate values, depending on the required results.

Speed Limit—When on, lets you set a maximum speed value that can be applied to particles' motion. When off, the software uses any necessary speed. Default=off, 600.0.

Range group

You can set a volume and falloff within which the force takes effect, either as absolute distances or relative to particle size. Default=Absolute Size.

Absolute Size—Choose this to set the core radius and falloff zone as absolute distances, in system units.

Core Radius—The distance from each particle's pivot point, in system units, within which the force is applied at full strength. Default=10.0

Falloff Zone—The distance beyond the core radius, in system units, over which the force diminishes from full strength to 0. Default=10.0

Relative to Particle Size—Choose this to set the core radius and falloff zone as percentages, relative to the radius of each particle. The particle radius is determined by measuring the distance from the pivot point to the farthest corner of its bounding box.

Core %—The distance from each particle, as a percentage of the particle radius, within which the force is applied at full strength. Default=200.0

Falloff %—The distance beyond the core radius, as a percentage of the radius, over which the force diminishes from full strength to 0. Default=100.0

Variation %—The amount by which range values can vary randomly, as a percentage of the specified values. The software uses the same random variation value for both Core and Falloff values, whether absolute or relative. Default=0.0

For example, if you choose Absolute Size, set Core Radius to 40 and Falloff Zone to 20, and set Variation % to 50, then each particle's actual core radius will be a random number in the range 20 to 60, and the falloff zone a random number in the range 10 to 30. For each particle, the software uses the same Variation % value, so if Core Radius is determined to be 20, then Falloff Zone would be 10.

Scope group

By default, the Keep Apart force keeps particles only in the current event (that is, the event that contains the Keep Apart operator) apart from each other when used locally, or particles in each event in the current *flow (page 3–942)* apart when used globally. These settings let you alternatively keep particles away from particles in other events or flows without affecting the behavior of latter. Default=Current Event.

Current Event—Keeps particles in the current event only away from each other.

When used globally, keeps particles in each event in the current flow apart from each other, but not apart from particles in other events in the flow.

Current Particle System—Keeps particles in the current event away from each other, and away from all other particles in the current *flow (page*

3–942). The particles not in the current event are not influenced by the Keep Apart operator.

When used globally, keeps apart and influences all particles in the current flow.

Selected Events—Keeps particles in the current event away from all particles in the events highlighted in the list below this choice. Particles not in the current event are not influenced by the Keep Apart operator.

When used globally, keeps all particles in the current flow apart from all particles in the events highlighted in the list below this choice, but influences only particles in the current flow.

After choosing this option, click items in the list to highlight them.

Selected Particle Systems—Keeps particles in the current event away from all particles in the *flows (page 3–942)* highlighted in the list below this choice. The particles not in the current event are not influenced by the Keep Apart operator.

When used globally, keeps all particles in the current flow apart from all particles in the flows highlighted in the list below this choice, but influences only particles in the current flow.

After choosing this option, click items in the list to highlight them.

Uniqueness group

The Uniqueness setting enables randomization of the *Range (page 2–174)* > Variation % calculations.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Script Wiring rollout

This rollout appears in the parameters panel below the main operator rollout after you highlight the operator, right-click it, and then choose Use Script Wiring. Thereafter, a check mark appears next to the Use Script Wiring in the right-click menu, and the rollout appears whenever you highlight the operator. To turn off script wiring, choose Use Script Wiring again from the right-click menu.

Script wiring lets you use a script to control the Force and Range parameters, which you normally specify in the operator's parameters. Place a *Script operator (page 2–208)* before the Keep Apart operator in the event, and then use it to define values in the particleFloat and particleVector channels. For an example of a script that sets particleFloat values, see *particleFloat Sample Script (page 2–206)*.

On the Script Wiring rollout, choose either of the following:

Use Script Float As—Choose either of the following:

- **Not Used—**Particle Flow uses the *Force setting* specified in the Parameters rollout.
- **Influence—**Particle Flow applies the script particleFloat value to the Force value.

Use Script Vector As—Choose one of the following. Choosing Absolute Size Range or Relative Size Range makes the Range settings on the Parameters rollout unavailable.

- **Not Used—**Particle Flow uses the *Range (page 2–174)* settings specified in the Parameters rollout.
- **Absolute Size Range—**Particle Flow applies the script particleVector values to the Absolute Size Range values. The X component of the vector value is used for the core radius, and the Y component for the falloff radius.
- **Relative Size Range—**Particle Flow applies the script particleVector values to the Relative Size Range values. The X component of the vector value is used for the core percent, and the Y component for the falloff percent.

Shape

Shape Operator

Particle View (page 2–125) > Click Shape in an event or add a Shape operator to the particle system and then select it.

Shape is the default operator for defining the geometry used in the particle system. You can use it to specify particles in the shape of pyramids, cubes, spheres, or vertices, as well as particle size.

Note: All primitive particles created with Shape have pivot points in their geometric centers.

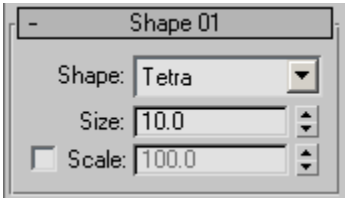
See also

Shape Instance Operator (page 2–178)

Shape Facing Operator (page 2–176)

Shape Mark Operator (page 2–183)

Interface



Shape—Sets the shape of the particles to Vertex or one of three geometric primitives. Default=Pyramid.

- **Vertex—**Use this option to create non-rendering particles whose sole attribute is their location. This can be useful when using script operators.
- **Tetra—**These tetrahedron-shaped particles have a four triangular sides. Having only four polygons, the tetra is the simplest of the rendering particle shapes.
- **Cube—**These box-shaped particles have six square sides, all the same size.
- **Sphere—**Creates spherical particles. The sphere has the most polygons of the rendering particle shapes, and its orientation isn't obvious.

Size—Sets the overall size of the particles in system units. Default=10.0

Setting Size=0 makes the particles invisible.

Scale—Turn on to set the size of the particles as a percentage of the Size value. Use the numeric field to set the Scale percentage value. Default=100.

Scale is animatable, while Size is not. Also, turning on Scale creates the scale data that any Scale operators further on in the flow can work with.

Shape Facing Operator

Particle View (page 2–125) > Click Shape Facing in an event or add a Shape Facing operator to the particle system and then select it.

The Shape Facing operator creates each particle as a rectangle that always faces a particular object, camera, or direction. For effects like smoke, fire, water, bubbles, or snowflakes, use Shape Facing with a material containing appropriate opacity and diffuse maps.

By default, the particles' top and bottom sides are parallel to the horizontal plane. The Orientation setting lets you change this default alignment.

Note: For the Shape Facing operator to be able to calculate the particle orientation, it must know the current particle position. For this reason, always place the Shape Facing operator *below* a Position operator in an event.

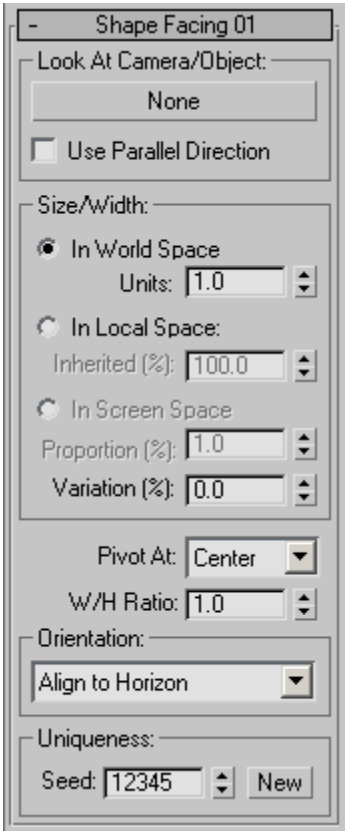
See also

Shape Operator (page 2–176)

Shape Instance Operator (page 2–178)

Shape Mark Operator (page 2–183)

Interface



Look At Camera/Object group

Use this control for defining the camera or object toward which the particle will face. This object is known as the *Look At object*. For the technically minded, the facing is maintained by keeping each particle's local Z axis pointed at the Look At object, or when Use Parallel Direction is on, aligned with the vector between the Particle Source icon and the Look At object.

Name—After picking a Look At object, its name appears here.

Pick Camera or Object—Click this button, and then select a camera or object in the scene to use as the Look At object. The mouse cursor changes to a cross shape when positioned over a valid object.

Use Parallel Direction—When off, all particles continuously rotate to stay facing toward the Look At object. Each particle's orientation varies because its location differs from those of the others. When on, all particles face in the same direction, defined by an imaginary line between the Particle Source gizmo center and the Look At object. Default=off.

Size/Width group

Use these settings to specify the coordinate system for setting the size, as well as the size parameters. The numeric settings in this group are not animatable.

In World Space—Sets the absolute size in system units, using the World coordinate system.

Units—With the In World Space option, sets the particle size in system units. Range=0 to 1000000000. Default=1.

In Local Space—Sets particle size relative to the existing size, in local space. The software uses the dimensions of the existing shape to determine the size of the "facing rectangle."

Inherited %—Sets the facing particles' size as a percentage of the existing size. Range=0 to 100. Default=100.

In Screen Space—Sets the facing particles' size as a percentage of the screen width. The actual size of each particle changes as necessary throughout the animation, depending its distance from the camera, to maintain a constant size from the camera's point of view.

This option is available only when the Look At object is a camera, and Use Parallel Direction is off.

Proportion %—Sets the particle size as a percentage of the screen width. Default=1.

Variation %—Sets the percentage by which particle size can vary. Default=0.

Pivot At—Specifies the part of the particle around which rotation is performed when maintaining the facing direction. The choices, available from the drop-down list, are Top, Center, and Bottom. Default=Center.

With Top and Bottom, the center of the corresponding side is used as the particle center for rotation.

This option is useful, for example, when particles are lying on a surface, and each particle is an explosion. In this situation, you'd probably want the entire particle rectangle to appear above the surface, so you'd set Pivot At to Bottom.

W/H Ratio—Defines the aspect (width-to-height) ratio of the shape rectangle.

Adjust this ratio to that of the particle-map imagery. For the commonly used TV aspect ratio of 4:3, set W/H Ratio=1.33. The parameter is not animatable. Range=0.001 to 1000. Default=1 (square particles).

This value is not animatable.

Orientation

Use this drop-down list to choose how particles rotate on the axes not specified by the Look At direction. Default=Align to Horizon.

Align to Horizon—Keeps the top edge aligned with the horizon (the world XY plane).

Align to Speed Follow—Uses the right side of each particle as the leading edge, so the top and bottom edges are aligned with the direction of particle motion.

Random—Orients the top edge at random.

Allow Spinning—To spin the particles, choose this option, and in the same event use a *Spin operator (page 2–154)*. In the latter, set Rotation Axis to Particle Space and use the default axis values: X=0, Y=0 , Z=1.

Uniqueness group

The Uniqueness setting enables changing the randomization of the size/width variation.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Shape Instance Operator

Particle View (page 2–125) > Click Shape Instance in an event or add a Shape Instance operator to the particle system and then select it.

Shape Instance lets you use any *reference object (page 3–1002)* in the scene as particles. You can effectively define only one reference object per event, but the object can comprise any number of sub-objects, each of which Particle Flow can treat as a separate particle. Also, using tests, you can divide the particle stream into multiple branches and define a different particle shape for each.

As the name of the operator indicates, reference geometry is instanced into the particle system. Thus, any physical changes you make to the original geometry are reflected instantly in the particle system. If you hide the original geometry, the particles still appear. However, if you delete the original geometry, the particles are deleted as well.

Using Transform-Animated Reference Objects

You can create animated particle shapes by animating the reference object with modifiers, and then turning on *Animated Shape*. However, Shape Instance ignores any rotational and positional transformations applied directly to the reference object; it uses only the pure scale component. For

example, if you scale an object non-uniformly using the View reference coordinate system, the result might skew the object's shape. Because the skewing is the by-product of a rotational transform, it wouldn't be reflected in the particles' shape. However, modifiers and comparable functions that contain transformations applied to the reference object are reflected in the instanced particles. For example, if you want particles to use rotational transforms applied to the reference object, use the Reset XForm function on the reference object. Reset XForm is available from the Utilities panel.

Tip: Particle Flow *can*, however, use rotational and positional animation applied to descendant objects in a hierarchy, when you use the entire hierarchy as a single particle. To do this, create a hierarchy, animate the descendant objects, and then designate the parent as the reference object. Do not turn on Separate Particles For > Object And Children. You'll find a procedure illustrating this, below.

When using an animated reference object, it is recommended that you hide the reference object after instancing it in the particle system.

Procedure

Example: To use the reference object's rotational and positional animation in the particle system:

1. Create an object you want to use as the particle shape, such as a teapot, and a Dummy helper object.
2. Animate the teapot using the Move and Rotate tools.
3.  Use the Align tool to center the teapot to the dummy.

   This step isn't necessary, but it can help to obtain more consistent results.
4.  Link the teapot as a child of the dummy (drag from the teapot to the dummy).

   You can use any object as the parent, but using a dummy, which doesn't render, lets you animate all instanced particles visible in the final animation.
5. Create a default Particle Flow system.
6. In Particle View, delete the Rotation operator.
7. Replace the Shape operator with a Shape Instance operator.
8. In the Shape operator parameters, designate the dummy as the Particle Geometry Object.
9. Turn on Animated Shape.
10. For sequential animation, which can be more interesting visually, in Animation Offset Keying, set Sync By to Particle Age or Event Duration.

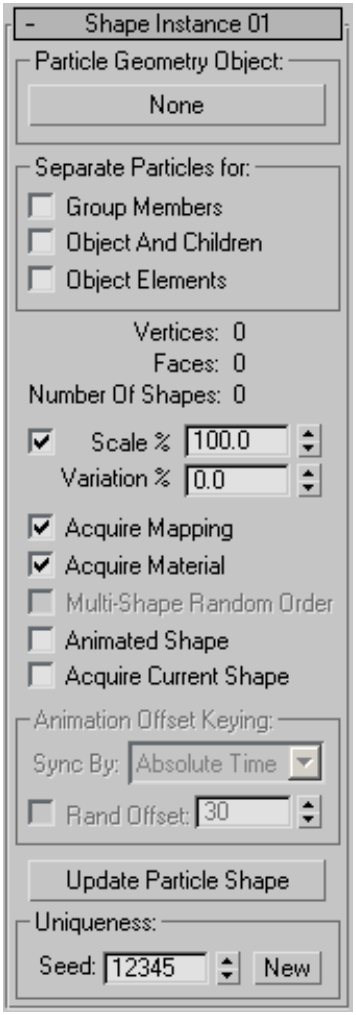
See also

Shape Operator (page 2–176)

Shape Facing Operator (page 2–176)

Shape Mark Operator (page 2–183)

Interface



Particle Geometry group

Use this control for defining the object to be used as the particle geometry. This object is known as the *reference object*.

[button]—Click this button, and then select a geometry object in the scene to be used as the particle shape. The mouse cursor changes to a cross shape when positioned over an object that's valid for use as a reference object. Valid objects include *multi-shape* objects such as groups and hierarchies; see *Separate Particles For group (page 2–181)*.

After picking a reference object, its name appears on the button.

You can use almost any geometry object as a reference object. By default, the operator automatically converts closed splines to rendering geometry by "filling in" the area defined by the shape outline. To use the shape outline instead, select the original shape and, on the Modify panel > Rendering rollout, turn on Display Render Mesh. You needn't turn on Renderable, but changes to all other Rendering rollout settings, such Thickness and Sides, are reflected in the rendered particles.

Particle Flow does not automatically "fill in" open shapes such as Line (non-closed) and Arc. To render these shapes when using them as particles, turn on Modify panel > Rendering rollout > Display Render Mesh.

If a Shape Instance operator is in effect for which the geometry object is not defined, and the viewport display type is set to Geometry, then the particles appear in the viewports as X characters.

Note: If a multi-shape object used as a reference object comprises objects both with and without materials applied, and Acquire Material is on, Particle Flow will apply the first available material to any objects without materials applied.

Note: You can use a NURBS surface as a reference object, but not a NURBS curve. To use a NURBS curve as particle geometry, first convert it to a surface using a method such as the Cap function.

Tip: For best results, when using an animated reference object, it's highly recommended that you hide the original object before rendering, or use Object Properties to turn off its Renderable check box.

Note: If using a group as a reference object, make sure the group is closed before selecting it.

Selecting an open group will add only the group member you click, not the entire group.

Separate Particles For group

3ds Max provides a number of methods for combining disparate objects into single entities, including grouping, building hierarchies with linking, and attaching. By default, when using such a *multi-shape* object as particle geometry, each particle comprises all member objects. Alternatively, you can instruct Particle Flow to treat each member object as a separate particle with these options. When on, the software uses each member object as a single particle, in left-to-right order according to its X-axis position, by default. So, for example, if you create text in the Front viewport, assign it as reference geometry, and turn on Object Elements, the letters come out in the proper order.

You can vary the order randomly by turning on *Multi-Shape Random Order*.

You can turn on any combination of Separate Particles items. All are off by default.

Group Members—When on, group members are treated as separate particles.

Object and Children—When on, linked objects are treated as separate particles.

Object Elements—When on, element sub-objects of a single mesh object are treated as separate particles.

Vertices/Faces—Show the number of vertices and triangles per particle. If the reference geometry is multi-shape and the appropriate check boxes in the Separate Particles are on, then these displays show an average vertex/face count.

Number of Shapes—Shows number of different particle shapes. This number is always 1 unless the reference geometry is multi-shape and the appropriate check boxes in the Separate Particles For group are on, in which case this displays the number of resulting different shapes.

Scale (%)—Specifies a uniform scaling factor for all particles. Range=0 to 100000. Default=on, 100.

The center of scaling for multi-shape object members treated individually depends on how the objects are combined. For grouped and linked objects, the scaling is done about the objects' pivot points. For object elements, the scaling center is the geometric center of each object; that is, the averaged location of all the vertices.

This value cannot be animated. To animate particle size, use the *Scale operator (page 1–441)*.

Variation (%)—Specifies a randomized percentage of scaling variation. Use the Uniqueness setting to change the randomization. Range=0 to 100. Default=0.

This value cannot be animated. To animate scaling variation, use the *Scale operator (page 2–156)* operator.

Acquire Mapping—When on, all mapping data from the reference object is transferred to particles. Default=on.

Acquire Material—When on, material data from the reference object is transferred to particles. Default=on.

If the reference object is a group, with different materials applied to the group members, the software creates a new Multi/Sub-Object material containing all of the materials and uses it as the particle material.

Tip: Because materials are "sticky" in Particle Flow, if you turn off Acquire Material after specifying a reference object with an attached material, the material remains applied to the particles. To avoid applying the instanced shape's material to the particles, turn off Acquire Material before specifying the reference object.

Multi-Shape Random Order—When on, assigns shapes to particles in random order. When off, Particle Flow emits each shape in the multi-shape object as a single particle, in the order of the shape's X coordinate. In other words, the shape with the lowest X-axis coordinate is emitted first, then the one with the next highest, and so on. The option is available only if at least one of Separate Particles For group check boxes is on. Default=off.

For example, if you want the particles to spell out a word or phrase, use extruded text created in the Front viewport as the reference object, turn on Separate Particles For > Group Elements, but leave Multi-Shape Random Order off.

Animated Shape—When on, particles use any animation in the reference object, including cyclic animation applied with the Parameter Curve Out-of-Range Types controls. When off, the particles are not animated.

With this option, you can use the Animation Offset Keying controls to specify how to synchronize the reference-object animation with the particles. See the next section.

Note: A particle that uses this option is animated only while in the event containing the Shape Instance operator. If it moves to another event that doesn't contain a shape operator, it keeps the same shape, but the animation stops. The easiest way to keep the animation going from event to event is to place the Shape Instance operator in the *global event (page 3–949)*. Otherwise, you need to place the operator in each event in which the particle should be animated.

Acquire Current Shape—When on, Particle Flow acquires the particle shape from the reference object as each particle enters the event. Thus, if the reference object is animated, particles that enter at different times get different shapes. However, these shapes are not animated. When the option is off, the shape is acquired from frame 0. Default=off.

This option is available only when Animated Shape is off.

Animation Offset Keying group

These controls are available only when Animated Shape is on.

Sync By—Lets you choose how to synchronize reference-object animation with the particles.

- **Absolute Time—**All particles have the same shape at any given moment.
- **Particle Age—**Animation of the reference object is synchronized with particle age: Frame 0 of the reference-object animation corresponds to the frame of each particle's birth.
- **Event Duration—**Frame 0 of the reference-object animation corresponds to the moment the particle enters the event.

Rand Offset—When on, randomly varies the start of each particle's animation. Use the numeric setting to specify the number of frames by which the animation start can vary.

Update Particle Shape—Refreshes instanced particle shapes from the reference object. Use this after changing an object in a hierarchical reference object.

In most cases, Particle Flow automatically updates instanced particles when changes are made to the reference object. However, in some cases, when you change an object deep in a hierarchy, the particle instances might not be updated. In such a case, click Update Particle Shape to refresh the instances.

Uniqueness group

The Uniqueness setting enables changing the randomization of the scale variation, animation offset, and the multi-shape random order.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Shape Mark Operator

Particle View (page 2–125) > Click Shape Mark in an event or add a Shape Mark operator to the particle system and then select it.

Use the Shape Mark operator to replace each particle with either a rectangle or a box cut out from the particle geometry with an image mapped onto it. The image can be animated and the animation can be synchronized with particle events.

A typical application of Shape Mark would be to leave marks after particles impact objects in the scene. For example, when a torpedo hits a boat and explodes, you could use Shape Mark to leave scorch marks on the boat surface.

Tip: By default, the mark left by Shape Mark is always rectangular, no matter what shape the particles are. To leave a differently shaped mark, choose the Shape group > Rectangle option, and use a *material (page 2–186)* in the same event with transparent areas to define the mark's outline. For example, apply a Mask map to the material's Opacity channel, and to the map's Mask channel, apply a bitmap containing an alpha channel. On the Bitmap Parameters rollout, set Mono Channel Output to Alpha.

Note: With its default settings, the Shape Mark operator can generate coplanar faces, which the mental ray renderer cannot render well. If you want to use mental ray rendering with Shape Mark, adjust the operator's settings as follows:

- In the Shape group, choose Box Intersection.
- Set Surface Offset to **0.01**.
- Set Offset Variation to **0.005**.
- Set Vertex Jitter to **0.002**.

With these settings, faces are no longer coplanar, and the mental ray renderer gives better results.

See also

Shape Operator (page 2–176)

Shape Instance Operator (page 2–178)

Shape Facing Operator (page 2–176)

Procedure

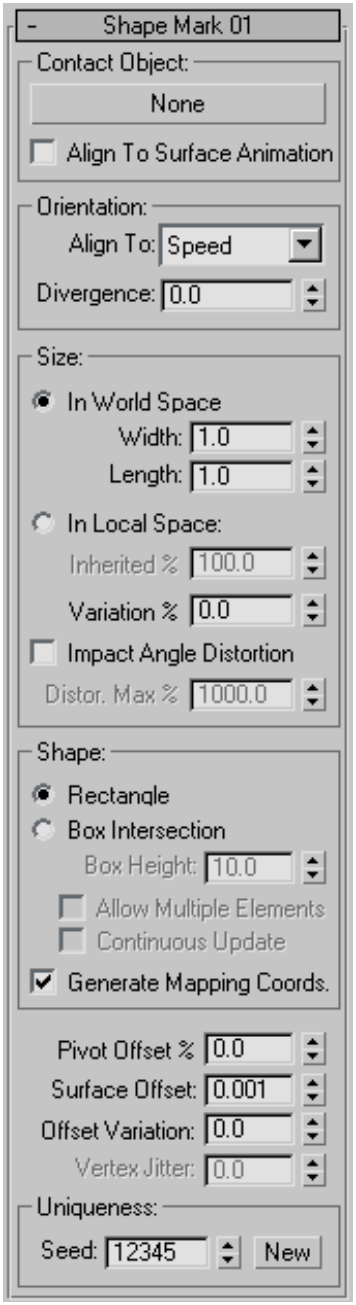
Example: To use Shape Mark:

1. Determine which object is to receive the marks; this will be the *contact object*. Apply a deflector to this object.
2. Set up your particle system with an event that causes particles to collide with the contact object deflector.
3. At the end of this event, add a *Collision test (page 2–212)*.
4. In the Collision test, designate the deflector from step 1.
5. Create a new event with the Shape Mark operator, and wire the Collision test to this event.
6. In the Shape Mark operator > Contact Object group, designate the object from step 1. Change the other Shape Mark settings as necessary.

   Now, when the particles strike the contact object, they disappear, leaving marks on the object.

   If you want the particles to bounce after leaving marks, rather than disappearing, use a *Collision Spawn test (page 2–215)* instead, and turn off its Delete Parent check box. The spawned particles become the marks, and the original particles remain in the first event.
7. Optionally, add a Material operator to define the surface characteristics of the marks.

Interface



Contact Object group

Use this control for defining the object on which marks are to be left.

[button]—Click this button, and then select an object in the scene to use as the contact object. The mouse cursor changes to a cross shape when positioned over a valid object.

After picking a contact object, its name appears on the button.

Align to Surface Animation—When on, Shape Mark takes into account surface changes due to vertex animation of the contact object. If *Shape (page 2–185)* is set to Rectangle, then the mark changes its orientation and position to appear be stuck to the surface of the contact object. If Shape is set to Box Intersection, then the mark changes shape along with that of the contact object. When off, only transformation of the contact object is taken into consideration. Default=off.

Turn this on only if there is significant vertex animation at the contact point, such as with an animated water surface.

Warning: This option requires significant CPU and memory resources.

Orientation group

Align To—Depending on the setting for *Shape (page 2–185)*, Shape Mark creates either a rectangle or a box cutout on the contact geometry. The Orientation setting specifies how the shape is oriented. In the mark's local coordinate system, the X axis is Length, the Y axis is Width, and with the box cutout, Z is height. The Z axis is perpendicular to the surface of the object at the contact point.

The alignment choices are as follows:

- **Speed—**The Length direction is parallel to the projection of the particles' speed vector onto the contact plane.

- **Particle X/Y/Z—**The Length direction is parallel to the projection of the particle's local coordinate axis X, Y, or Z, respectively, as the particle moves toward the contact object.
- **Random—**Uses a random Length direction in the contact plane.

Divergence—Applies a range of random variation, in degrees, to the orientation of the Length direction. Unavailable when using the Random option.

Size group

Use these settings to specify the coordinate system for setting the size of the mark, as well as the size parameters. The numeric settings in this group are not animatable.

In World Space—Sets the absolute size of the mark in system units, using the World coordinate system.

Width/Length—With the In World Space option, sets the particle dimensions in system units. Range=0 to 1000000000. Default=1.0.

In Local Space—Sets the mark size relative to the existing particle size, in local space. The software uses the dimensions of the existing shape to determine the size of the "facing rectangle."

Inherited %—Sets the percentage of the mark size, relative to the existing particle size. Range=0 to 100. Default=100.0.

Variation %—Sets the percentage by which particle size can vary. Default=0.0.

Impact Angle Distortion—When on, increases the Length value of the mark according to the particle's angle of approach. This effectively stretches the mark shape if particle approaches the contact geometry at a low angle. Available only when Align To is set to Speed. Default=off.

For example, if a drop of paint hits a surface perpendicularly, it creates a circular mark, but if it hits the surface at a lower angle, the resulting shape is an ellipse.

Distor(tion) Max %—Sets the maximum percentage by which the software may stretch the mark. Available only when Impact Angle Distortion is on. Default=1000.

With very low angles of approach, the stretching factor can become very high. For example, value 500% means that the stretching factor cannot exceed 5.

Shape group

These settings let you specify the mark-making object as a rectangle or a box. Default=Rectangle.

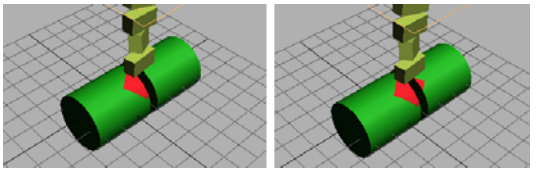
Rectangle—The mark shape is a two-faced rectangle.

When using a material with Shape Mark, always choose this option.

Box Intersection—With this option, the software creates a box for each particle that leaves a mark, and derives the mark shape from a Boolean intersection between the contact object and the box.

Box Height—Sets the height of the box used with the Box Intersection method. Available only with Box Intersection. Default=10.0.

Allow Multiple Elements—When on, particles can leave marks on all parts of contact objects that contain multiple elements. When off, a particle marks only the first element it collides with. Available only with Box Intersection. Default=off.



Particles falling onto a two-element cylinder

Left: Allow Multiple Elements is off; Right: Allow Multiple Elements is on.

Continuous Update—When on, the shape of the mark is recalculated at each frame, according to the current positions of the particle and the contact surface. This option can consume a great deal of CPU time. Available only with Box Intersection.

Generate Mapping Coords.—Allows correct application of the shape mark when using a mapped material. Default=on.

If you're not using a mapped material, you can save memory by turning this off.

Pivot Offset %—Shifts the position of the shape mark along its length dimension with respect to the pivot of the impacting particle. Default=0.0. Range=-50.0 to 50.0.

By default, the center of the mark's length dimension coincides with the point where the particle's pivot strikes the contact object. This setting lets you offset the mark's position to anywhere along its length.

Note: The width dimension and the Box Intersection's height dimension are always centered at the intersection of the particle pivot and the contact object's surface.

Surface Offset—Specifies the distance of the shape mark above the contact object's surface. Default=0.001.

The mark is slightly elevated above the contact geometry to achieve the visual effect of the mark spot overlapping the contact geometry. This parameter is not animatable.

Offset Variation—Specifies the maximum extent of a random variation in the actual surface offset among particles. Default=0.0.

Adjusting this value can help to alleviate rendering artifacts with overlapping marks.

Vertex Jitter—Specifies the maximum extent of a random variation in the positions of vertices of marks created using the Box Intersection method. Available only with the Box Intersection method. Default=0.0.

Adjusting this value can help to alleviate rendering artifacts with overlapping marks.

Uniqueness group

The Uniqueness setting enables changing the randomization of the size/width variation.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Material and Mapping

Particle Flow provides three operators for applying materials to particles. To give the same appearance to all particles throughout an event, use *Material Static Operator (page 2–187)*. If you're using a compound material such as *Multi/Sub-Object (page 2–1594)*, you can assign different sub-materials to different particles with the *Material Frequency operator (page 2–189)*. And to assign materials that change in appearance over time, use the *Material Dynamic operator (page 2–191)*.

Related to this is the *Mapping operator (page 2–195)*, which lets you give the same mapping coordinates to the entire surface of each particle in an event, thus using a single pixel from a material

to color the particles. By animating the mapping coordinates, you can cause the particles to change color over time. This is particularly effective with a gradient material.

The *Bitmap map (page 2–1631)*, used in conjunction with Material Dynamic, lets you assign different frames from an image sequence to particles based on the particle age, among other effects.

Following are some items to keep in mind when using materials with Particle Flow:

- A material is a static property of an event. It does not travel along with the particles from one event to the next. A particle's material ID does, but its material does not. If you want particles always to use the same material, define the material in the *global event (page 3–949)* with a Material operator or a *Shape Instance operator (page 2–178)*. Otherwise, you need to define it in each local event.
- The primitive particle shapes available with the *Shape operator (page 2–176)* do not have mapping coordinates. The *Mapping operator (page 2–195)* applies the same mapping coordinates to each particle's entire surface, so it's not suitable in this situation. If you want to apply image-based materials to particles, use *Shape Instance (page 2–178)* instead.
- If you use a Material operator with *Shape Instance (page 2–178)*, be sure to apply mapping coordinates to the *reference object(s) (page 3–1002)*. You can do this by making sure the object's Generate Mapping Coords option is on (if available), or applying a *UVW Map modifier (page 1–922)* or *Unwrap UVW modifier (page 1–878)*. If you don't apply mapping coordinates, the system generates a Missing Map Coordinates warning when you render the scene.
- If you use an object with a material already applied as a *reference object (page 3–1002)* for instanced particles, you don't need a Material operator in the same event. However, the material appears only in the event containing the *Shape Instance operator (page 2–178)*; it does not persist from event to event.
- If you use the *Cache operator (page 2–197)* with Update set to Always, toggling the Material Editor > Show Map In Viewport switch causes Particle Flow to recalculate the cache.
- You can drag a material from the Material Editor to a Particle Flow source icon, but the material will not have any effect on the system. You must use a Material operator or Shape Instance to apply materials to particles in Particle Flow.
- Once you've assigned a material to a Material operator, the material shows up in the Material Editor as "hot"; that is, triangles appear in the corners of its sample slot. However, because of the nature of the Particle Flow data structure, the Material Editor functions Select By Material and Get Material > Browse From Selected do not work correctly with Particle Flow systems. You can, however, use Get Material > Browse From Scene.

Material Static Operator

Particle View (page 2–125) > Click Material Static in an event or add a Material Static operator to the particle system and then select it.

The Material Static operator lets you give particles material IDs that remain constant throughout the event. It also lets you assign a material to each particle based on its material ID. The operator can assign the same material ID to all particles, or different IDs to successive particles on a cyclical or random basis. The most common usage of this latter capability is with a Multi/Sub-Object material, for applying a different material to each particle.

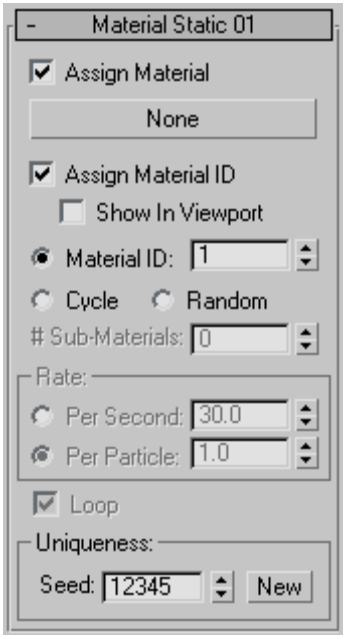
See also

Materials and Mapping in Particle View (page 2–186)

Material Frequency Operator (page 2–189)

Material Dynamic Operator (page 2–191)

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Assign Material—When on, the operator assigns a material to the particles. Default=on.

[button]—Use this button to assign a material to the operator. Click the button and then use the Material/Map Browser to choose the material. Alternatively, drag the material from a Material Editor sample slot to the button. After you assign a material to the operator, its name appears on the button.

Assign Material ID—When on, the operator defines a material ID number for each particle. Default=off.

Show In Viewport—When on, the material is shown applied to the particles in the viewports.

Assignment Method

The Material Static operator gives you a choice of three different methods for assigning Material IDs to particles:

Material ID—Assigns the same material ID to all particles. Use the numeric field to set the ID value.

Cycle—Assigns each particle a material ID in the range 1 to N, where N=# Sub-Materials, in increasing sequential order. The first ID assigned is 1, then 2, then ..., then N, then 1 again, and so on.

Random—Assigns each particle a random material ID in the range 1 to N, where N=# Sub-Materials.

 # Sub-Materials—The highest ID number assigned to particles using the Cycle or Random option.

In general, set this to the same number of sub-materials in the Multi/Sub-Object material. If you set it to a smaller number, the operator will use only that many sub-materials, starting with the first and counting upward.

Note: The software automatically sets this value to the number of sub-materials in the material, *once only*, the first time you apply the material to the operator. Any subsequent changes in the material itself, or applying a different material to the operator, will not change or update the setting.

Rate group

These settings let you choose the basis on which the operator changes material ID assignments, and specify the rate of change.

Per Second—Sets the number of times per second that the assigned material ID is incremented.

If this value is the same as the rate at which particles enter the event, then one ID is assigned per particle. If it's lower, then multiple particles are given the same ID, or if it's higher, then the software increments the assigned ID faster than 1 per particle.

For instance, if particles enter the event at intervals of 1/60 of a second, and Per Second=30, then each pair of particles will be assigned the same ID. Or if particles enter the event at intervals of 1/15 of a second, and Per Second=30, then the ID is incremented (or changed randomly) twice per particle.

Per Particle—Sets the number of particles that must appear before material ID assignment changes. For example, If you set Per Particle=3, the material ID changes every three particles.

If you set Per Particle to a number less than 1.0, the software then moves through the sub-material list more rapidly than one (or more) particle per ID. That is, the software divides this value into 1.0, and adds the result to the current material ID to obtain the next one. For example, with eight sub-materials, if you set Per Particle=0.33, and use the Cycle option, the following series of IDs will result: 1, 4, 7, 2, 5, 8, 3, 6, 1, ... In general, this option is useful only with the Cycle option.

Loop—When on, and the last ID has been assigned, the software loops back around to the first ID and continues the cycle. When off, the software assigns the last cycle ID to all subsequent particles. Available only with the Cycle assignment method. Default=on.

For example, say you want the first eight particles that enter the event to use different materials, and all subsequent particles to use a ninth material. To do so, you would create a nine-sub-material Multi/Sub-Object material and assign it to the Material Static operator. Turn on Assign Material ID, choose the Cycle assignment method, and set # Sub-Materials=9. For Rate, use the default settings of Per Particle and 1.0. Lastly, turn off Loop.

Uniqueness group

The Uniqueness setting varies the sequence of assigned IDs with the Random option.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Material Frequency Operator

Particle View (page 2–125) > Click Material Frequency in an event or add a Material Frequency operator to the particle system and then select it.

The Material Frequency operator lets you assign a material to an event, and specify the relative frequency with which each sub-material appears on the particles. Typically, the material is a Multi/Sub-Object or other compound material, and you specify the frequency by setting a percentage for each of up to 10 different sub-materials (or material ID). The software assigns IDs to particles in a random sequence, based on these percentages. You can also use other materials that use sub-materials, such as Double Sided and Top/Bottom.

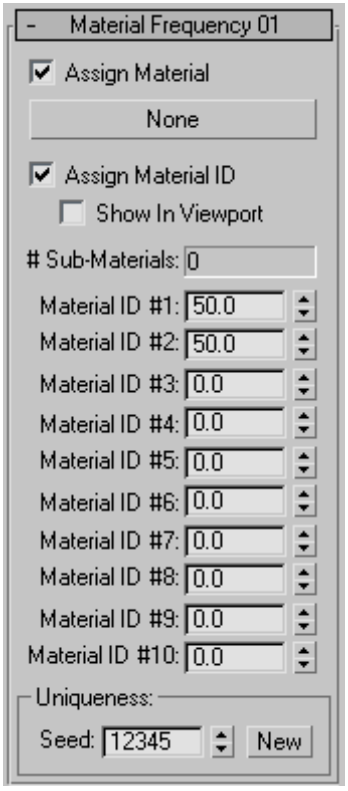
See also

Materials and Mapping in Particle View (page 2–186)

Material Static Operator (page 2–187)

Material Dynamic Operator (page 2–191)

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Assign Material—When on, the operator assigns a material to the particles. Default=on.

[button]—Use this button to assign a material to the operator. Click the button and then use the Material/Map Browser to choose the material. Alternatively, drag the material from a Material Editor sample slot to the button. After you assign a material to the operator, its name appears on the button.

Assign Material ID—When on, the operator defines a material ID number for each particle, and enables the remaining parameters. Default=on.

In general, this should remain on. Particle Flow uses the material ID with compound materials to know which sub-material to assign to a particle.

Show In Viewport—When on, the material is shown applied to the particles in the viewports.

 # Sub-Materials—Displays the number of sub-materials in the assigned material.

Material ID #1–10—Specifies the relative likelihood of particles to be assigned the corresponding material ID. Assign values for all IDs, or sub-materials, in the material that you want to have applied to the particles. So, for example, with a Multi/Sub-Object material containing five sub-materials, you'd set values for Material IDs #1-5.

This value is not absolute, but relative to the other settings. To follow the previous example, if you wanted all five materials to appear with equal frequency, you'd set the same non-zero value for Material IDs #1-5; the actual value wouldn't matter. On the other hand, if you wanted the materials to appear with decreasing frequency, you'd set the lower Material ID settings to relatively higher values; say 100, 80, 50, 33, and 10. In this case, each particle would twice as likely to be assigned material ID 1 as it would material ID 3, and one-tenth as likely to be assigned ID 5 as it would ID 1.

The actual sequence of material ID assignments is random, and can be varied by changing the Uniqueness Seed setting.

Uniqueness group

The Uniqueness setting varies the random sequence of assigned IDs.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Material Dynamic Operator

Particle View (page 2–125) > Click Material Dynamic in an event or add a Material Dynamic operator to the particle system and then select it.

The Material Dynamic operator lets you give particles material IDs that can vary during the event. It also lets you assign a different material to each particle based on its material ID. When used with an *animated texture (page 3–908)*, this lets you assign a different frame or map to each particle based on its total age or the amount of time it has spent in the current event. In this context, one example of an animated texture is a material that uses a multi-frame bitmap, such as an AVI file, as the Diffuse *Bitmap map (page 2–1631)*. Other examples of animated textures are materials that use the *Particle Age map (page 2–1675)* or the *Particle MBlur map (page 2–1676)*. Alternatively, you can use different sub-materials from a compound material such as Multi/Sub-Object.

Note: When using Object Motion Blur, if an event contains a Material Dynamic operator that uses a material with a Particle Age/MBlur/Bitmap map, the event should not also contain a Delete operator, or a Spawn or Collision Spawn test. Also, the event should not contain any tests that are wired to another event. The only exception to this is the Age Test operator when set to Absolute Time without any variation; that is, all particles leave the event at the same time. This applies to the use of Object Motion Blur only; there are no restrictions with Image Motion Blur.

See also

Materials and Mapping in Particle View (page 2–186)

Material Static Operator (page 2–187)

Material Frequency Operator (page 2–189)

Procedure

Example: To assign animation frames to particles based on age:

1. Prepare a file to be used as the animated bitmap. This can be a multiple-frame format such as AVI, or an image file list (IFL file) that points to a sequence of still images. With the latter, you can use images in a format such as Targa that contain predefined alpha channels to specify particle opacity selectively.
2.  Open the Material Editor, and assign a Bitmap map as the Diffuse map.
3. Use the Select Bitmap Image File dialog to assign the file from step 1 as the bitmap. This dialog appears automatically when you first assign a Bitmap map; alternatively, click the Bitmap button on the map's Bitmap Parameters rollout.
4. On the map's Time rollout, turn on Sync Frames To Particle Age.
5. If you want to use the image background or alpha channel to define transparency, on the Maps rollout, copy this map to the Opacity slot, and set the parameters accordingly.
6.  At the material level, turn on Show Map In Viewport.
7. Add a camera to the scene and set it up as desired. Activate the Perspective viewport and press the C key to set the viewport to show the camera view.
8. Create a default Particle Flow system.
9. Open Particle View.
10. Replace the Shape operator with a Shape Facing operator.

    Using this operator makes it easier to see the animation.

11. Click the Shape Facing operator, and in its rollout, click the Look At Camera/Object button, and then select the camera.
12. In the Size/Width group, increase In World Space > Units to about **15**.
13. Add a Material Dynamic operator to Event 01, and assign it the material from the beginning of this procedure.
14. Turn on Show In Viewport.
15. In the Animated Texture group, make sure Same As Particle ID is chosen, and turn on Reset Particle Age.

    Always choose Same As Particle ID with an animated texture, and one of the Sub-Material Rotoscoping options when using a compound material.

    Turning on Reset Particle Age causes the software to set particles to age 0 as they enter the event. In this example, the particles are born in the event, so technically it's not necessary to turn on Reset Particle Age. However, it's a good habit to get into to ensure that the animation always plays from the first frame.
16. Play the animation.

    In the viewport, the bitmap animation advances as the particles move, but all the particles show the same frame, regardless of age. This is anomalous behavior related to the limitations of viewport interactivity. The particles render correctly, however.
17. Render the animation.

    As each particle is born, it begins displaying the animation from the first frame. At each frame, each particle's age is incremented, and it displays the next frame from the applied map.

Example: To use the Particle Age map:

The *Particle Age map (page 2–1675)* applies up to three different colors or maps to particles throughout their life span, gradually changing from one to the next as the particles age. This effect can be used, for example for sparks flying from a fire: At first they're yellow; then, as they cool down, they turn red, and finally they become gray ashes. In order for Particle Age to know how far a particle has progressed through its life span, the particle has to be given a finite life. You do this using the *Delete operator (page 2–146)*.

1. Start or reset the program, and add a Particle Flow system.
2. Open Particle View and the Material Editor. Position them side by side.
3. In Particle View, add a Material Dynamic operator and a Delete operator to Event 01.

   For Particle Age to work, the Delete operator must be in the same event as the Material Dynamic operator. Alternatively, you can add the Delete operator to the *global event (page 3–949)* so that it affects every event.
4. Click the Delete operator, and in the parameters panel, choose By Particle Age, and set Life Span=**100** and Variation=**0**.

   This gives each particle a life span of 3 1/3 seconds.
5. Click the Material Dynamic operator.
6. In the Material Editor, assign a Particle Age map as the Diffuse map. On the Particle Age Parameters rollout, set three different colors, such as red, green, and blue. Also change the Age percentage values as necessary. For example, if you want each particle to show the second color a third of the way through its life instead of halfway, change Age #2 to **33**.
7. Drag the active material from its sample slot to the material button on the Material Dynamic parameters rollout in Particle View. When the Instance (Copy) dialog appears, click OK to accept the default choice: Instance.
8. In the Material Dynamic parameters, make sure Assign Material ID is on.

If it isn't, the particles all change color at the same time.

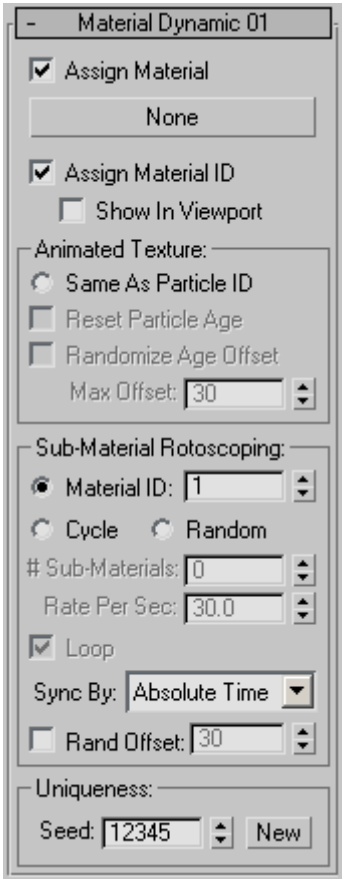
There's no need to turn on Show In Viewport; the Particle Age map doesn't appear in the viewports.

9. Render the animation, or a few representative frames.

   As each particle falls, it gradually changes color, with the oldest particles changing first.

   Tip: You needn't actually delete the particles to use this method. There are several ways to avoid this. You could set Life Span to a higher number than the length of the animation, and then, in the Particle Age map parameters, lower the Age #2 and Age #3 settings. Or, if you're using a local Delete operator, you could use an *Age test (page 2–211)* to move the particles into another event just before they're scheduled to be deleted. In that case, to avoid an abrupt color change, you might want to add to subsequent events a *Material Static operator (page 2–187)* with a material that uses the same final color or map as the Particle Age map.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Assign Material—When on, the operator assigns the specified material (see next parameter) to the particles. Default=on.

[button]—Use this button to assign a material to the operator. Click the button and then use the Material/Map Browser to choose the material. Alternatively, drag the material from a Material Editor sample slot to the button.

After you assign a material to the operator, its name appears on the button.

Assign Material ID—When on, the operator defines a material ID number for each particle. Default=on.

In general, this should remain on. Particle Flow uses the material ID with the Particle Age map to find the particle's other properties, including its life span and current age. And it uses the material ID with compound materials to know which sub-material to assign to a particle.

Show In Viewport—When on, the material is shown applied to the particles in the viewports when the particles are displayed as geometry.

Assignment Method

The Material Dynamic operator lets you assign Material IDs to particles in several different ways, depending on whether you're using an animated texture or a compound material. Default=Same As Particle ID.

Animated Texture group

Same As Particle ID—Assigns the same material ID to a particle as its particle ID. Choose this when using a material containing an animated texture, such as a Bitmap, Particle Age, or Particle MBlur map.

Note: Particle Flow assigns Particle IDs consecutively to particles at birth, starting with 0. Although the highest possible Particle ID is over 2,000,000,000, the highest possible *material ID* is 65535. Thereafter, the numbering sequence starts again at 0. Thus, when using a Particle Age map in a material assigned to the Material Dynamic operator, for best results, use a total of 65,536 particles or fewer.

Reset Particle Age—When on, sets each particle's age to 0 when it enters the event.

When using an animated material with a Bitmap map, turn this on to ensure that the animation always plays from the first frame.

Randomize Age Offset—When on, the software varies the difference between the particle age and the starting material ID at random. The maximum difference is determined by the Max Offset parameter.

Max Offset—The maximum number of frames by which the software can randomly vary particle age.

Sub-Material Rotoscoping group

These settings let you choose the basis on which the operator changes material ID assignments when using a compound material such as Multi/Sub-Object, and specify the rate of change.

Material ID—Assigns the same material ID to all particles. Use the numeric field to set the ID value.

Cycle—Assigns each particle a material ID in the range 1 to N, where N=# Sub-Materials, in increasing sequential order. The first ID assigned is 1, then 2, continuing to increment each by 1 until N; then 1 again, and so on.

Random—Assigns each particle a random material ID in the range 1 to N, where N=# Sub-Materials.

 # Sub-Materials—The highest ID number assigned to particles using the Cycle or Random option.

In general, set this to the same number of sub-materials in the Multi/Sub-Object material. If you set it to a smaller number, the operator will use only that many sub-materials, starting with the first and counting upward.

Note: The software automatically sets this value to the number of sub-materials in the material, *once only*, the first time you apply the material to the operator. Any subsequent changes in the material itself, or applying a different material to the operator, will not change or update the setting.

Rate Per Sec(ond)—Sets the number of times per second that the assigned material ID is incremented. If this value is the same as the rate at which particles enter the event, then one ID is

assigned per particle. If it's lower, then multiple particles are given the same ID, or if it's higher, then the software increments the assigned ID faster than 1 per particle.

For instance, if particles enter the event at intervals of 1/60 of a second, and Per Second=30, then each pair of particles will be assigned the same ID. Or if particles enter the event at intervals of 1/15 of a second, and Per Second=30, then the ID is incremented (or changed randomly) twice per particle.

Loop—When on, and the last ID has been assigned, the software loops back around to the first ID and continues the cycle. When off, the software assigns the last cycle ID to all subsequent particles. Available only with the Cycle assignment method. Default=off.

For example, say you want the first eight particles that enter the event to use different materials, and all subsequent particles to use a ninth material. To do so, you would create a nine-sub-material Multi/Sub-Object material and assign it to the Material Dynamic operator. Choose the Cycle assignment method, and set # Sub-Materials=**9**. Set the Rate Per Sec value to the rate at which particles enter the event. Lastly, turn off Loop.

Sync By—Choose the time frame for applying animated parameters. Available only with the Material ID and Cycle options.

- **Absolute Time—**Any keys set for parameters are applied at the actual frames for which they're set.
- **Particle Age—**Any keys set for parameters are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for parameters are applied to each particle starting when it first enters the event.

Rand Offset—When on, the software varies the difference between the particle age and the assigned material ID at random. The maximum difference is determined by the numeric parameter. Available only with the Material ID and Cycle options.

Uniqueness group

The Uniqueness setting varies the sequence of assigned IDs with the Random option, and the offset with the Randomize Age Offset option.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Mapping Operator

Particle View (page 2–125) > Click Mapping in an event or add a Mapping operator to the particle system and then select it.

The Mapping operator lets you assign a constant UVW mapping to the entire surface of the particles. It works in conjunction with a map specified in a material operator in the current event. By animating the mapping coordinates, you can vary the location on the material map from which the particle color is taken, thus changing the particle color in a predictable way over time.

The Mapping operator was designed primarily to be used with gradient maps, although you can use it with any map you like. The procedure below describes a recommended method for using the Mapping operator.

See also

Material Static Operator (page 2–187)

Material Frequency Operator (page 2–189)

Material Dynamic Operator (page 2–191)

Procedure

Example: To animate particle coloring using the Mapping operator:

1.  Open the Material Editor, and create a material that uses Gradient Ramp as the Diffuse map.
2. On the map's Gradient Ramp Parameters rollout, make sure Gradient Type is set to Linear.
3. Create a colorful gradient. The Mapping operator uses the colors in left-to-right order as the U value increases from 0.0 to 1.0.
4. Add a Particle Flow Source object to the scene.
5. Open Particle View.
6. In Event 01, click the Speed operator and set Speed to **100**.
7. In Event 01, click Display and set Type to Geometry.
8. Add a Material Static operator to Event 01, and then click the operator in the event.
9. Drag the material you created from its Material Editor sample slot to the button (labeled "None") on the Material Static parameters rollout in Particle View.
10. Add a Mapping operator to Event 01, and then click the operator in the event.
11. On the Mapping rollout in Particle View, turn on Show Map In Viewport.
12. 

Go to frame 100 and turn on Auto Key.
13. On the Mapping rollout in Particle View, set Map Values > U=**1.0**.
14. For Sync By, choose Particle Age.
15. 

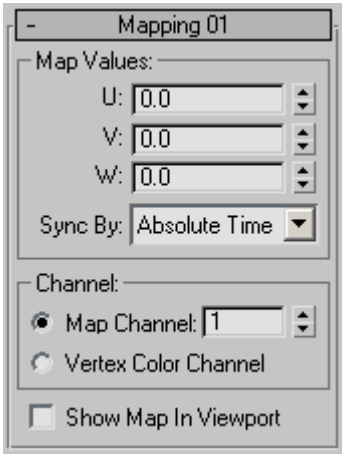
Turn off Auto Key, and click Play Animation.

As each particle is born and falls, its U mapping coordinate gradually changes from 0.0 to 1.0, while its color changes to match the corresponding position across the Gradient Ramp map.

To vary the effect, try changing the Gradient Type setting, and animate the V and W values as well. Or, for more sophisticated effects, assign animation controllers such as Noise Float to the Mapping coordinate values, or wire them to other changing parameters in the scene (Animation menu > Wire Parameters > Parameter Wire Dialog).

Tip: In the Material Editor, it's easier to see the result of the Gradient Type choice if you set Sample Type to Cube.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Map Values group

U/V/W—Set the coordinates on the map from which the particle color is taken. One pixel from this location is used to color the entire particle.

In general, U and V represent the horizontal and vertical dimensions, respectively, beginning at the

upper-left corner of the map, and W is useful only with three-dimensional maps, in which context it specifies depth.

Sync By—Choose the time frame for applying animated parameters. For further information, see *Animation Offset Keying group (page 2–158)*.

- **Absolute Time—**Any keys set for parameters are applied at the actual frames for which they're set.
- **Particle Age—**Any keys set for parameters are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for parameters are applied to each particle starting when it first enters the event.

Channel group

The Channel setting determines whether the mapping is applied to a map channel or to a vertex color channel.

Map Channel—The mapping is applied to the specified map channel.

Vertex Color Channel—The mapping is applied to the particles' vertex color channel.

Show Map In Viewport—The map coloring is visible in the viewports.

Additional Operators

Cache Operator

Particle View (page 2–125) > Click Cache in an event or add a Cache operator to the particle system and then select it.

The Cache operator records and stores particle states to memory. When it's in effect, the first time you play or go to a frame, the particle motion up to and including the frame is calculated and recorded in the cache. Thereafter, playing the frame or any previous frames uses the data in the cache, rather than recalculating the particle action. This is particularly useful with large or complex particle systems, where playback, and particularly backtracking, are slow because of the large amount of processing required. With caching on, the particle system needs to be calculated only once for each frame, and thereafter playback and moving among frames is significantly faster.

To cache an entire flow, add a Cache operator to the global event. To cache an individual local event, add a Cache operator to it.

To determine how best to implement caching in your particle system, turn on Particle View > Track Update > Update Progress, and then play the animation or go to a frame relatively distant from the current frame. Observe the particle diagram, and watch for events with actions that highlight longer than a fraction of a second, or events with more activity than the rest. Any such events could benefit from caching.

Note: Use no more than one Cache operator per event, unless you're using two: one set up for viewports, and the other for rendering. Similarly, don't use global and local Cache operators in the same flow, unless one is set up for viewports, and the other for rendering.

Tip: When using real-time playback, if playback isn't fast enough, you can use the Cache operator to speed it up. For best results, set the Viewport integration step to the same value as the real-time playback speed, and set cache sampling to Integration Step. For example, if the real-time playback speed is set to 1/4x, then the Viewport Integration Step for viewport should be also set to 1/4 Frame, and the Cache operator Sampling parameter should be set to Integration step. Thus, because real-time playback is showing four

times as many frames per second, and the Cache operator has stored four times as many frames per second, the two are synchronized. If the real-time playback speed is set to 2x or 4x, set the Viewport Integration Step to Frame.

Procedure

Example: To use the Cache operator:

1. Start or reset the software, and then add a Particle Flow system.
2. Open Particle View.
3. In Particle View, go to Options menu > Track Update and turn on Update Progress.

   This lets you monitor calculation of the particle system by highlighting actions in Particle View as the system executes them.
4. In Particle View, Add a *Keep Apart operator (page 2–172)* to Event 01.

   Keep Apart is a fairly calculation-intensive operator.
5. Go to frame 100 by clicking the right end of the time slider track.

   There is a delay as the system calculates all particle motion between frames 0 and 100. This is necessary because Particle Flow is a history-dependent system. At the same time, the actions in the system highlight briefly in Particle View as they're executed at each frame.
6. Go to frame 50 by clicking the center of the time slider track.

   There is a another delay as the system calculates all particle motion between frames 0 and 50.
7. In Particle View, drag a Cache operator from the depot to Event 01. Insert it anywhere in the event.

   When you release the mouse button, there is a delay as the Cache operator automatically caches particle motion from the start of the animation to the current frame. Meanwhile, each action highlights briefly at each frame.
8. Go to frame 20 or so.

   The only operators that highlight are Cache (very fast) and Display; there is no delay for calculation. All frames between 0 and 50 have been cached.
9. Go to frame 100.

   The delay this time is a result of caching particle motion for frames 50 to 100.
10. Jump to different frames, and drag the time slider.

    All particle motion is now cached, so no delays occur. But if you change a setting's value, the Cache operator automatically recalculates and stores the particle motion.
11. Go to frame 100, and then, in Particle View, click the Keep Apart operator and use the keyboard to change the Falloff Zone value to **8.0**.

    By default, the Cache operator automatically updates the cache when you change any parameters it stores. Thus, when you change the Falloff Zone value at frame 100, it recalculates and recaches the entire animation.

    Next, you'll briefly explore how manual caching works.
12. Click the Cache operator and set Update to Manually.

    The Cache operator no longer updates the stored data automatically when you change a parameter.
13. Click the Keep Apart operator, change the Falloff Zone value to **9.0**, and then drag the time slider.

    There is no delay, because the animation is still playing back from the cached data. However, the cached data is now invalid, because you changed a parameter in the particle system.

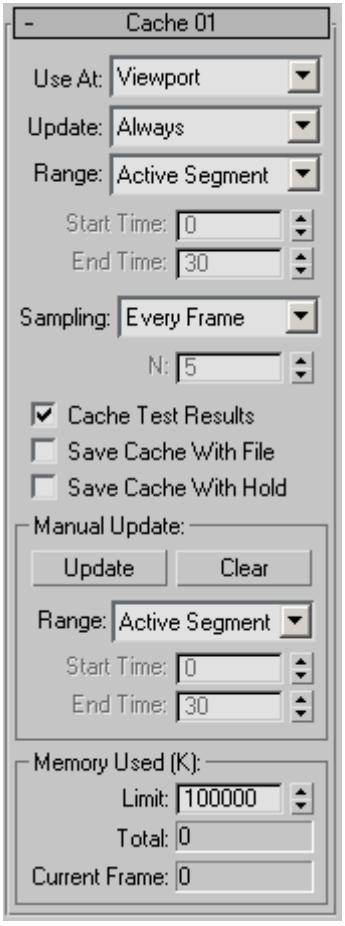
14. Click the Cache operator, and in the Manual Update group, click Update.

    The software closes Particle View, updates the cache for the active segment, and then reopens Particle View. The cached data is now accurate.

    Other manual update options let you update the cache for the entire animation or a custom frame range.

    As you can see, the Cache operator is quite powerful. When present and active, by default it overrides recalculation of the particle system, except when you change any action parameters, whereupon it automatically updates the cached data. Used appropriately, it can save a good deal of time in setting up and testing particle systems.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Use At—Caches particle motion when playing back in the viewports, or at render time, or both. Default=Viewport.

Important: Choose the Viewport/Render option only when using the same number of particles in the viewports and for rendering (see *Quantity Multiplier group (page 2–137)*). Otherwise, unpredictable results can occur.

Update—If you change a parameter in the particle system, the cached data might become invalid.

This setting determines whether Particle Flow should update the cache automatically, or let you do it manually. Default=Always.

- **Always—**Changing any parameter under the Cache operator's purview causes it automatically to update stored data from the start of the current range (see next parameter) to the current frame.
- **Manually—**Updating occurs only when you click the Update button, or clear the cache and then move to a different frame.

Range—Sets the frame range within which the Cache operator operates. Default=Active Segment.

- **Complete—**The software caches the entire animation.
- **Active Segment—**The software caches only frames in the active segment, as defined by the Start Time and End Time settings on the Time Configuration dialog. This is the frame range shown on the track bar. You can also change the active segment by holding down Alt and Ctrl and dragging the track bar with the left, middle, or right mouse button.
- **Custom—**The software caches only frames in the custom range, as defined by the Cache operator's Start Time and End Time settings (see following).

  Note: If you cache only part of the animation, Particle Flow calculates particle behavior in subsequent, non-cached frames using the cached data. For example, if you cache frames 0 to 50, and then jump to frame 60, Particle flow will calculate frame 51 based on the cached data, and frames 52 to 60 based on each previous frame.

Start/End Time—The first and last frames of the range considered for caching when Range=Custom (see above). Default=0, 30.

Note: The frame range time frame is in absolute time; that is, in terms of the entire animation. If you use a Cache operator locally, and specify a frame range during which no particles are present in the event, Particle Flow won't use the cache.

Sampling—Determines how often the Cache operator samples and caches the animation. Default=Every Frame.

- **Every Frame—**The software caches animation data once per frame.
- **Integration Step—**The software caches animation data at each integration step, using the Integration Step setting as specified in the Use At setting (see above). If Use At is set to Viewport/Render, it uses lower of the two Integration Step values. For example, if Viewport is set to Half Frame, and Render to 1/8 Frame, the sampling rate would be eight per frame.
- **Every Nth Frame—**The software caches animation data at frame intervals specified by the N value, below.

N—Determines the frame interval for caching when Sampling (above) is set to Every Nth Frame. Default=5.

For example, with N set to the default value of 5, the cache stores animation data for every fifth frame.

Cache Test Results—When caching particle data, this caches the results of test actions as well. Default=on.

This is important if Cache is used as a local operator, and the next event doesn't have a Cache operator. For the next event to work properly, it should receive particles from the current event. Those particles result from the activity of a test action. The Cache operator can record the test activity to play it back later.

If the Cache operator is used as a global operator, there is no need to cache the test results. This is because the system has cache data for every event,

and is able to jump to an arbitrary frame without the need for test results.

Save Cache with File—When on, the software Includes the cached data with scenes that you save to disk. This can significantly increase the size of saved files, but saves the time of recalculating the particle motion upon reloading the file. Default=off.

Normally, the cache data is saved only in disk files that you create with the Save or Save As commands. You can also instruct the software to include cached data with other types of files using the two following options.

Save Cache with Hold—Saves cached data in the Hold file, created with Edit menu > Hold. Default=off.

Manual Update group

These controls let you update or clear the cache manually within a frame range, or clear the entire cache.

Update—Calculates particle motion within the current range (defined below) and stores it in the cache, replacing any existing cached data. Available only when Range is set to Active Segment or Custom.

During the update, the software temporarily closes Particle View and displays a Cache Update Progress bar, while moving the time slider through the cached range. At the same time, by default, the animation plays in the viewports. To prevent this, and speed calculation, turn off Update Viewports.

If the cache runs out of memory during a manual update, Particle Flow halts the update operation and displays an alert. Click OK to continue, and then, if possible, increase the Limit value before updating the cache again.

Clear—Deletes any cached data.

After you clear the cache buffer, if you go to a different frame or change any action parameters, the software automatically caches particle motion even if Update is set to Manually.

Range—Sets the frame range within which the Cache operator recalculates data when you click Update. Default=Active Segment.

- **Complete—**The software *clears* the cache for entire animation. Choosing Complete makes the Update button unavailable; this option is for clearing the cache only.

  To update the entire animation, set the active segment or the custom range to encompass all frames, and then use the corresponding choice and click Update.

- **Active Segment—**The software updates the cache only for frames in the active segment, as defined by the Start Time and End Time settings on the Time Configuration dialog. This is the frame range shown on the track bar. You can also change the active segment by holding down Alt+Ctrl and dragging the track bar with the left, middle, or right mouse button.

- **Custom—**The software updates the cache only for only frames in the custom range, as defined by the Start Time and End Time settings (see following).

Start/End Time—The first and last frames of the range that's updated when Range=Custom (see above). Defaults=0, 30.

Update Viewports—When on, the animation plays in the viewports during manual updating of the cache. Turn this off to disable playing the animation in the viewports during manual caching; this can speed up the caching process, especially with large or complex particle systems. Default=on.

Memory Used (K) group

The Cache operator stores data in system memory; you can specify an upper limit for the amount of memory it uses. If the Limit setting and the amount of cached data exceeds the available free memory, the computer system might use virtual (hard disk-based) memory instead, which slows down the caching. If Particle Flow fills the cache, any remaining frames are calculated on the fly.

This group also lets you monitor the amount of memory used for caching data.

Limit—The maximum amount of system memory used to cache particle data, in kilobytes. Default=100,000, or 97.6 MB.

Total—The amount of memory currently used by the cached data, in kilobytes. Read-only.

Note: Even animation frames with no particles will probably consume a certain amount of cache memory. The reason for this is that the cache also stores states for randomly calculated values such as Variation, to ensure that particle activity is consistent across a rendering network, and with machines that might not have regular access to all frames.

Current Frame—The amount of memory used by the data cached for the current frame, in kilobytes. Read-only.

Display Operator

Particle View (page 2–125) > Click Display in an event or add a Display operator to the particle system and then select it.

The Display operator lets you specify how particles appear in the viewports. The default display mode is Ticks, which is the simplest, and thus the fastest to display. It's useful for animations that use a large number of particles. At the opposite end of the complexity spectrum is the Geometry option, which lets the software depict particles as their actual shapes. In addition, the Display operator provides a variety of simple shapes that provide fast feedback in testing animation, as well as the ability to easily distinguish among particles in different events. It also lets you set the percentage of visible particles.

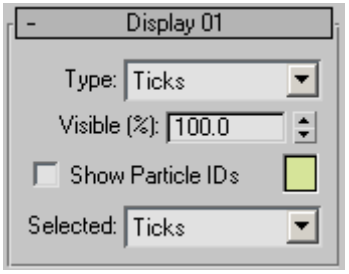
By default, Particle Flow automatically inserts a new Display operator in each local event you add to the system. Alternatively, you can choose Particle View > Options menu > Default Display > Global. With this option, the software automatically inserts a Display operator in new global events, but does not add one to new local events.

Particles in any local event that doesn't contain a Display operator don't appear in the viewports, unless an associated global event contains a Display operator (that is, a global Display operator). If a particle is affected by multiple Display operators (for example, both global and local operators) simultaneously, the software generates all viewport particle representations at the same time.

See also

Render Operator (page 2–206)

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Type—Choose how particles appear in the viewports. The two-dimensional particle markers depict position only. Geometry shows how

particles will actually render, in three dimensions; Lines shows speed and direction of motion; and Bounding Boxes reflects scale and orientation.

In the following list, each choice's description is preceded by the number of dimensions used by each particle representation:

- **None—**Particles do not appear in the viewports.
- **Dots—**(0D) Each particle appears as a single pixel.
- **Ticks—**(2D) Each particle appears as a + sign.
- **Circles—**(2D) Each particle appears as a small circle.
- **Lines—**(1D) Each particle appears as a line, one pixel thick. The line's length indicates particle speed, and its orientation reflects the direction of motion. Use this option for fast, accurate feedback when experimenting with the Speed operators.

  Tip: When using the Lines display type with the OpenGL display driver, slow-moving particles might not appear properly in the viewports. In such cases, to see all particles, add a second Display operator in the same event and set its Type to Dots.
- **Bounding Boxes—**(3D) Each particle appears as a bounding box. Use this option for a good representation of the final animation, at a slight cost in computational speed.
- **Geometry—**(3D) Each particle appears as its actual geometry. Use this option for the best representation of the final animation, at the greatest cost in computational speed.
- **Diamonds—**(2D) Each particle appears as a diamond.
- **Boxes—**(2D) Each particle appears as a small square.
- **Asterisks—**(2D) Each particle appears as an asterisk (*).
- **Triangles—**(2D) Each particle appears as a small triangle.

Visible %—Specifies the percentage of particles visible in the viewports. This option lets you speed up viewport redrawing by reducing the number of visible particles.

Show Particle IDs—When on, each particle's unique index number is visible in the viewports. Particles are numbered in the order of their birth, starting with 1 for the first particle born.

[color swatch]—Shows the color for particles displayed using options other than Geometry, as well as for the particle IDs. The software chooses a different color at random for each Display operator added to the system. In addition to the Type options, using different colors helps to distinguish among particles in different events.

To change the color, click the color swatch and use the Color Selector dialog to choose a new color.

If a system has a global Display operator, and you select its Source icon in a viewport, then you can change the color of the global Display operator from the color swatch on the Modify panel. However, changing the global Display operator's color in Particle View does not change the color of the swatch on the Modify panel.

Note: When a Particle Flow source icon is selected, all of its non-selected particles, other than those shown as geometry, are colored white in the viewports. To see all assigned particle colors, deselect the particle system.

Selected—Choose how *selected particles (page 2–138)* appear in the viewports. The choices are the same as for Type, above.

Force Operator

Particle View (page 2–125) > Click a Force operator in an event or add a Force operator to the particle system and then select it.

The Force operator lets you influence particle motion with one or more space warps from the Forces category. Use this operator along with different forces to simulate the effects of wind, gravity, and so on. The following force space warps work with the Force operator:

- *Displace (page 2–76)*
- *Drag (page 2–66)*
- *Gravity (page 2–73)*
- *Motor (page 2–61)*
- *PBomb (page 2–68)*
- *Push (page 2–59)*
- *Vortex (page 2–63)*
- *Wind (page 2–75)*

Note: By default, the influence of these space warps on Particle Flow particles is equivalent to their influence on the 3ds Max 2 particle systems: PArray, Super Spray, Blizzard, and PCloud. To obtain an influence on Particle Flow particles equivalent to that of the 3ds Max 1 particle systems Snow and Spray, set *Influence* to **100.0**.

Tip: To make the particles follow a path, use the *Speed By Icon operator (page 2–162)* and use Path Constraint to assign its icon to the path. For a procedure, see *Example: To send particles along a path: (page 2–163)*.

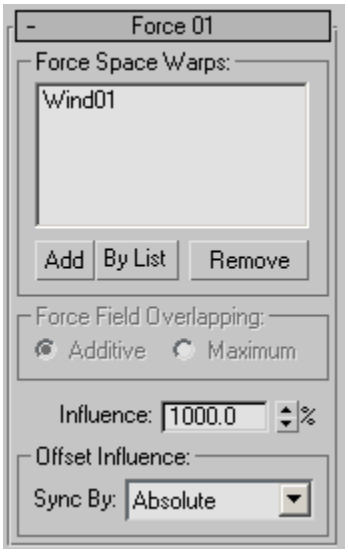
To employ deflectors for particle dynamics, with or without the Force operator, use the *Collision test (page 2–212)* and *Collision Spawn test (page 2–215)*.

Procedure

To affect particle motion with force space warps:

1. Add one or more force space warps to the scene, and set them up as necessary.
2. In Particle View, add a Force operator to any events in which particles are to be affected by the forces. To affect particles in all events, add the Force operator to the PF Source instead.
3. Highlight the Force operator, and then use the Add or By List button to apply the force space warps to the operator.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Force Space Warps group

This group displays the forces currently in effect, and let you add and remove forces.

[List]—Shows the forces that apply to this operator. If more than three forces apply, a scroll bar appears at the right side of the list.

If you delete a listed space warp from the scene, its name is replaced in the list by the entry "<deleted>".

Note: Particle Flow applies the forces to particle motion in the order in which the space warps appear in the list; the effect is cumulative in top-to-bottom order. First, the topmost space warp is applied to particle motion, then the second space warp is applied to the result of the first space warp, and so on. Changing the order can alter the final result.

Add—Click this button, and then select a force space warp in the scene to add it to the end of the list.

By List—Click this button, and then use the Select Force Space Warps dialog to add one or more space warps to the list. The space warps must already exist in the scene.

The software adds space warps to the list in same order in which they appear in the dialog. To effect a different order, use the Add button to add them one at a time.

Remove—Highlight a space warp in the list, and then click this button to remove it from the list. Any removed space warps remain in the scene.

Force Field Overlapping—Determines how multiple forces that occupy the same volume of space affect the particles. With Additive, the forces are combined according to their relative strengths. With Maximum, only the force with the greatest strength affects the particles.

For example, you might apply Wind and Gravity space warps to particles, and set their Strength parameters to 1.5 and 1.0, respectively. If you choose Additive, the Wind space warp will have approximately 50 percent more influence over the particles than the Gravity space warp. But if you choose Maximum, only the Wind space warp will affect the particles.

Influence—Specifies the strength with which the force or forces are applied to the particles as a percentage. Default=1000.0.

A negative Influence value reverses the force effects.

Note: By default, the influence of the Force space warps on Particle Flow particles is equivalent to their influence on the 3ds Max 2 particle systems PArray, Super Spray, Blizzard, and PCloud. To obtain an influence on Particle Flow particles equivalent to that of the 3ds Max 1 particle systems Snow and Spray, set Influence=**100.0**.

Offset Influence group

Choose the time frame for applying animated parameters. For an explanation, see *Animation Offset Keying group (page 2–158)*.

Sync By—Choose the time frame for applying animated parameters:

- **Absolute Time—**Any keys set for parameters are applied at the actual frames for which they're set.
- **Particle Age—**Any keys set for parameters are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for parameters are applied to each particle starting when it first enters the event.

Script Wiring rollout

This rollout appears in the parameters panel below the main operator rollout after you highlight the operator, right-click it, and then choose Use Script Wiring. Thereafter, a check mark appears next to the Use Script Wiring in the right-click menu, and the rollout appears whenever you highlight the operator. To turn off script wiring, choose Use Script Wiring again from the right-click menu.

Script wiring lets you use a script to control parameters that you normally specify in the operator's parameters. Place a *Script operator (page 2–208)* before the Force operator in the event, and then use it to define values in the particleFloat channel. You'll find an example script below.

Use Script Float As—Choose either of the following:

- **Not Used—**Particle Flow uses the *Influence setting* specified in the Parameters rollout.
- **Influence—**Particle Flow applies the script to the Influence setting.

particleFloat Sample Script

Note: See this topic in the online *User Reference* for the particleFloat sample MAXScript code.

Notes Operator

Particle View (page 2–125) > Click Notes in an event or add a Notes operator to the particle system and then select it.

The Notes operator lets you add a textual comment to any event. It doesn't have any direct effect on the particle system, but it helps you keep track of the overall function of each event.

Note: You can also add a comment directly to an event or action by right-clicking it and choosing Comments.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

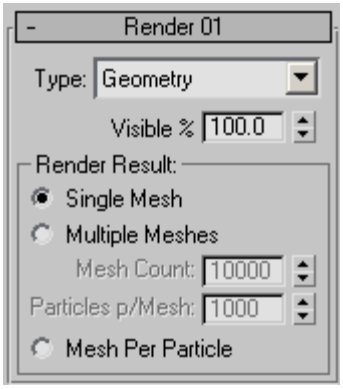
The interface is a simple text box. Click in the box and enter your text.

Render Operator

Particle View (page 2–125) > Click a Render operator in an event or add a Render operator to the particle system and then select it.

The Render operator provides controls related to rendering particles. You can specify the form that rendered particles are to take, and how to convert the particles to individual mesh objects for rendering purposes.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Type—Lets you render particles as bounding boxes or geometry, or disable the particle system at render time, or enable the particle system but prevent it from sending renderable particles to the 3ds Max renderer. Default=Geometry.

- **None—**The particle system (or event, if used locally) is considered non-renderable, and therefore is never evaluated during render.
- **BoundingBox—**Particles render as boxes whose sizes represent the particle geometry extents.

  Use this option with complex particle systems to speed up test renders.
- **Geometry—**Particles render with their actual geometry.

  Use this option with simpler particle systems, and for the final render.
- **Phantom—**Use this option when you want Particle Flow to function normally, calculating the animation at each frame, but not to render geometry directly. An example would be when you're using a *Final Step Update script (page 2–140)* to pass data from the particle system to other objects in the scene.

Visible %—The percentage of particles that render. Range=0 to 100. Default=100.

Lower this value for faster rendering of complex particle systems.

You can also reduce the number of particles in the system at render time with the *Quantity Multiplier > Render setting (page 2–137)*.

Render Result group

These settings determine how the system converts particles to mesh format for rendering. By default, the software renders all particles as a single mesh per event, which provides for the most efficient operation under most conditions. However, in some situations each particle should be converted to an individual mesh object by turning on Mesh Per Particle. For example, with particles each of whose face count exceeds 10,000, it's most efficient to render a mesh per particle. Alternatively, you can opt to combine particles into multiple meshes by specifying a face count and number of particles per mesh.

Certain renderers might require a non-default setting for Render Result. Consult the renderer documentation for further information.

Note: The number of particles Particle Flow can handle per frame is limited only by system resources, but each single mesh is allowed a maximum of 5,000,000 faces or vertices. If the total number of faces or vertices in a mesh exceeds 5,000,000, Particle Flow ignores particles beyond this limit.

Single Mesh—Sends the renderer one mesh object comprising all particles in the system.

Multiple Meshes—Sends the renderer the specified number of mesh objects, each containing the specified number of particles. If the result of dividing the total number of particles by the Particles p/Mesh value is less than the specified Mesh Count value, some of the meshes might contain few or no faces.

This is a compromise method of operation, and can be used with renderers that cannot handle all particles in a single mesh, but can deal with groups of particles of a certain size.

Mesh Count—The maximum number of mesh objects that Particle Flow will send to the renderer.

Particles p/Mesh—The number of particles that each mesh object will comprise.

Mesh Per Particle—Sends the renderer a separate mesh for each particle.

This is the least efficient method of operation, but might be required by certain renderers.

Note: With the default scanline and mental ray renderers, this method does not support rendering of particles born after the start of the rendered frame sequence. In general, use Mesh Per Particle only with renderers that require it.

Script Operator

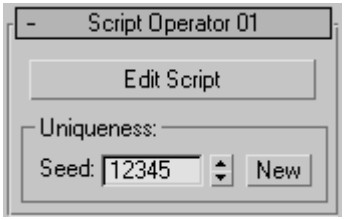
Particle View (page 2–125) > Click Script Operator in an event or add a Script Operator to the particle system and then select it.

The Script operator enables control of particles within the Particle Flow system using a MAXScript script. The script can use any program functionality available to MAXScript.

When you add a new Script operator, it contains a default script that slows particles, and then, when they are slow enough, splits off the first 50 particles into a stream traveling in a negative direction on the world X axis, and the remainder traveling in the opposite direction.

Tip: You can use MAXScript to align particle scale with the underlying bitmap, thus providing a "scale bitmap" function. This requires a sandwich with three operators: two Script operators, and a Speed By Surface operator in between. The first Script operator reads the current speed vector into the MXVector channel, thus caching the current speed. The Speed By Surface operator changes the speed according to the underlying bitmap. And the second Script operator reads the speed channel into a temporary variable, restores the speed from the MXVector channel (the cached value), and uses the temporary variable to define the scale. This way, the original speed is restored, and the scale value is defined by the bitmap.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Edit Script—Click this button to open the current script in a MAXScript Editor window.

For detailed information about the MAXScript utility, open the MAXScript Reference, available from Help menu > MAXScript Reference.

Uniqueness group

The Uniqueness setting provides a randomization seed that the script can use or ignore.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Flows

The flows category contains two operators for creating initial particle system setups. These are:

Empty Flow (page 2–209)

Standard Flow (page 2–209)

See also

Operators (page 2–142)

Tests (page 2–210)

Empty Flow

Particle View (page 2–125) > depot

Empty Flow provides a starting point for a particle system consisting of a single global event containing a Render operator. This lets you build a system completely from scratch, without having first to delete the default operators provided by the Standard Flow system.

To use Empty Flow, drag it from the depot to the event display. In Particle View, this creates a global event containing a single Render operator. If the Global default display option is active in the Particle View > Options menu, the global event will also contain a Display operator. Adding an Empty Flow also creates a Particle Flow Source icon in the viewports, at the world origin (0,0,0).

Note: If an orthogonal viewport is active when you add an Empty Flow to the system, the software orients the new source icon parallel to the plane of the active viewport, with the default emission direction pointing forward. For example, if the Front viewport is active, the icon is oriented parallel to the XZ plane in the world coordinate system, with the default emission direction along the positive Y axis. If a Camera or Perspective viewport is active, Particle Flow uses the default orientation: parallel to the XY plane, pointing in the negative Z direction.

See also

Standard Flow (page 2–209)

Standard Flow

Particle View (page 2–125) > depot

Standard Flow provides a starting point for a particle system consisting of a global event containing a Render operator, wired to a birth event containing a Birth, a Position, a Speed, a Rotation, a Shape, and a Display operator, with all parameters set to default values. This is the same system that the software creates automatically when you add a Particle Flow icon to the viewport.

To use Standard Flow, drag it from the depot to the event display. In Particle View, this creates the particle system described above. If the Global default display option is active in the Particle View > Options menu, the global event will also contain a Display operator. Adding a Standard Flow also creates a Particle Flow Source icon in the viewports, at the world origin (0,0,0).

Note: If an orthogonal viewport is active when you add a Standard Flow to the system, the software orients the new source icon parallel to the plane of the active viewport, with the default emission direction pointing forward. For example, if the Front viewport is active, the icon is oriented parallel to the XZ plane in the world coordinate system, with the default emission direction along the positive Y axis. If a Camera or Perspective viewport is active, Particle Flow uses the default orientation: parallel to the XY plane, pointing in the negative Z direction.

See also

Empty Flow (page 2–209)

Tests

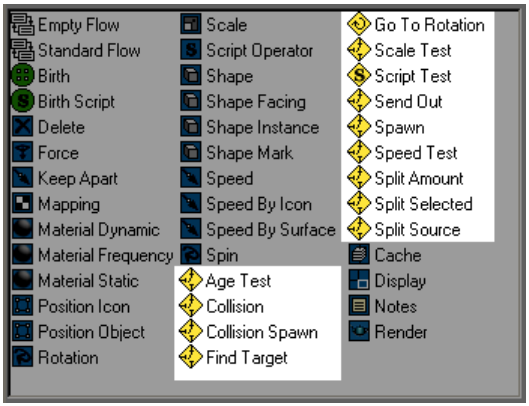
The basic function of a test in Particle Flow is to determine whether particles satisfy one or more conditions, and if so, make them available for sending to another event. When a particle passes a test, it is said to "test True." To send eligible particles to another event, you must *wire* the test to that event. Particles that don't pass the test ("test False") remain in the event and are repeatedly subjected to its operators and tests. Or, if the test isn't wired to another event, all particles remain in the event. You can use several tests in an event; the first test checks all particles in the event, and each test after the first checks only particles that remain in the event.

One test, Spawn, doesn't actually perform a test, but simply creates new particles from existing ones, and sets the new particles' test result to True so they're automatically eligible for redirection to another event. And the Send Out test simply sends all particles to the next event by default.

Some tests can also serve as operators, in that they contain parameters that modify particle behavior. If you don't wire a test to another event, it functions only as an operator; the test aspect doesn't affect particle flow.

Tip: Always place a test at the end of its event, unless you have specific reasons for placing it elsewhere. That way, all preceding actions can take effect during each integration step before the test is evaluated.

All the tests are grouped together in the Particle View depot, and are listed in alphabetical order. All test icons are the same: a yellow diamond containing a simplified diagram of an electrical switch.



The Particle Flow tests in the Particle View depot

The tests are:

Age Test (page 2–211)

Collision Test (page 2–212)

Collision Spawn Test (page 2–215)

Find Target Test (page 2–218)

Go To Rotation Test (page 2–224)

Scale Test (page 2–227)

Script Test (page 2–229)

Send Out Test (page 2–230)

Spawn Test (page 2–230)

Speed Test (page 2–233)

Split Amount Test (page 2–234)

Split Selected Test (page 2–235)

Split Source Test (page 2–236)

See also

Operators (page 2–142)

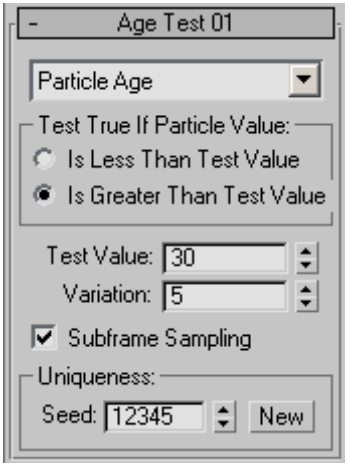
Flows (page 2–208)

Age Test

Particle View (page 2–125) > Click Age Test in an event or add Age Test to the particle system and then select it.

Age Test lets the particle system check whether a specific amount of time has passed since the start of the animation, or how long a particle has existed, or how long a particle has been in the current event, and branch accordingly.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

The first interface element is a drop-down list that lets you choose the type of age to test. Default=Particle Age.

- **Absolute Age—**Tests the current frame number in the animation.
- **Particle Age—**Tests the current age of each particle in frames. This is the default test type.
- **Event Age—**Tests the current duration of the current event in frames.

Test True if Particle Value—Lets you specify whether the test passes particles on to the next event if the age test succeeds or fails. Default=Is Greater Than Test Value

By default, Age Test returns True if the value tested for exceeds the Test Value quantity, but you can alternatively choose Is Less Than Test Value. For example, if you use the Absolute Age test type and set Test Value=60 and Variation=0, and choose Is Less Than Test Value, then particles will move to the next event only until frame 60. After frame 60, any remaining particles stay in the current event unless another test returns True.

Test Value—The specific frame number, particle age (in frames), or event duration (in frames) to test for. Default=30.

This value cannot be animated.

Variation—The number of frames by which the value tested for can vary randomly. Default=5.

This value cannot be animated.

To obtain the actual test value for each particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the Test Value setting. For example, if Test Value=300 and Variation=10, then tested value for each particle would be between 290 and 310.

Subframe Sampling—Turning this on helps avoid particle "puffing" when passing particles to the next event by testing the time at a much higher subframe resolution (that is, throughout each frame), rather than using the relatively coarse frame resolution. Default=on.

"Puffing" is the effect of creating separate "puffs" or clusters of particles, rather than a continuous stream.

Turn off Subframe Sampling to cause the test to be executed exactly at frame times.

Uniqueness group

The Uniqueness setting enables randomization of the test value variation.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Collision Test

Particle View (page 2–125) > Click Collision in an event or add Collision to the particle system and then select it.

Collision tests for particles that collide with one or more specified Deflector space warps. It can also test whether a particle has slowed down or sped up after one or more collisions, has collided more than once, and even whether it will collide with a deflector in a specified number of frames.

The Collision test supports all deflector space warps except the DynaFlect deflectors:

- *POmniFlect (page 2–78)*
- *SOmniFlect (page 2–84)*
- *UOmniFlect (page 2–85)*
- *SDeflector (page 2–87)*
- *UDeflector (page 2–89)*
- *Deflector (page 2–90)*

Tip: When testing for collisions with multiple deflectors, for best results, place all the deflectors in a single Collision test. This tests for collisions with all the deflectors simultaneously, and helps avoid possible missed collisions.

See also

Collision Spawn Test (page 2–215)

Procedure

Example: To test for particles slowing down after one or more collisions:

In the real world, particles bouncing repeatedly against a surface lose kinetic energy at each collision, and slow down gradually. Rather than testing for a specific number of bounces, you can use the Is Slow After Collision(s) to test whether particle speed has sunk below a specific level.

1. Start or reset 3ds Max. Set the animation length to **500** frames.
2. Add a Gravity space warp and a Deflector space warp. Decrease the deflector's Bounce setting below 1.0, and increase the Variation and Chaos values above 0.0. Set the deflector's Width and Length to **500**.
3. Create a default Particle Flow system. Position the emitter directly above the deflector.
4. Add a *Force operator (page 2–204)* to the end of Event 01 and add the Gravity space warp to the Force operator.
5. Create a new event with a *Display operator (page 2–202)*, and change the display type to a different choice than is used in Event 01.
6. Add a Collision test to Event 01, below the Force operator, and wire it to the new event.
7. In the Collision test settings, add the deflector. Choose Is Slow After Collision(s), and set Speed Min=**100** (assuming you're using the default initial speed of 300).
8. Play the animation. You might need to adjust one or more settings before seeing the expected behavior.

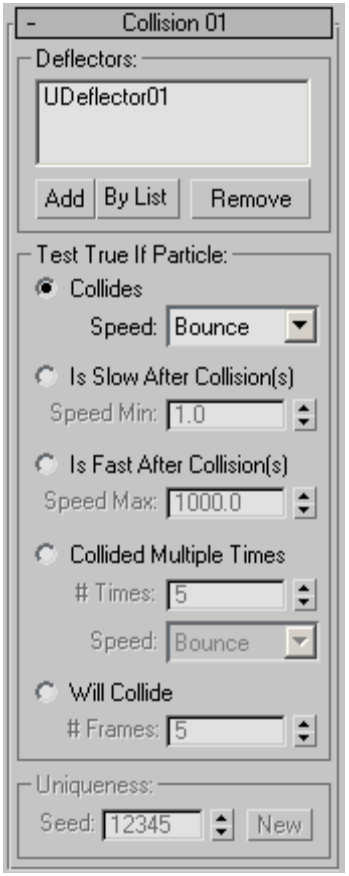
   After several bounces, the particles change in appearance and move steadily away from the deflector, indicating that they've entered the second event. Of course, you can set any behavior you like in this event.

In the next step, you'll see what happens when actions in an event are not in the right order.

9. In Event 01, move the Collision test above the Force operator, and then play the animation.

   Quite a few particles leak through the deflector. This happens because the software first tests the particles for a collision, and then applies the Gravity force. The particles that are approaching the deflector and are very close to it are being tested for a collision, which tests False because they haven't actually struck the deflector yet. The software then applies the gravity, which pushes them past the deflector, making them no longer eligible for testing for collision. Generally speaking, it is best to keep Force operators above Collision tests in each event to ensure that particles don't leak through the deflector.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Deflectors group

This group displays the deflectors currently in effect, and let you add and remove deflectors.

[list]—Shows the deflectors that apply to this operator. If more than three deflectors apply, a scroll bar appears at the right side of the list.

If you delete a listed space warp from the scene, its name is replaced in the list by the entry "<deleted>".

Add—Click this button, and then select a Deflector space warp in the scene to add it to the list.

By List—Click this button, and then use the Select Deflectors dialog to add one or more space warps to the list. The space warps must already exist in the scene.

Remove—Highlight a deflector in the list, and then click this button to remove it from the list. Any removed space warps remain in the scene.

Test True If Particle group

Choose the condition under which the test will pass particles on to the next event, and then adjust the associated setting or settings. Default=Collides.

Collides—Choose this option, and then set Speed according to how particle speed should be affected by the collision.

Speed—Choose one of the following. Default=Bounce.

- **Bounce—**The speed and direction after collision is determined by the deflector properties.
- **Continue—**Particle speed and direction are unaffected by the collision.
- **Stop—**Particle speed is set to 0 after the collision.
- **Random—**Particles bounce off the deflector in random directions.

Is Slow After Collision(s)—The test succeeds if, after collision, particle speed is less than the Speed Min value.

With this choice, particle behavior with respect to the deflector(s) is the same as with the Collides > Bounce option.

Speed Min—Particles traveling at less than this speed, in system units per second, test True and become eligible for moving to the next event. Default=1.0.

Is Fast After Collision(s)—The test succeeds if, after collision, particle speed is greater than the Speed Max value.

With this choice, particle behavior with respect to the deflector(s) is the same as with the Collides > Bounce option.

Speed Max—Particles traveling faster than this speed, in system units per second, test True and become eligible for moving to the next event. Default=1000.0.

Collided Multiple Times—The test becomes True after a particle collides a specific number of times. The particle is moved to the point of the last collision and then redirected to the next event.

 # Times—The number of times a particle must collide in order to test True.

Speed—Determines speed and direction after the specified number of collisions. See above for explanations of the choices.

Will Collide—The software extrapolates particle motion in a linear fashion, based on the current direction and speed, and the test becomes True if the results suggest that the particle will collide with the deflector during a specified time interval. The particle is redirected to the next event without updating its speed or position.

 # Frames—The number of frames ahead during which the software looks for an impending collision.

Uniqueness group

The Uniqueness setting enables randomization of the Random options in the Speed drop-down lists.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Collision Spawn Test

Particle View (page 2–125) > Click Collision Spawn in an event or add Collision Spawn to the particle system and then select it.

Collision Spawn creates new particles from existing ones that collide with one or more Deflector space warps. You can specify different post-collision behavior for the colliding particles and their offspring. Each spawned particle is born at the same location as its parent, and has the same orientation and shape. Collision Spawn can give the spawned particles a different speed and scaling factor. If you wire the Collision Spawn test to another event, spawned particles are sent to that event, where you can specify different properties for the new particles.

Examples of Collision Spawn usage include marks or explosions resulting from collisions between particles and objects. To achieve these effects, you can use Collision Spawn in conjunction with the *Shape Mark (page 2–183)* and *Shape Facing (page 2–176)* operators.

The Collision Spawn test supports all deflector space warps except the DynaFlect deflectors:

- *POmniFlect (page 2–78)*
- *SOmniFlect (page 2–84)*
- *UOmniFlect (page 2–85)*
- *SDeflector (page 2–87)*
- *UDeflector (page 2–89)*
- *Deflector (page 2–90)*

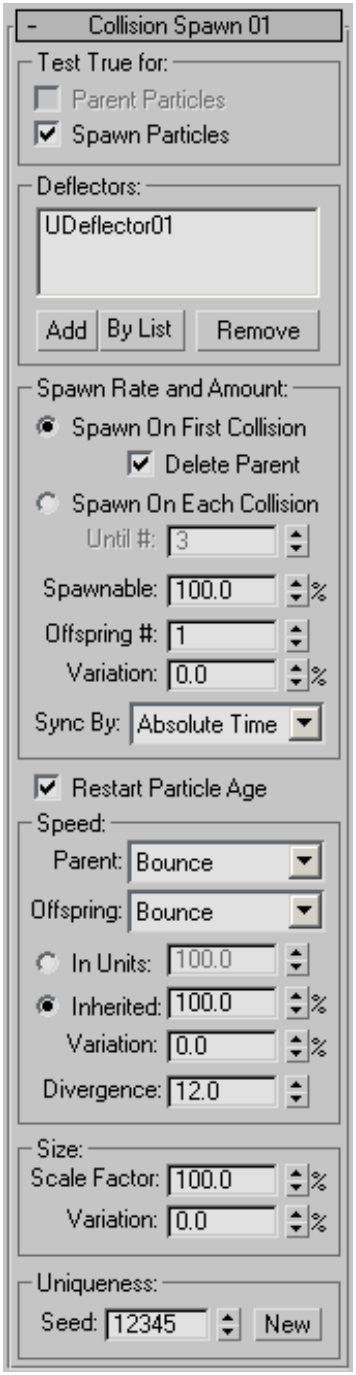
Tip: When testing for collisions with multiple deflectors, for best results, place all the deflectors in a single Collision Spawn test. This tests for collisions with all the deflectors simultaneously, and helps avoid possible missed collisions.

See also

Collision Test (page 2–212)

Spawn Test (page 2–230)

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

In the context of Collision Spawn, a *parent* is the original particle from which new particles are spawned.

Test True for group

These check boxes let you specify which particles, if any, should become eligible for redirection to the next event upon satisfaction of the test conditions.

Note: After a particle tests True, the Collision Spawn test no longer tests the particle for collision. If you simply want particles to spawn every time they collide, turn either or both of these off and use a different test for redirection.

Parent Particles—When on, parent particles qualify for redirection when the test conditions are met. Default=on.

When using the *Spawn On First Collision option*, Parent Particles is available only when Delete Parent is off.

This option is also available when using the *Spawn On Each Collision option*. In this case, parent particles test True only after colliding the number of times specified by the Until # value.

Spawn Particles—When on, spawned particles qualify for redirection when the test conditions are met. Default=on.

Deflectors group

This group displays the deflectors currently in effect, and let you add and remove deflectors.

[list]—Shows the deflectors that apply to this operator. If more than three deflectors apply, a scroll bar appears at the right side of the list.

If you delete a listed space warp from the scene, its name is replaced in the list by the entry "<deleted>".

Add—Click this button, and then select a Deflector space warp in the scene to add it to the list.

By List—Click this button, and then use the Select Deflectors dialog to add one or more space warps to the list. The space warps must already exist in the scene.

Remove—Highlight a deflector in the list, and then click this button to remove it from the list. Any removed space warps remain in the scene.

Spawn Rate And Amount group

Use these settings to specify when particles are to spawn and other values related to how many particles are spawned.

Spawn On First Collision—Particles spawn only the first time they collide with a deflector.

Delete Parent—When on, deletes each original particle from which a new one is spawned. Available only with the Spawn On First Collision option.

Spawn On Each Collision—Particles spawn on each of multiple collisions, up to a limit specified with the Until # parameter.

Until #—The maximum number of collisions by the parent particles that produce spawned particles. Available only with the Spawn On Each Collision option. Default=3.

When Test True For > Parent Particles is on, parent particles test True only after colliding the number of times specified by the Until # value.

Spawnable—The percentage of particles in the current event that will spawn new particles. This is determined once for each particle, when it enters the event. However, the parameter is animatable. Default=100.0.

For values other than 100.0, Spawnable uses a randomized selection process, which is affected by the Uniqueness Seed value. For example, with five parent particles, Offspring #=1, and Spawnable=80.0, you might get any number of spawned particles between two and five for each collision. The average per spawning would be four, however.

Offspring #—The number of new particles the system creates from each parent particle for each spawning event. Default=1.

Variation—The amount by which the Offspring # value can vary randomly. Default=0.0.

To obtain the actual test value for each particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then applies the result as a percentage of the Offspring # setting. For example, if Offspring #=20 and Variation=10, then the actual number of offspring for each particle would be between 18 and 22.

Sync By—Choose the time frame to use when animating Offspring # and Variation:

- **Absolute Time—**Any keys set for parameters are applied at the actual frames for which they're set.
- **Particle Age—**Any keys set for parameters are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for parameters are applied to each particle starting when it first enters the event.
- **Restart Particle Age—**When on, sets the age of each newly spawned particle at 0. When off, each spawned particle inherits its parent's age. Default=on.

Speed group

Lets you specify the behavior of spawned particles, as well as their speed in absolute terms or relative to the parents' speed, with optional random variation. The direction of a spawned particle is always in relation to that of its parent, but you can

set a Divergence so they eventually spread out. Default=Inherited.

Parent—Specify the parent's behavior. Available only when Delete Parent is off.

- **Bounce—**The speed and direction after collision is determined by the deflector properties.
- **Continue—**Particle speed and direction are unaffected by the collision.

Offspring—Specify the behavior of the newly spawned particles.

- **Bounce—**The speed and direction of offspring after collision is determined by the deflector properties.
- **Continue—**Speed and direction of offspring are unaffected by the collision.

In Units—Choose this to specify the speed of spawned particles in system units per second. Default=100.0.

A positive value inherits the parent's direction; a negative value reverses it.

Inherited—Choose this to specify each spawned particle's speed as a percentage of its parent's speed. Default=100.0

A positive value inherits the parent's direction; a negative value reverses it.

Variation—The amount by which a spawned particle's speed can vary randomly. Default=0.0.

To obtain the actual speed for each spawned particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the particle's speed as specified or inherited. For example, if a particle's speed is 100 units/second and Variation=20, then the tested value for each particle would be between 80 and 120 units/second.

Divergence—When on, spreads out the stream of spawned particles. Use the numeric setting to define the extent of the divergence in degrees. Range=0 to 180. Default=0.

Size group

Scale Factor—The amount of uniform scaling to apply to each spawned particle, as a percentage of its parent's size. Default=100.0.

Variation—The amount by which a spawned particle's scale can vary randomly. Default=0.0.

To obtain the actual scaling for each spawned particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the Scale Factor value. For example, if Scale Factor=100 and Variation=20, then each spawned particle would be between 80 and 120 percent of its parent's size.

Uniqueness group

The Uniqueness setting enables randomization of the Spawnable result, when less than 100.0, as well as of the Variation values.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Find Target Test

Particle View (page 2–125) > Click Find Target in an event or add Find Target to the particle system and then select it.

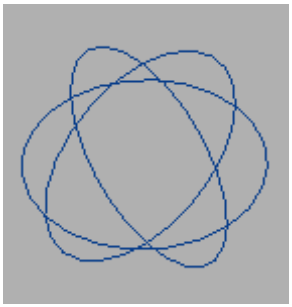
Select a Find Target icon in the viewport. > Modify panel > Parameters rollout

Create panel > Helpers > Particle Flow > Find Target

By default, Find Target sends particles to a specified target or targets. Upon reaching a target, the particles become eligible for redirection to another event. You can specify that the particles should use a particular speed or time frame in

moving toward the target. You can also specify where on the target the particles should go.

Alternatively, you can use Find Target as a simple proximity test: If a particle comes within a certain distance of its target, it becomes eligible for redirection to the next event.



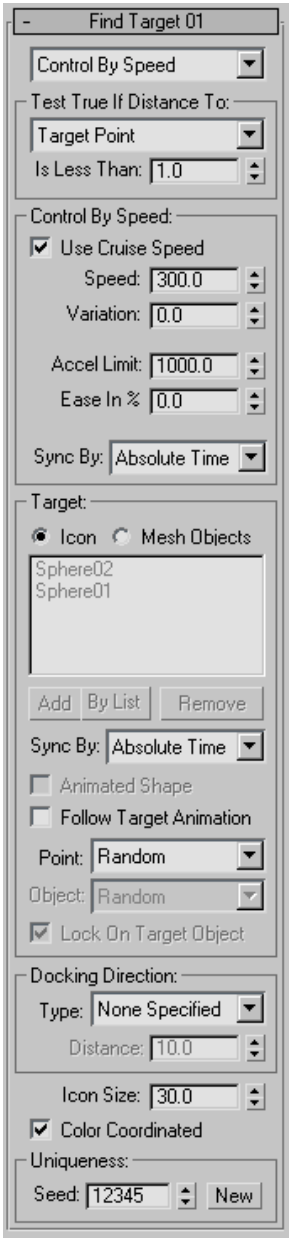
Find Target icon

When you add a Find Target test to the particle system in *Particle View (page 2–125)*, a spherical Find Target icon appears in the scene at the world origin (0,0,0). You can use this icon as a target, or you can use one or more mesh objects in the scene as targets. To display the Find Target parameters on the Modify panel, select the icon. If you delete the icon, the software also deletes the test.

Note: If you add Find Target from the Create panel, Particle Flow creates a separate event for the test in the particle diagram.

Tip: In a particle loop, all particles end up at their respective starting positions, enabling seamless repetition of the resulting animation. You can make particle loops with Particle Flow using a Script operator and a Find Target operator. At the start of the loop, the Script operator should read all particle positions and write them into the MXS Vector channel. Then, at the end of the loop, set the Find Target operator to Control by Time, set Timing to Absolute Time, set Time to the end of the loop, and in the Target group, set Point to By Script Vector. Particle Flow will direct particles to the previously cached position at the specified frame.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Control By ...—The drop-down list at the top of the parameters panel lets you choose whether to send particles to a target by specifying the speed and acceleration, or by specifying the amount of time they should take. Alternatively, by choosing No Control, you can test particles' distance from a target.

- **Control By Speed—**Specify speed and acceleration settings for the particles to follow while traveling to the target. Use the Control By Speed group to set the parameters.
- **Control By Time—**Specify time-based settings for the particles to follow while traveling to the target. Use the Control By Time group to set the parameters.
- **No Control—**With this option, Find Target functions only as a proximity test. When a particle comes within the specified distance of the target, the particles become eligible for redirection to the next event.

  Note: When you choose No Control, the Find Target test does not affect particle speed or direction.

Test True If Distance To group

These settings let you choose which distance the Find Target test measures and specify the measured distance. You can choose Target Pivot or Target Point, and specify a distance with the Is Less Than setting.

Target Pivot—Measures the distance between the particle and the target pivot.

If the particles are directed to the outside of the target and the Is Less Than value is small, this condition might never be satisfied.

Target Point—Measures the distance between the particle and the *target point*.

Is Less Than—When particles are closer than this distance from target pivot or point, they test True and become eligible for redirection to the next event. Measured in scene units.

Note: If you set Is Less Than to 0, the particles might never test True. This might be desirable for an animation of, for example, bees buzzing around a flower but not landing on it. In this case, you might want to use a low value for Accel Limit so the bees don't circle too close to the flower.

Control By Speed group

Use these settings to specify speed and acceleration parameters when using Control By Speed or Speed Then Time. This group is available only when using Control By Speed.

Use Cruise Speed—When on, the software gives you explicit control over particle speed and speed variation. When off, the software calculates particle speed automatically using the Accel Limit value. Default=on.

Speed—Particle speed in scene units per second. Default=300.0.

Variation—The amount by which the actual Speed value can vary randomly. Default=0.0.

To obtain the actual speed for each particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then applies the result as a percentage of the Speed setting. For example, if Speed=200 and Variation=10, then the actual speed for each particle would vary randomly between 190 and 210.

Accel Limit—Sets the acceleration limit. This value impacts the inertia and speed of the particles. Default=1000.0.

The default acceleration limit value is based on the default Speed value of 300.0. If you change the Speed value, it is recommended that you also change the Accel Limit value proportionately.

Tip: Use a lower Accel Limit value for smooth motion, and a higher value when greater accuracy

is needed, such as when the particles should hit a small target. You can animate this setting (use Sync By > Event Duration) to specify different appropriate values, depending on the required results.

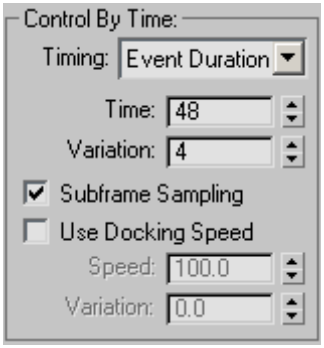
Ease In %—Controls the rate by which particles slow down when it approaches the target point.

The software calculates the final speed with this formula: (100% - Ease In) * Speed. Therefore, if the Ease In value is 100%, a particle should approach the target with a speed of 0, and if Ease In is 0%, the particle doesn't slow down at all when approaching the target. At intermediate values, the speed is calculated according to the distance to the target point, as linear interpolation between the initial (cruise) speed, and the final speed. When a particle enters the event, the distance to the target point is calculated and later used for the interpolation. Default=0.0.

Sync By—Choose the time frame to use when animating Speed, Variation, and Accel Limit:

- **Absolute Time—**Any keys set for parameters are applied at the actual frames for which they're set.
- **Particle Age—**Any keys set for parameters are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for parameters are applied to each particle starting when it first enters the event.

Control By Time group



Lets you specify the amount of time particles should take to reach the target.

This group is unavailable when using Control By Speed.

Timing—Determines how the software applies the specified timing, defined by the Time and Variation values. The possible options are:

- **Absolute Time—**Time refers to the overall time of the system. Each particle will reach its target at the frame number specified by Time.
- **Particle Age—**Time refers to the time elapsed since the birth of the particle. Each particle will reach its target when its age reaches the value specified by Time.
- **Event Duration—**Time refers to the time elapsed since the particle entered the current event. Each particle will reach its target when it has been in the current for the number of frames specified by Time.

Time—The number of frames particles should take to reach the target. Default=60.

Variation—The number of frames by which Time can vary randomly. Default=5.

To obtain the actual time to the target for each spawned particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the Time

value. For example, if Time=60 and Variation=20, then the time to target for each particle would be between 40 and 80 frames.

Subframe Sampling—Turning this on helps avoid particle "puffing" by timing particles at a much higher subframe resolution (that is, throughout each frame), rather than using the relatively coarse frame resolution. Default=on.

"Puffing" is the effect of clustering particles, rather than producing a continuous stream. This effect is especially noticeable when the emitter is animated.

Use Docking Speed—Lets you specify particles' speed when they reach the target.

A particle might be required to reach the target from a specific direction at a specific speed. When Use Docking Speed is off, the software calculates a path for particles to reach the target in the shortest distance with the least acceleration along the path. When Use Docking Speed is on, the software calculates particles' final speed when they reach the target using the Speed and Variation parameters. Therefore if you want a "smooth landing," set Speed to 0.0. Default=off.

Speed—The speed of particles when they reach the target in scene units per second. Default=100.0.

Variation—The amount by which the actual Speed value can vary randomly. Default=0.0.

Target group

By default, Find Target uses its own target icon, but you can use these controls to designate other scene objects as targets instead.

Icon—Use the Find Target icon as the target.

Each Find Target test has its own target icon. Even if you don't use it as a target, it still influences particle behavior if you set the docking type to Parallel, Spherical, or Cylindrical.

Mesh Objects—Use one or more scene mesh objects as targets. Choose this to make the list and list-management buttons available.

If you designate more than one target, the target used by each particle is determined by the choice in the *Object drop-down list* in this group.

[list]—Shows the mesh objects used as targets. If more than three target objects apply, a scroll bar appears at the right side of the list.

Add—Click this button, and then select a mesh object in the scene to add it to the list as a target.

By List—Click this button, and then use the Select Target Objects dialog to add one or more mesh objects to the list. The objects must already exist in the scene, and the dialog shows only eligible objects.

Remove—To remove a target object from the list, first highlight it and then click this button. Any removed objects remain in the scene.

Sync By—Choose the time frame to use when sending particles toward animated objects when using Animated Shape or Follow Target Animation:

- **Absolute Time—**Animation derived from the target is applied to particles immediately.
- **Particle Age—**Animation derived from the target is applied to particles at the corresponding frames of each particle's existence.
- **Event Duration—**Animation derived from the target is applied to particles is applied to each particle starting when it first enters the event.

Animated Shape—Turn on to allow particles to target the surface of an object whose form is animated by scaling, by morphing, or with modifiers. This requires more computation, because the destination must be updated at every integration step.

Follow Target Animation—Turn on to allow particles to follow a moving target; that is, a target whose location is animated. This requires more computation, because the destination must be updated at every integration step.

Point—Lets you specify where on its target a particle should land.

- **Random—**Each particle targets a random point on the target.
- **Closest Surface—**Each particle targets the nearest point of the target's surface.
- **By Script Vector—**The target point or points are defined using a Script operator that defines values in the particleVector channel. The Script operator can be anywhere upstream of the Find Target test.

  Note: When using By Script Vector to target absolute positions, such as vertex locations, be sure to set Target to Icon. If you set it to Mesh, the positions specified by the script will be relative to the position of the mesh object. This latter option is useful for placing scripted targets on the surface of a moving object.

Object—With multiple mesh targets, lets you specify how the software should choose among them. Available only when targeting more than one object.

- **Random—**For each particle, the software chooses a target object at random.
- **Closest Pivot—**For each particle, the software chooses as its target the object whose pivot is nearest to the particle.
- **Closest Surface—**For each particle, the software chooses as its target the object whose surface is nearest to the particle.
- **Least Deviation—**For each particle, the software chooses as its target the object that requires the least change in its current direction (or resteering) to reach.
- **By Script Integer—**For each particle, the choice of a targeted object is defined by a script operator that sets an index. This index corresponds to a target-object entry position in the target list. See *Script Operator Example (page 2–224)*.

Lock On Target Object—When on, the software calculates each particle's target object once: when the particle enters the event. Thereafter, the particle is "locked on" to its target object. When off, the software can continually recalculate the target object for each particle. Available only when multiple target objects are designated.

For example, if you set particles to target the closest surface, due to the target animation and particle movement, the definition of the closest surface is constantly changing. Thus the particle may change the target object due to the circumstances.

Note: Each time Find Target sets a target object, it "locks on" to a specific point on that object. This point can change only if the target object changes. Thus, with a single target object, the target point always remains constant *relative to the target object*. That is, if the target object or its surface is animated, and Follow Target Animation or Animated Shape is on, the absolute coordinates of the target point may change.

Note: When Lock On Target Object is off, more calculation is required because the system might have to recalculate each particle's optimal target point in each frame.

Docking Direction group

Docking type—Lets you specify from which direction particles should approach targets.

- **None Specified—**No docking constraints. Particles reach their targets in the most efficient way, based on their assigned parameters and their current attributes.

- **Along Icon Arrow—**The final direction is the same as the Find Target icon arrow.

  Note: When using this option, arrows appear on the Find Target icon to indicate the direction particles will use for docking. You can change the docking direction by reorienting the icon. This applies even when using mesh objects as targets.

- **Icon Spherical—**The final direction points toward the center of the operator icon.

  Note: When using this option, arrows appear on the Find Target icon to indicate the directions particles will use for docking. You can change the docking directions by reorienting the icon. This applies even when using mesh objects as targets.

- **Icon Cylindrical—**The final direction points toward the icon arrow as a line, thus forming a cylindrical field with the arrow as the cylinder's main axis.

  Note: When using this option, arrows appear on the Find Target icon to indicate the directions particles will use for docking. You can change the docking directions by reorienting the icon. This applies even when using mesh objects as targets.

- **Surface Normals—**Each particle reaches its target point from a direction perpendicular to the surface at that point.

  Note: When using this option, arrows appear on the Find Target icon to indicate that particles will use surface normals for docking; the actual directions they will use depends on the target surface. This applies even when using mesh objects as targets.

Distance—The distance from the target at which particles begin docking behavior. This includes the docking direction, and, when using Control By Time, the docking speed.

Icon Size—Set the size of the Find Target icon. This affects particle behavior when using the icon as the target.

Color Coordinated—When on, the Find Target icon uses the color of the event containing the test as defined by its local Display operator, if one exists. This applies even if the Display operator is turned off. When off, the Find Target icon uses the default Test Gizmos color as defined in Customize User Interface > Colors > Particle Flow. Default=on.

Turn on Color Coordinated to make it easier to spot the Find Target icon, because the particles in the event use the same color as the icon.

Uniqueness group

The Uniqueness setting enables randomization of the Random options and the Variation options.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Script Operator Example

Note: See this topic in the online *User Reference* to view the script operator sample MAXScript code.

Go To Rotation Test

Particle View (page 2–125) > Click Go To Rotation in an event or add Go To Rotation to the particle system and then select it.

The Go To Rotation test enables a smooth transition in the rotational component of a particle, so that the particle can gradually rotate to a specific orientation over a specific period. An example of its usage would be with falling leaves, which spin chaotically as they fall, but land on a flat side rather than an edge. The test aspect lets you direct the particle to a new event when the transition period ends.

To set a target orientation, place the Go To Rotation test *before* an orientation-type operator (*Rotation (page 2–153)* or a *Script operator (page 2–208)* if it defines the rotation channel) in the same event. In this situation, the Go To Rotation test can grab the particle rotational component before the orientation-type operator overwrites it. The Go To Rotation operator modifies the particle orientation and spinning in the post-evaluation cycle. For an example, see the procedure below.

Warning: The Go To Rotation test is not compatible with the Spin, Shape Facing, and Shape Mark operators. Do not use any of these operators in the same event with a Go To Rotation test.

Note: You can define the transition period only in terms of time. You cannot set the test to come to a specific rotation by the time of another test, such as a collision test. Also, limited control is provided over the axis of spinning when a particle comes to the final rotation.

Procedure

Example: To make falling particles land smoothly, face up:

This procedure assumes a basic knowledge of Particle Flow usage.

1. In the Perspective viewport, add a Particle Flow system and raise its icon about 80 units on the Z axis.
2. In Event 01, make the following changes:
   - Birth > Amount=**50**
   - Speed=**100**
   - Shape=Cube (or use Shape Instance with a custom object, such as a coin-shaped cylinder)
   - Display > Type=Geometry

   This reduces the number of particles and slows them down, making it easier to see what's going on.
3. Add a *Spin operator (page 2–154)* to Event 01, and set both Spin Rate and Variation to **150**.

   This gives the particles random spinning behavior as they fall.
4. Add an Age test to the end of Event 01, and set Test Value=15.

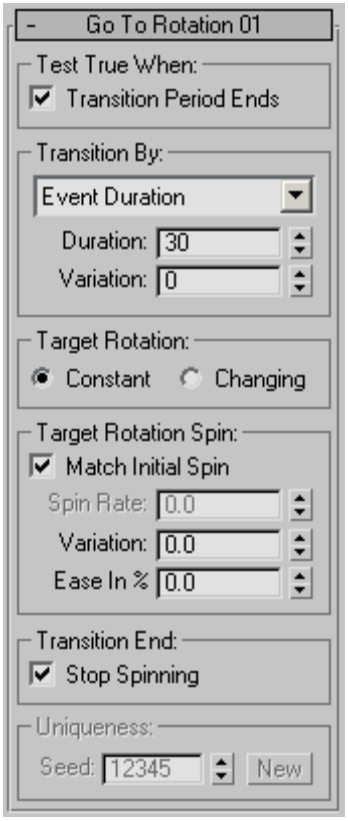
   This allows each particle to fall and tumble for 15 frames before Go To Rotation takes effect.
5. From the Depot, drag a Go To Rotation test to an empty area in Event Display. Set Duration=**15** and keep all other default settings.
6. Insert a Rotation operator immediately *after* the Go To Rotation test. Choose Random Horizontal as the orientation matrix. Keep the other default settings.

   The Go To Rotation test will use this as the final orientation for the particles.
7. Use a Speed event to create a third event. Set Speed=**0.0**.

   This stops the particles at the end of the animation.
8. In both new events, set Display > Type=Geometry.
9. Wire the Age Test in Event 01 to Event 02, and then wire the Go To Rotation test in Event 02 to Event 03.
10. Play the animation.

    The particles tumble chaotically as they fall for about 30 frames, and then come to a smooth stop, facing up.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Test True When group

Transition Period Ends—When on, the particles test True at the end of the specified transition period, and become eligible for redirection to the next event. When off, particles will not go to the next event, even if wired. Turn off to disable the test aspect of Go To Rotation. Default=on.

Transition By group

Determines how the software applies the specified timing, defined by the Time and Variation values. Default=Event Duration.

The possible options are:

- **Absolute Time—**Time refers to the overall time of the system. Each particle will reach its target orientation at the frame number specified by Time.
- **Particle Age—**Time refers to the time elapsed since the birth of the particle. Each particle will reach its target orientation when its age reaches the value specified by Time.
- **Event Duration—**Time refers to the time elapsed since the particle entered the current event. Each particle will reach its target orientation when it has been in the current for the number of frames specified by Time.

Duration—The number of frames particles should take to reach the target orientation. Default=30.

Variation—The number of frames by which Duration can vary randomly. Default=0.

To obtain the actual time to the target orientation for each particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the Duration value. For example, if Duration=60 and Variation=20, then the time to target orientation for each particle would be between 40 and 80 frames.

Target Rotation group

When you use an orientation-type operator with Go To Rotation, this setting lets you determine whether the test sets the target orientation on a one-time or ongoing basis. Default=Constant.

Constant—Defines that the orientation-type operator sets a constant orientation for a particle.

For example, when using a Rotation operator set to an orientation matrix other than Speed Space Follow, the Go To Rotation operator would acquire the target orientation only once, and then use it as its goal.

Changing—Defines that the orientation-type operator sets a changing rotation for a particle. At each frame the desirable final rotation may be different.

For example, if you use the test with a Rotation operator set to Speed Space Follow, the test will adjust the particle rotation constantly to aim at the changing final rotation.

Target Rotation Spin group

Defines the angular velocity for each particle when it reaches the target orientation.

Match Initial Spin—Sets the angular velocity at the end of the transition period to the same as the angular velocity when particle enters the event.

Note: The axis of rotation might still be different, because it is calculated on the fly to let the particle come to the target orientation.

Spin Rate—Defines each particle's angular velocity, in degrees per second, when it reaches the target orientation. Available only when Match Initial Spin is off. Default=0.0.

For a smooth transition to the target orientation, set to **0.0**.

Variation—Defines a random variation in the Spin Rate value when a particle reaches the target orientation. Default=0.0.

To obtain the final angular velocity for each particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the Spin Rate value. For example, if Spin Rate=6.0 and Variation=1.0, then the final angular velocity for each particle would be between 5.0 and 7.0 degrees per second.

Ease In %—Defines a curve in achieving the final spin rate. Default=0.0

When set to 0.0, the test produces a linear interpolation between the initial and final spin rate, and when set to 100.0, the final spin rate is achieved earlier. For the smoothest approach to the target orientation, set Spin Rate and Variation to 0.0, and Ease In to 100.0.

Transition End group

Stop Spinning—When on, halts particle rotation when it reaches the target orientation. Default=on.

Even with Spin Rate spinner set to 0.0, we recommend that you keep this check box on to avoid slow spinning at the end, due to computational averaging and marginal errors.

Uniqueness group

The Uniqueness setting enables randomization of the Variation options.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

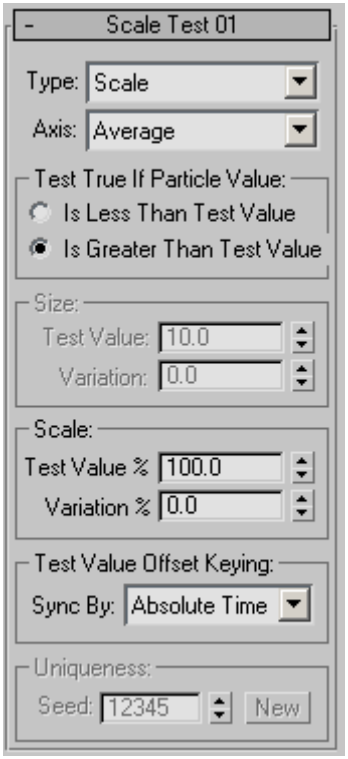
Scale Test

Particle View (page 2–125) > Click Scale Test in an event or add Scale Test to the particle system and then select it.

Scale Test lets the particle system check particle scaling, or particle size before or after scaling, and branch accordingly. The test provides a variety of axis options for measuring scale or size.

You can use this test to cause a change in behavior based on size. For example, a bubble could grow to a certain size, and then pop. Or an object could shrink in size until it falls under the influence of a force, like wind.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Type—Choose the type of measurement to test. You can test actual scaling, or the size before or after scaling. Default=Scale.

For example, if the X-axis size of a particle's bounding box starts out at 10 system units, and you then use the Scale operator to scale it by 150% on the X axis, the pre-scale size is 10, and the post-scale size is 15. And, of course, the scale is 150.

Note: The software measures particle size based on each particle's bounding-box dimensions in the particle's local coordinate space. With the low-polygon Shape options Tetra and Sphere, the result of this measurement might not be the same as the Shape operator's Size value.

- **PreScale Size—**Tests the size before scaling.
- **PostScale Size—**Tests the size after scaling.
- **Scale—**Tests the scaling percentage.

Axis—Choose the axis to measure. Default=Average.

- **Average—**Obtains an average measurement by adding the sizes on all three axes and then dividing by three.
- **Minimum—**Uses the smallest dimension.
- **Median—**Uses the middle dimension in order of size. For example, if the particle dimensions are X=5, Y=6, Z=12, then the number compared to Test Value would be 6.
- **Maximum—**Uses the largest dimension.
- **X/Y/Z—**Uses the specified dimension.

Test True if Particle Value—Lets you specify whether the test passes particles on to the next event if the speed test succeeds or fails. Available for all tests except True When Accelerates/Decelerates. Default=Is Greater Than Test Value.

By default, Scale Test returns True if the value tested for exceeds the Test Value quantity, but you can alternatively choose Is Less Than Test Value. For example, if you set Type to Scale and set Test Value=150 and Variation=0, and choose Is Less Than Test Value, a particle will move to the next event only when its scaling factor is less than 150%.

Size group

These settings are available when Type is set to PreScale Size or PostScale Size.

Test Value—The specific size to test for. Default=10.0.

Variation—The amount by which the value tested for can vary randomly. Default=0.0.

To obtain the actual test value for each particle, the system multiplies the Variation value by a random

number between -1.0 and 1.0, and then adds the result to the Test Value setting. For example, if Test Value=10 and Variation=5, then the tested value for each particle would be between 5 and 15.

Scale group

These settings are available when Type is set to Scale.

Test Value—The specific scaling factor to test for. Default=100%.

Variation—The amount by which the value tested for can vary randomly. Default=0.0%.

To obtain the actual test value for each particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the Test Value setting. For example, if Test Value=100% and Variation=10%, then the tested value for each particle would be between 90% and 110%.

Test Value Offset Keying group

Sync By—Choose the time frame to use when animating Test Value and Variation:

- **Absolute Time—**Any keys set for parameters are applied at the actual frames for which they're set.
- **Particle Age—**Any keys set for parameters are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for parameters are applied to each particle starting when it first enters the event.

Uniqueness group

The Uniqueness setting enables randomization of the test value variation. Available only when either Variation value exceeds 0.0.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

Script Test

Particle View (page 2–125) > Click Script Test in an event or add a Script Test operator to the particle system and then select it.

Script Test lets you test particle conditions using a MAXScript script. The script can use any program functionality available to MAXScript.

The default test script tests for the presence of all particles within a spherical volume of radius 20.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Edit Script—Click this button to open the current script in a MAXScript Editor window.

For detailed information about the MAXScript utility, open the MAXScript Reference, available from Help menu > MAXScript Reference.

Uniqueness group

The Uniqueness setting provides a randomization seed that the script can use or ignore.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

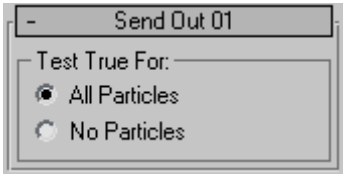
Send Out Test

Particle View (page 2–125) > Click Send Out in an event or add Send Out to the particle system and then select it.

The Send Out test simply sends all particles to the next event, or, conversely, keeps all particles in the current event. Use Send Out when you simply want to send particles to another event without any conditions.

Tip: You can temporarily convert any test to Send Out. To specify that a test should send all particles out without any conditions, click the left side of its icon in Particle view; the icon changes to a green light bulb to indicate that all particles automatically test True. Or, if you click the right side of the icon, it changes to a red light bulb, indicating that all particles test False and thus will stay in the current event. To revert to the test's original function, click its icon again.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Test True For—Lets you specify whether the test passes all particles on to the next event or keeps them in the current event. Default=All Particles.

- **All Particles—**All particles are passed on to the next event.
- **No Particles—**All particles are retained in the current event.

Spawn Test

Particle View (page 2–125) > Click Spawn in an event or add Spawn to the particle system and then select it.

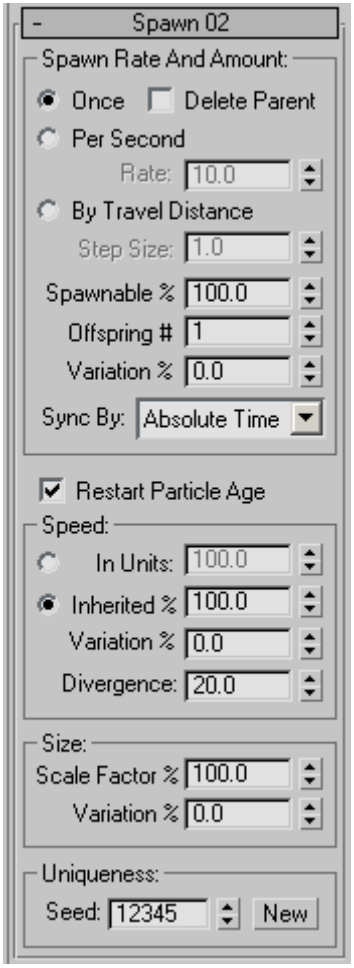
Spawn creates new particles from existing ones. Each spawned particle is born at the same location as its parent, and has the same orientation and shape. Spawn can give the spawned particles a different speed and scaling factor. If you wire the Spawn test to another event, spawned particles are sent to that event, where you can specify different properties for the new particles.

Spawn is a test only in that it sends the spawned particles to another event (if wired); it doesn't actually test any properties. All particles that encounter Spawn are immediately affected by it. Thus, if you want particles to spawn based on the results of a test, use a different test that branches to an event containing the Spawn. In such a case, you might want to then send the spawned particles to yet another event, or the particles will continually respawn. Alternatively, to spawn particles after a collision, use *Collision Spawn Test (page 2–215)*.

See also

Collision Spawn Test (page 2–215)

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

In the context of Spawn, a *parent* is the original particle from which new particles are spawned.

Spawn Rate And Amount group

Use these settings to specify how often particles are to spawn, the measurement system to use, and other values related to how many particles are spawned.

Once—Particles spawn one time only. For each existing particle, one new one is born.

Delete Parent—When on, deletes each original particle from which a new one is spawned. Available only with the Once option.

Per Second—Lets you specify a number of particles to spawn every second. For example, if you use the default Rate setting of 10.0, at 30 fps a new particle is born every three frames.

Rate—The number of particles to spawn per second. Available only with the Per Second option.

By Travel Distance—Lets you spawn new particles at regular intervals over the path of a moving parent particle.

Step Size—The system spawns a new particle every time the parent moves this distance, in system units.

Spawnable—The percentage of particles in the current event that will spawn new particles. This is determined once for each particle, when it enters the event. However, the parameter is animatable. Default=100.0.

For values other than 100.0, Spawnable uses a randomized selection process, which is affected by the Uniqueness Seed value. For example, with five parent particles, Offspring #=1, and Spawnable=80.0, you might get any number of spawned particles between two and five. The average per spawning would be four, however.

Offspring #—The number of new particles the system creates from each parent particle for each spawning event. Default=1.

Variation—The amount by which the Offspring # value can vary randomly. Default=0.0.

To obtain the actual test value for each particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then applies the result as a percentage of the Offspring # setting.

For example, if Offspring #=20 and Variation=10, then the actual number of offspring for each particle would be between 18 and 22.

Sync By—Choose the time frame to use when animating Rate, Step Size, Offspring #, and Variation:

- **Absolute Time—**Any keys set for parameters are applied at the actual frames for which they're set.
- **Particle Age—**Any keys set for parameters are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for parameters are applied to each particle starting when it first enters the event.

Restart Particle Age—When on, sets the age of each newly spawned particle at 0. When off, each spawned particle inherits its parent's age. Default=on.

Speed group

Lets you specify the speed of spawned particles in absolute terms or relative to the parents' speed, with optional random variation. The direction of a spawned particle is always in relation to that of its parent, but you can set a Divergence so they eventually spread out. Default=Inherited.

In Units—Choose this to specify the speed of spawned particles in system units per second. Default=100.0.

A positive value inherits the parent's direction; a negative value reverses it.

Inherited—Choose this to specify each spawned particle's speed as a percentage of its parent's speed. Default=100.0

A positive value inherits the parent's direction; a negative value reverses it.

Variation—The amount by which a spawned particle's speed can vary randomly. Default=0.0.

To obtain the actual speed for each spawned particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the particle's speed as specified or inherited. For example, if a particle's speed is 100 units/second and Variation=20, then the tested value for each particle would be between 80 and 120 units/second.

Divergence—When on, spreads out the stream of spawned particles. Use the numeric setting to define the extent of the divergence in degrees. Range=0 to 180. Default=0.

Size group

Scale Factor—The amount of uniform scaling to apply to each spawned particle, as a percentage of its parent's size. Default=100.0.

Variation—The amount by which a spawned particle's scale can vary randomly. Default=0.0.

To obtain the actual scaling for each spawned particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the Scale Factor value. For example, if Scale Factor=100 and Variation=20, then each spawned particle would be between 80 and 120 percent of its parent's size.

Uniqueness group

The Uniqueness setting enables randomization of the Spawnable result, when less than 100.0, as well as of the Variation values.

Seed—Specifies a randomization value.

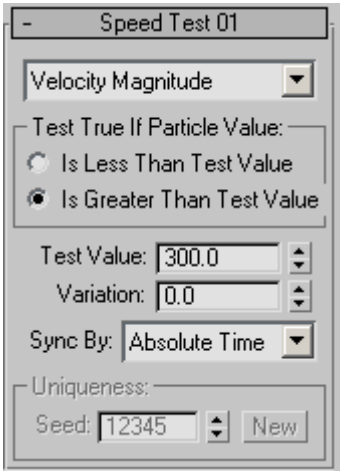
New—Calculates a new seed using a randomization formula.

Speed Test

Particle View (page 2–125) > Click Speed Test in an event or add Speed Test to the particle system and then select it.

Speed Test lets the particle system check particle speed, acceleration, or the rate of circular travel, and branch accordingly. The test provides a number of variants that let you test speed or acceleration on specific axes, or simply whether the particle is accelerating or decelerating.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

The first interface element is a drop-down list that lets you choose the type of measurement to test:

- **Velocity Magnitude—**Tests the particle velocity, in system units per second, without consideration of direction. This is the default test type.
- **Velocity X/Y/Z—**Tests the particle velocity on the specified axis, in system units per second, using the world coordinate system.

  An example of this option would be with fireworks: As the particles move upward, their velocity on the world Z axis is positive. When the reach the top of their trajectory and begin to move downward, their velocity on the world Z axis slows to 0, and then becomes negative. If you choose Velocity Z and Is Less Than Test Value, and set Test Value to 0.0, you can send particles to another event at the moment they begin to move downward.
- **Acceleration Magnitude—**Tests the particle acceleration (change in velocity), in system units per second per second, without respect to direction.
- **Acceleration X/Y/Z—**Tests the particle acceleration (change in velocity) on the specified axis, in system units per second per second, using the world coordinate system.
- **Steering Rate—**Tests the circular component of particle travel in degrees per second, without consideration of rotation or spinning.

  For example, if a particle travels along a parabolic path, its motion has both linear and circular components. The circular component is greatest at the top of the parabola. If a particle travels in a full circle in one second, the rate is 360; if it travels in a half circle, the rate is 180.

  Potential usage: When a particle is forced to turn too sharply, it might explode or change its type of movement. For example, missiles chase a jet fighter, which maneuvers to elude the missiles. The missiles are forced to change their course rapidly, but the missile construction cannot stand the fast change in steering, so the missiles blow up or disintegrate.

  Tip: You can test steering rate by setting the particle speed with *Speed By Icon (page 2–162)*, and linking the Speed By Icon operator icon to a circular path.
- **True When Accelerates—**Returns a True value when the particle velocity is increasing in value.

- **True When Decelerates—**Returns a True value when the particle velocity is decreasing in value.

Test True if Particle Value—Lets you specify whether the test passes particles on to the next event if the speed test succeeds or fails. Available for all tests except True When Accelerates/Decelerates. Default=Is Greater Than Test Value.

By default, Speed Test returns True if the value tested for exceeds the Test Value quantity, but you can choose Is Less Than Test Value as well. For example, if you use the Velocity Magnitude test type and set Test Value=200 and Variation=0, and choose Is Less Than Test Value, then particles will move to the next event only when they travel faster than 200 units per second. Any particles traveling 200 units per second or slower stay in the current event unless they later exceed that speed or another test returns True.

Test Value—The specific speed or acceleration to test for. The unit of measurement depends on the type of test; see above. Default=300.0.

Variation—The amount by which value tested for can vary randomly. Default=0.0.

To obtain the actual test value for each particle, the system multiplies the Variation value by a random number between -1.0 and 1.0, and then adds the result to the Test Value setting. For example, if Test Value=300 and Variation=10, then tested value for each particle would be between 290 and 310.

Sync By—Choose the time frame to use when animating Test Value and Variation. For further information, see *Animation Offset Keying group (page 2–158)*.

- **Absolute Time—**Any keys set for parameters are applied at the actual frames for which they're set.
- **Particle Age—**Any keys set for parameters are applied at the corresponding frames of each particle's existence.
- **Event Duration—**Any keys set for parameters are applied to each particle starting when it first enters the event.

Uniqueness group

The Uniqueness setting enables randomization of the test value variation.

Seed—Specifies a randomization value.

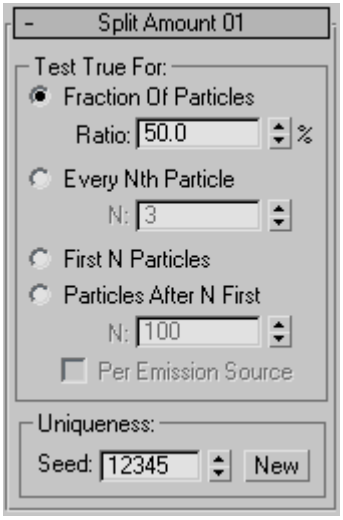
New—Calculates a new seed using a randomization formula.

Split Amount Test

Particle View (page 2–125) > Click Split Amount in an event or add Split Amount to the particle system and then select it.

The Split Amount test lets you send a specific number of particles to the next event, keeping all remaining particles in the current event. You can split the particle stream by a specific number or percentage, or by every Nth particle. With a specific number of particles, the splitting takes place once per event, but you can animate the percentage and "every Nth" values to vary the amount of split-off particles over time.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Test True For group

Lets you choose how to split the particle stream. Default=Fraction Of Particles, 50%.

Fraction Of Particles—Split the particle stream on a percentage basis, specified with the Ratio value. With Ratio greater than 0.0, whether a particular particle is split off depends on a randomization factor; change this with the Uniqueness Seed value.

Ratio—Specify the percentage of particles that will test True. Animatable. Default=50.0.

Available only with the Fraction Of Particles option.

Every Nth Particle—Splits off a regular sample of particles.

N—Specify how often to split off a particle. Animatable. Default=3.

For example, enter 8 to split off every eighth particle.

Available only with the Every Nth Particle option.

First N Particles—Splits off the number of particles specified as the N value (below), starting with the first particle to enter the event, and retains the rest.

Particles After N First—Splits off all particles starting with the first one after N particles, as specified with the N parameter (below). All particles starting with the first to enter the event up to N are retained in the event.

N—Specify the number of particles to split off, with First N Particles, or to retain in the event, with Particles After N First.

Available only with the First N Particles or the Particles After N First option.

Per Emission Source—When on, the software applies the N value for the options First N Particles and Particles After N First separately for each emission source.

Use this option with multiple Particle Flow sources converging into a single Split Amount test.

Uniqueness group

The Uniqueness setting enables randomization of particle retention with the Fraction Of Particles option.

Seed—Specifies a randomization value.

New—Calculates a new seed using a randomization formula.

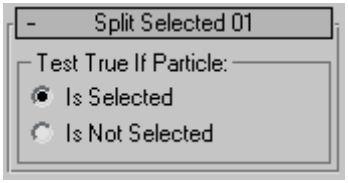
Split Selected Test

Particle View (page 2–125) > Click Split Selected in an event or add Split Selected to the particle system and then select it.

The Split Selected test lets you split the particle stream based on particles' selection status. For information about selecting particles, see *Selection rollout (page 2–138)*.

Note: This test considers only particles selected at the Particle sub-object level. To use all particles in an event at a specific frame, go to that frame, go to the Event sub-object level, and highlight the event. Then go to the Particle sub-object level, and on the Selection rollout, click Get From Event Level.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Test True If Particle group

Lets you choose how to split the particle stream. Default=Is Selected.

Is Selected—All selected particles are eligible for redirection to another event.

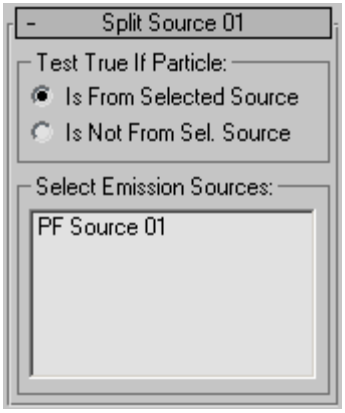
Is Not Selected—All non-selected particles are eligible for redirection to another event.

Split Source Test

Particle View (page 2–125) > Click Split Source in an event or add Split Source to the particle system and then select it.

The Split Source test lets you split the particle stream based on particles' origin. You can specify that particles from one or more specific Particle Flow sources should or should not be eligible for redirection to the next event.

Interface



The user interface appears in the parameters panel, on the right side of the Particle View dialog.

Test True If Particle group

Lets you choose how to split the particle stream. Default=Is From Selected Source.

Is From Selected Source—All particles from sources highlighted in the Selected Emission Sources list are eligible for redirection to another event.

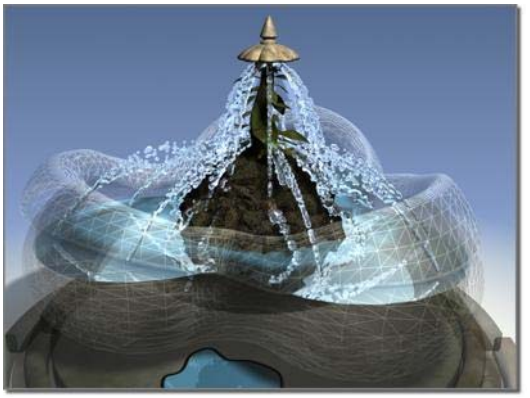
Is Not From Sel(ected) Source—All particles from sources *not* highlighted in the Selected Emission Sources list are eligible for redirection to another event.

Selected Emission Sources—Lists all emission sources in the system. Highlight the sources to be considered by the test.

Non-Event-Driven Particle Systems

Create panel > Geometry > Particle Systems

Create menu > Particles



Fountain spray created as a particle system

Non-event-driven particle systems provide relatively simple, straightforward methods for generating particle sub-objects over time for the purpose of simulating snow, rain, dust, and so on. You use particle systems primarily in animations.

3ds Max provides six built-in, non-event-driven particle systems:

Spray Particle System (page 2–244)

Snow Particle System (page 2–246)

Super Spray Particle System (page 2–249)

Blizzard Particle System (page 2–251)

PArray Particle System (page 2–256)

PCloud Particle System (page 2–253)

Note: You can take advantage of the AutoGrid feature to orient and position new particle systems with respect to existing objects. For details, see *AutoGrid (page 2–7)*.

Interface



Snowstorm created as a particle system

This topic describes only the general properties of particle systems. Other plug-in particle systems might be available in your configuration.

The built-in particle systems share these controls:

Emitter—Specifies where in the scene the particles are generated. The emitter is the particle system's main sub-object. It doesn't render. Particles appear on the surface of the emitter and fall (or drift, drop, flurry, spray) from the emitter in a particular direction.

Timing—The timing parameters control the dynamics of particles in the system. They specify how quickly particles appear, how quickly they disappear, whether the emission rate is constant, and so on.

Particle-specific parameters—These parameters are specific to the kind of particle system. Examples are particle size and speed.

Rendering properties—These parameters are also specific to the kind of particle system. There are options for displaying particles in viewports and rendering them in scenes and animations. Particles do not necessarily appear the same in renderings as they do in viewports.

You can modify and animate particle system parameters. You can also affect particle system behavior with *space warps (page 2–55)*. In addition, you can deform particle systems with the *Mesher compound object (page 1–374)*.

Note: Particles can participate in dynamics simulations.

See also

Particle System Usage (page 2–238)

Creating a Particle Emitter (page 2–239)

Particle Flow (page 2–109)

Using Particle Systems

You create particle systems when you want to model an object or effect that can best be described as a large collection of similar objects behaving in a similar fashion. Obvious examples of such effects include rain and snow, but other equally valid examples include water, smoke, ants, and even crowds of people.

On the Create panel, click Spray, Snow, Super Spray, Blizzard, PArray, or PCloud to create a particle system. Spray and Snow exist primarily for compatibility with earlier releases of 3ds Max, and are superseded by Super Spray and Blizzard.

To create a particle system, first choose Create menu > Particles > Spray or Snow.

These are the basic steps for creating a particle system:

1. Create a particle emitter. All particle systems require an emitter. Some particle systems use the particle system icon as the emitter while others use an object you select from the scene as the emitter.
2. Determine the number of particles. You set parameters such as birth rate and life span to control how many particles can exist at any given time.
3. Set particle shape and size. You can select from many standard particle types (including metaballs) or you can select an object to be emitted as a particle.
4. Set initial particle motion. You can set the speed, direction, rotation, and randomness of particles as they leave the emitter. Particles can also be affected by animation of the emitter.
5. Modify particle motion. You can further modify the motion of particles after they leave the emitter by binding the particle system to a space warp in the Forces group, such as Path Follow, or make them bounce off a deflector in the Deflectors space warp group, such as UDeflector.

   Important: When you use forces and deflectors together, always bind the forces before the deflectors.

   Tip: If the particles don't follow the emitter after it's moved, then change any *Path Follow parameter (page 2–71)*. The motion will be applied to the particles.

Rain and Snow

Create rain and snow using *Super Spray (page 2–249)* and *Blizzard (page 2–251)*. These particle systems are optimized for droplet (Super Spray) and tumbling flake (Blizzard) effects. Add space warps such as *Wind (page 2–75)* to create spring rains or winter storms.

Bubbles

Create bubbles by using the Bubble Motion options of *Super Spray (page 2–249)*. If you require good rendering speed, consider using constant or tetra particles. If you require bubble detail, consider using opacity-mapped facing particles, instanced spheres, or metaparticles.

Flowing Water

You generate flowing liquid effects by setting *Super Spray (page 2–249)* to generate closely packed

metaparticles. The metaparticles blob together forming a stream. Add a *Path Follow (page 2–71)* space warp to send the stream down a trough.

Explosions

Particle Array (PArray) (page 2–256) uses another object as its particle emitter. You can set the particle type to use fragments of the emitter object to simulate the object exploding.

Volume Effects

Particle Cloud (PCloud) (page 2–253) constrains its particles within a specified volume. You can use Particle Cloud to generate bubbles in a glass of soda, or bees buzzing inside a jar.

Crowds

Super Spray (page 2–249), *Blizzard (page 2–251)*, *Particle Array (page 2–256)*, and *Particle Cloud (page 2–253)* can use instanced geometry as their particle type. You can create a stream of ants, a flock of birds, or a cloud of dandelion seeds using instanced geometry particles.

Creating a Particle Emitter

After you choose a particle system type to add, you create the particle system icon in the scene. The icon serves different purposes depending on the type of particle system.

As an emitter: The icon defines the starting location and direction of the particles. Spray, Snow, Super Spray, Blizzard, and Particle Cloud use the icon as the particle emitter.

As a placeholder: The icon serves only to hold the parameters for the particle system. The particles are emitted from another selected object. Particle Array and Particle Cloud use the icon as a placeholder.

You can choose whether Particle Cloud uses its icon or another selected object as the emitter.

Procedures

To create emitter icons:

- Drag in a viewport to set the size and orientation of the particle emitter icon for Spray, Snow, Super Spray, Blizzard, and Particle Cloud.

  Spray, Snow, Blizzard, and Particle Cloud use the icon size as the area of particle emission. Super Spray emits particles from its center regardless of its icon size.

  All particle systems align the particle direction with the Z axis of the creation grid.

To create placeholder icons:

1. Drag in a viewport to set the size of the placeholder icon for Particle Array and Particle Cloud.

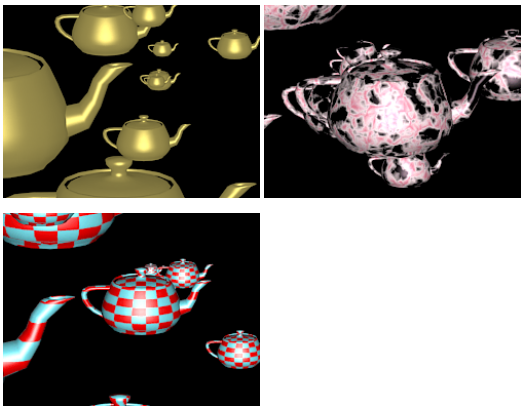
   The size and location of the particle system icon has no effect on the particles.

2. After placing the particle system icon, click Pick Object on the Basic Parameters rollout to select the object to use as the particle emitter.

Using Materials with Particle Array

The material that appears on the particles comes from one of three places:

- The particle system itself
- The material assigned to the distribution object (PArray only)
- The material assigned to the instanced objects



Particles with various materials assigned to them

You make this choice in the Material Mapping and Source group near the bottom of the Particle Type rollout. Choose Icon to use the material assigned to the particle system, choose Picked Emitter to use the material assigned to the distribution object, or choose Instanced Geometry to use the material assigned to the instanced objects. Note that this third option is available only when Instanced Geometry is the current particle type (selected in the Particle Types group at the top of the Particle Type rollout).

Note: SuperSpray and Blizzard don't have distribution objects. PCloud has a distribution object but you can't get the material from it. In these cases, a radio button lets you use the material from the instanced geometry. Only PArray lets you get the material from the distribution object.

Important: A particle-system object, like any other object, can carry only a single material at any time. Thus, if you choose either Picked Emitter or Instanced Geometry, an instance of the chosen material is actually copied to the particle system, overwriting the material assigned to the emitter. If you want to restore the emitter material, choose Icon, and then reassign the material from the Material Editor or the Browser.

No matter what choice you make, if the material used is not mapped, then all particles take on the surface properties of the material, regardless of which object is used as the source of the material.

Achieving Particle Motion Blur

Particle motion blur is actually the result of varying the opacity and the length of particles based on their speed. To accomplish this requires coordination between material assignment and the settings in the particle systems. Follow these instructions:

- Use the *Particle MBlur map (page 2–1676)* in the material that you assigned to the particles. For best results, assign it as an opacity map.
- Make sure that the particle system, PArray, PCloud, and Super Spray, or Spray, supports the Particle MBlur map.
- Choose Rotation and Collision rollout > Spin Axis Controls group > Direction of Travel/Mblur option.
- In this same group, set the Stretch spinner greater than 0 to stretch the particles as a percent of their length based on the particle Speed setting.
- Use the correct type of particle. MBlur works on all particle types except Constant and Facing.

  Note: Instanced objects with multi/sub-object materials cannot be image motion blurred.

Using Mapped Materials with Particle Systems

Mapped materials affect particles differently, depending on the source of the material.



Particles with a diffuse-mapped material

When you choose Icon, the mapping coordinates of the material are applied across the V (vertical) axis, from V=0 (the bottom edge of the map) to V=1 (the top edge of the map). The bottom edge of the map is applied at the birth of the particle, and the top edge at either the death of the particle (if Time is on) or the distance of the particle at its death (if Distance is on).

The Time spinner specifies the number of frames from birth that it takes to complete one mapping of a particle. Thus, if set to 15, the particle uses the bottom edge of the map at its birth, and moves through to the top edge of the map at frame 15.

The Distance spinner specifies the distance, in units, from birth that it takes to complete one mapping of a particle. Thus, if set to 50, as the particle moves along the normal vector, it displays the bottom edge of the map at birth, and the top edge at 50 units along the normal vector.

The one exception to this is when you use the Tetra particle type. In this case, each particle is always constantly mapped with V=0 at the head and V=1 at the tail.

When you choose Picked Emitter, the particles take on the color of the object at the point at which they're created. If the mapped surface is yellow where the particle emerges, then the particle is yellow.

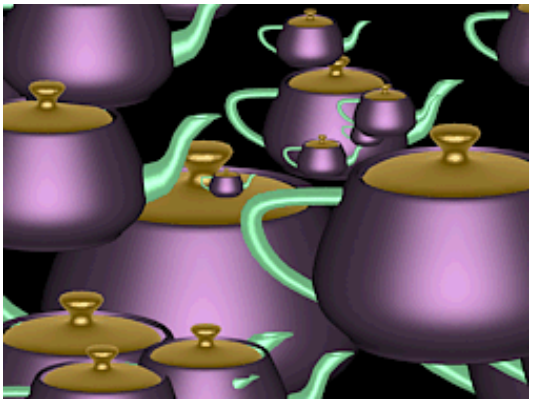
Again, Tetra particles are an exception and the distribution-object material is mapped from head to tail.

Fragment particles use the same technique, with additional options when the Thickness setting is greater than 0.

When Thickness is 0, all faces in the fragment are mapped the same as the portion of the object surface from which they're derived.

When Thickness is greater than 0, the outer faces of the fragment copy the surface of the distribution object, and are assigned the material ID specified in the Outside ID spinner in the Particle Type rollout > Fragment Materials group. The thickness edges use the Edge ID number, and the inner faces use Backside ID. Thus, by assigning a multi/sub-object material to the object-based emitter, you can specify a different material for the outer fragment surfaces, the edges, and the inner surfaces. Note that the Outside ID spinner defaults to a value of 0, which means "use whatever material is currently assigned." Change this to a specific sub-material number to assign a sub-material to the outside edges of the fragments.

Using Multi/Sub-Object Materials with Particle Systems



Particles with a Multi/Sub-Object material

If the assigned material is a Multi/Sub-Object material, particles are affected in the following ways, depending on the source of the material:

- **Icon:** In most cases, each particle, at its birth, is assigned a different sub-material, cycling through each available sub-material. For example, if there are only three sub-materials, the first particle receives sub-material #1, the second #2, the third #3, the fourth #1, and so on. The exceptions to this are as follows:
- MetaParticles use only the first sub-material.
- Object fragments are born at once, so all of them use only the first sub-material.
- **Picked Emitter:** When used with object fragments as particles, each particle uses the three ID numbers in the Fragment Materials group. If this source is used with the other particle types, the particles are assigned sub-materials in the same way as when Icon is chosen.
- **Instanced Geometry:** When used with Instanced Geometry particles, each particle is assigned a sub-material in exactly the same way as the source object, so each particle looks just like the source object. When used with other particle types, the particles are assigned sub-materials in the same way as when Icon is chosen.

  Note: Instanced objects with Multi/Sub-Object materials cannot be image motion blurred.

Using Spawned Particles

The examples in this topic start you out with a very simple spawning using the Super Spray particle system.

Example: Setting up the particle system:

1. In the Perspective viewport, create a pyramid that's **5** units in all dimensions.
2. Create a cylinder with a radius and height of **3**, and **6** sides.
3. In the center of the Perspective viewport, create a Super Spray emitter.
4. Set Speed in the Particle Motion group (Particle Generation rollout) to **3.0**.
5. On the Particle Type rollout, choose Instanced Geometry, then in the Instancing Parameters group click Pick Object, and select the pyramid.
6. In the Viewport Display group of the Basic Parameters rollout, choose Mesh and set the Percentage of Particles to **100**.
7. On the Particle Generation rollout, choose Use Rate, and set the spinner to **1**.
8. Set Emit Start and Emit Stop to **0**, set Display Until to **100**, set Life to **10**, and leave Variation at **0**.
9. Under Particle Size, set Grow For and Fade For to **0**.
10. Drag the time slider between frames 0 and 15.

    A single pyramid particle emerges from the emitter, moves up, and then dies and disappears at frame 10.

Example: Adding spawning effects:

1. On the Particle Spawn rollout, choose Spawn on Death.
2. Leave Spawns spinner at 1, and set the Multiplier to **2**.
3. Set the Direction Chaos spinner to **50**.
4. Drag the time slider slowly over frames 8 to 25 (approximately).

   At frame 10, two pyramids appear at the death of the original particle (because of the Multiplier setting), and move off in different directions (because of the Direction Chaos setting). All particles die after frame 20 because only one generation of spawned particles is specified.
5. Set Spawns to **4**, and the Multiplier to **4**.
6. Drag the time slider over frames 8 to 50.

   With each spawned generation the particles increase exponentially.

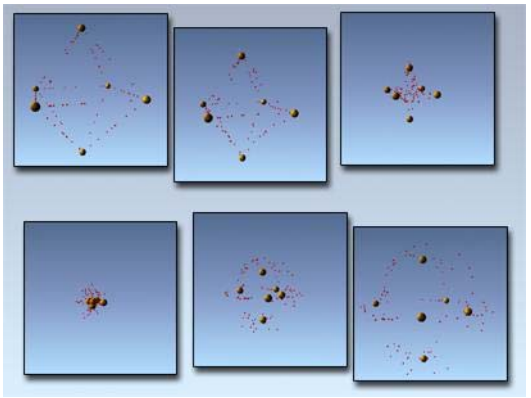
Example: To mutate the objects:

1. In the Object Mutation Queue group, click Pick Object, and then select the cylinder.

   The name of the cylinder appears in the list window.
2. Click Pick Object and select the pyramid. Then click Pick Object again, and select the cylinder. Your list now reads: Cylinder01, Pyramid01, Cylinder01.
3. Drag the time slider over frames 0 to 50.

   The pyramid appears over frames 0 to 10, then changes to four cylinders at frame 11, then 16 pyramids at the next spawning, and so on.

Using Interparticle Collision



Particles colliding and then rebounding

You can set up particles to detect collisions with each other. This can be useful when the particles are meant to model solid objects such as marbles.

Procedure

Example: To create particles that collide with each other:

1. Create a Super Spray particle emitter, and place a *Deflector space warp (page 2–90)* a short distance from it with the surface of the deflector perpendicular to the stream of the particles. Bind the deflector to the Super Spray.
2. Set the Super Spray values as follows:
   - Off Axis: **1**, Spread: **1**, Off Plane: **180**, Spread: **180**
   - Mesh: **chosen**, Percent of Particles: **100**
   - Particle Quantity: Use Rate **chosen** and set to: **1**
   - Speed: **3**, Variation: **100%**
   - Emit Start: **0**, Emit Stop: **5**, Display Until: **100**, Life: **100**
   - Size: **4.0**, Grow For: **0**, Fade For: **0**
   - Particle Type: **Sphere**

3. Drag the time slider so you can see the particle spheres bounce off the deflector. Note that the rebounding particles move through each other.
4. On the Rotation & Collision rollout, turn on Enable in the Interparticle Collisions group. View the animation again. This time, the particles bounce off each other.

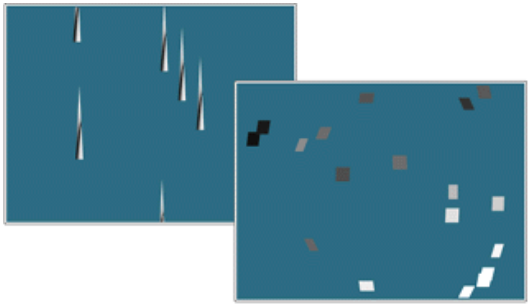
   Tip: InterParticle Collisions, Deflector Binding, and Bubble Noise do not get along well together. Particles may leak through the deflector when these three are used together. Instead of bubble motion use animated mapping. Use facing particles with an animated map of a bubble, where the bubble is smaller than the map size. The bubble is animated moving around the map. This simulates bubble motion at the map level.

Spray Particle System

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > Spray

Create menu > Particles > Spray

Spray simulates water drops such as rain, a fountain, the spray from a garden hose, and so on.



Kinds of Spray

Note: *SuperSpray (page 2–249)* is a more powerful and advanced version of Spray. It provides all the functionality of Spray, plus additional features.

Tip: To animate particles following a path through space, use the *Path Follow space warp (page 2–71)*.

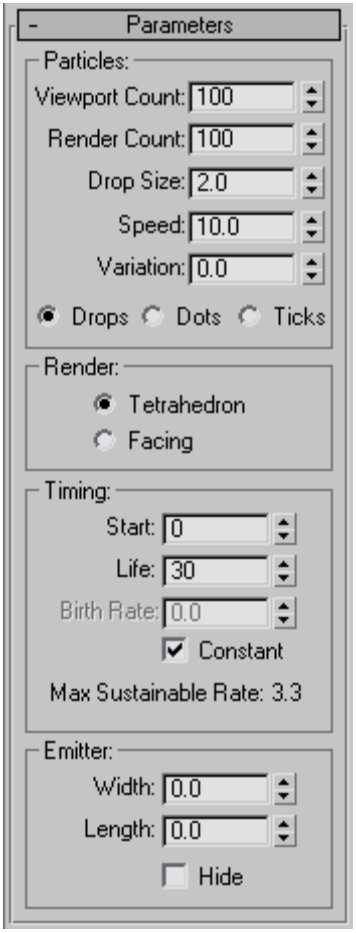
Procedure

To create spray:

1.  On the Create panel, make sure the Geometry button is active and Particle Systems is chosen in the object category list, then click Spray.
2. Drag in a viewport to create the Spray emitter; see *Creating a Particle Emitter (page 2–239)*.

   The emitter's direction vector points in the negative Z direction of the active construction plane. For example, if you create the emitter in the Top viewport, the particles will move downward in the Front and Left viewports.

Interface



Particles group

Viewport Count—Maximum number of particles displayed in viewports at any given frame.

Tip: Setting the viewport count less than the render count can improve viewport performance.

Render Count—Maximum number of particles that can appear in a single frame when you render it. Works in combination with the particle system's timing parameters.

- When the number of particles reaches the Render Count value, particle creation is suspended until some particles die.
- When enough particles die, particle creation resumes until Render Count is reached again.

Drop Size—Size of a particle in the active units.

Speed—Initial velocity of each particle as it leaves the emitter. Particles travel at this speed unless they are affected by a particle system space warp.

Variation—Varies the initial speed and direction of particles. The greater the Variation, the stronger and broader the spray.

Drops, Dots, or Ticks—Choose how particles are displayed in viewports. The display setting does not affect how particles are rendered. Drops are streaks that appear like raindrops, dots are points, ticks are small plus signs.

Render group

Tetrahedron—Particles are rendered as long tetrahedrons, with the length you specify in the Drop Size parameter. Tetrahedron is the default setting for rendering. It provides a basic simulation of water drops.

Facing—Particles are rendered as square faces whose width and height equals the Drop Size. Facing particles always face the camera (or the user's perspective). They are provided especially for use with material maps. Use with an appropriate opacity map for bubbles or snowflakes.

Note: Facing works correctly only in a perspective or camera view.

Timing group

Timing parameters control the "birth and death" rates of emitted particles.

At the bottom of the Timing group is a line that displays the maximum sustainable rate. This value

is based on the Render Count and the lifetime of each particle. To be precise:

maximum sustainable rate=Render Count/Life

Because the number of particles in a frame never exceeds Render Count, if the Birth Rate exceeds the maximum rate, the system will run out of particles, pause until some die off, and then start again, generating particles in bursts or spurts.

Start—Number of the first frame where particles appear.

Life—The lifetime of each particle, in number of frames.

Birth Rate—The number of new particles born per frame.

If this is less than or equal to the maximum sustainable rate, the particle system generates an even flow of particles. If it is greater than the maximum rate, the particle system generates particles in bursts.

You can animate the Birth Rate parameter.

Constant—When on, Birth Rate is unavailable and the birth rate used equals the maximum sustainable rate. When off, Birth Rate is available. Default=on.

Turning Constant off does not mean that the birth rate varies automatically; the birth rate remains constant unless you animate the Birth Rate parameter.

Emitter group

The emitter specifies the area where particles appear in the scene. It has a geometry you can display in viewports, but it isn't renderable.

The emitter is displayed as a rectangle with a vector pointing out of one side. The vector shows the direction in which the system emits particles.

Width and Length—You implicitly set the initial value of these parameters when you drag in a viewport to create the emitter. You can adjust the values in the rollout.

The space occupied by the particle system at any given time is the result of a combination of its initial parameters (such as size of the emitter, and speed and variation of emission) and any space warps that have been applied.

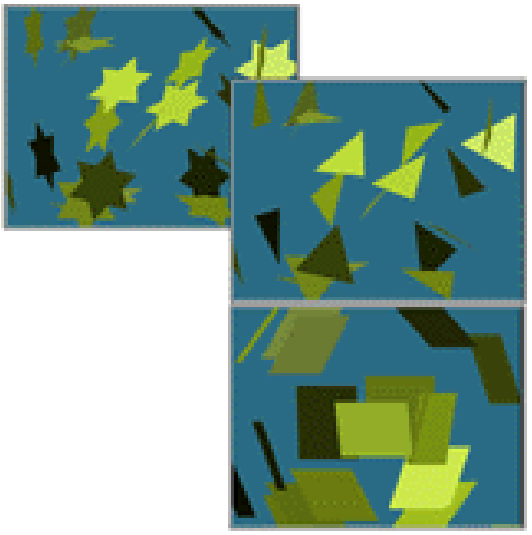
Hide—Turn on to hide the emitter in viewports. When Hide is off, the emitter is displayed in viewports. The emitter is never rendered. Default=off.

Snow Particle System

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > Snow

Create menu > Particles > Snow

Snow simulates falling snow or confetti. The snow system is similar to Spray, but it has additional parameters to generate tumbling snowflakes, and its rendering options are different.



Kinds of Snow

Note: *Blizzard (page 2–251)* is a more powerful and advanced version of Snow. It provides all the functionality of Snow, plus additional features.

Tip: To animate particles following a path through space, use the *Path Follow space warp (page 2–71)*.

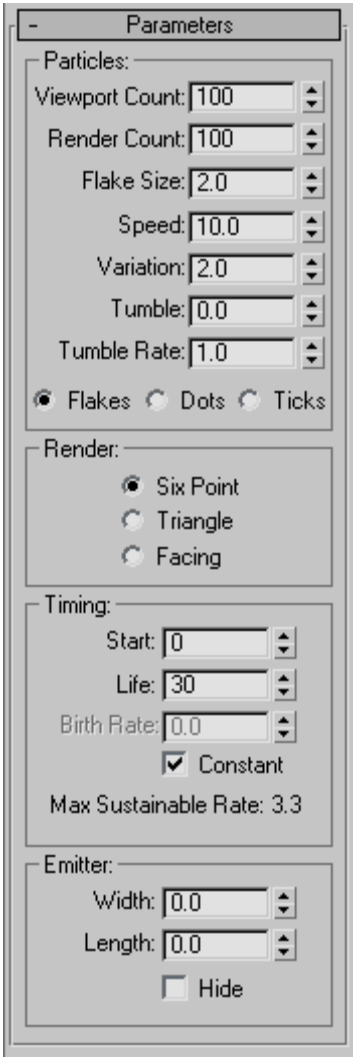
Procedure

To create snow:

1.  In the Create panel, make sure the Geometry button is active and Particle Systems is selected from the object category drop-down list, then click Snow.
2. Drag in a viewport to create the Snow emitter; see *Creating a Particle Emitter (page 2–239)*.

   The emitter's direction vector points in the negative Z direction of the active construction plane.

Interface



Particles group

Viewport Count—Maximum number of particles displayed in viewports at any given frame.

Tip: Setting the viewport count less than the render count can improve viewport performance.

Render Count—Maximum number of particles that can appear in a single frame when you render it.

Works in combination with the particle system's timing parameters.

- When the number of particles reaches the Render Count value, particle creation is suspended until some particles die.
- When enough particles die, particle creation resumes until Render Count is reached again.

Flake Size—Size of a particle in the active units.

Speed—Initial velocity of each particle as it leaves the emitter. Particles travel at this speed unless they are affected by a particle system space warp.

Variation—Varies the initial speed and direction of particles. The greater the Variation, the broader the area of snowfall.

Tumble—Amount of random rotation for snowflake particles. This parameter can range from 0 to 1. At 0, flakes do not rotate; at 1, they rotate the most. The axis of rotation is generated randomly for each particle.

Tumble Rate—Speed at which snowflakes rotate. The greater the Tumble Rate, the faster the rotation.

Flakes, Dots, or Ticks—Select how particles are displayed in viewports. The display setting does not affect how particles are rendered. Flakes are star-shaped snowflakes, dots are points, ticks are small plus signs.

Render group

Six Point—Each particle is rendered as a six-pointed star. Each side of the star is a face to which you can assign a material. This is the default setting for rendering.

Triangle—Each particle is rendered as a triangle. Only one side of the triangle is a face to which you can assign a material.

Facing—Particles are rendered as square faces whose width and height equals the Drop Size. Facing particles always face the camera (or the user's perspective). They are provided especially for use with material maps. Use with an appropriate opacity map for bubbles or snowflakes.

Note: Facing works correctly only in a perspective or camera view.

Timing group

Timing parameters control the "birth and death" rates of emitted particles.

At the bottom of the Timing group is a line that displays the maximum sustainable rate. This value is based on the Render Count and the lifetime of each particle. To be precise:

maximum sustainable rate=Render Count/Life

Because the number of particles in a frame never exceeds Render Count, if the Birth Rate exceeds the maximum rate, the system will run out of particles, pause until some die off, and then start again, generating bursts or spurts of particles.

Start—Number of the first frame where particles appear.

Life—The lifetime of a particle, in number of frames.

Birth Rate—The number of new particles born per frame.

If this is less than or equal to the maximum sustainable rate, the particle system generates an even flow of particles. If it is greater than the maximum rate, the particle system generates particles in bursts.

You can animate the Birth Rate parameter.

Constant—When on, Birth Rate is unavailable and the birth rate used equals the maximum sustainable rate. When off, Birth Rate is available. Default=on.

Turning Constant off does not mean that the birth rate varies automatically; the birth rate remains constant unless you animate the Birth Rate parameter.

Emitter Group

The emitter specifies the area where particles appear in the scene. It has a geometry you can display in viewports, but it isn't renderable.

The emitter is displayed as a rectangle with a vector pointing out of one side. The vector shows the direction in which the system emits particles.

You set emitter parameters in the Emitter group of the particle system's Parameters rollout.

Width and Length—You implicitly set the initial value of these parameters when you drag in a viewport to create the emitter. You can adjust the values in the rollout.

The space occupied by the particle system at any given time is the result of a combination of its initial parameters (size of the emitter, and speed and variation of emission) and any space warps that have been applied.

Hide—Turn on to hide the emitter in viewports. When off, the emitter is displayed in viewports. The emitter is never rendered. Default=off.

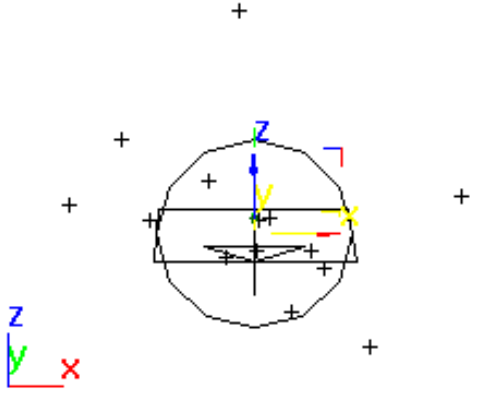
Super Spray Particle System

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > Super Spray

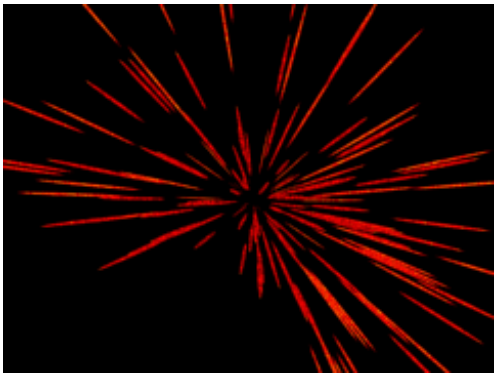
Create menu > Particles > Super Spray

Super Spray emits a controlled spray of particles. This particle system is like the simple Spray particle system with the added power provided by all the newer particle systems.

Tip: To animate particles following a path through space, use the *Path Follow space warp (page 2–71)*.



Super Spray viewport icon (emitter)



Particles emerging from a super spray system

Procedure

To create a super spray particle system:

1.  On the Create panel, make sure the Geometry button is active and Particle Systems is chosen in the object category list, then click Super Spray.
2. Drag in any viewport to create the Super Spray emitter icon; see *Creating a Particle Emitter (page 2–239)*.

   The icon appears as an intersecting plane and circle with an arrow. The initial direction of the

spray (based on the orientation of the emitter icon and indicated by the icon arrow) depends on the viewport in which you create the icon. Generally, the particles spray toward you when created in an orthographic viewport, or spray upward when created in a Perspective viewport.

3. Adjust the various parameters to alter the spray effect.

Interface

Note: This section describes the Particle Formation group in the Basic Parameters rollout, and the Particle Motion group in the Particle Generation rollout. These are the only controls unique to Super Spray. The other Super Spray rollouts and their contents are identical with those in Particle Array, except that Object Fragments and associated settings are not available on the Particle Type rollout. See *PArray Particle System (page 2–256)* for details or choose from the following list for rollout information.

Particle Generation Rollout (page 2–260)

Particle Type Rollout (page 2–262)

Rotation and Collision Rollout (page 2–268)

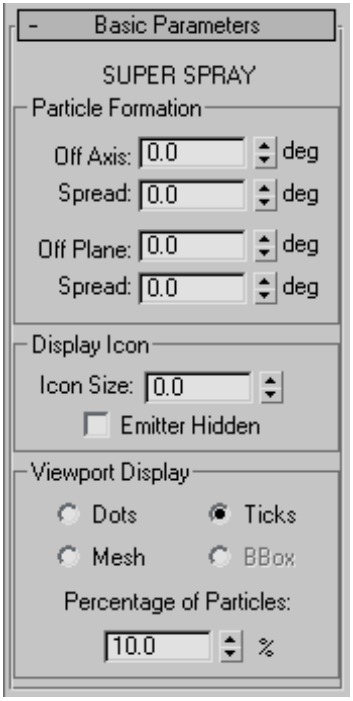
Object Motion Inheritance Rollout (page 2–269)

Bubble Motion Rollout (page 2–270)

Particle Spawn Rollout (page 2–271)

Load/Save Presets Rollout (page 2–274)

Basic Parameters rollout > Particle Formation group



These items control the direction and spread of particles.

Off Axis—Affects the angle of the particle stream off the Z axis (along the plane of the X axis).

Spread—Affects the spread of the particles away from the emission vector (along the plane of the X axis).

Off Plane—Affects the angle of emission about the Z axis. This has no effect if Off Axis is set to 0.

Spread—Affects the spread of the particles about the Off Plane axis. This has no effect if Off Axis is set to 0.

Particle Generation rollout > Particle Motion group

Speed—The speed of the particle at birth, in units per frame.

Variation—Applies a percentage of variation to the speed of emission for each particle.

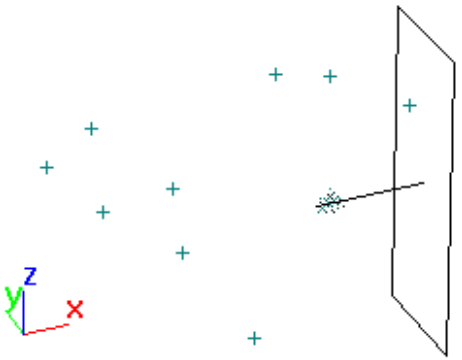
Blizzard Particle System

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > Blizzard

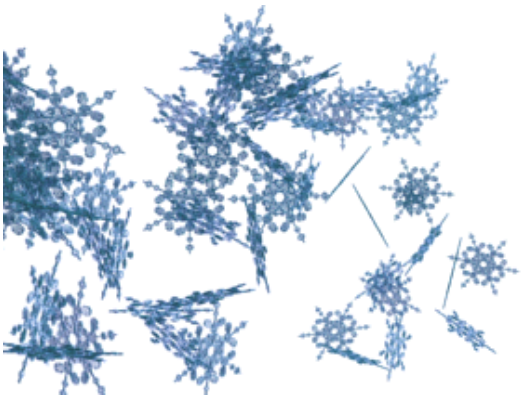
Create menu > Particles > Blizzard

This is an advanced version of the original Snow particle system.

Tip: To animate particles following a path through space, use the *Path Follow space warp (page 2–71)*.



Blizzard viewport icon (emitter)



Snowflake particles in a blizzard

Procedures

To create a blizzard particle system:

1.  On the Create panel, make sure the Geometry button is active and Particle Systems is chosen in the object category list, then click Blizzard.
2. Drag in a viewport to create the Blizzard emitter; see *Creating a Particle Emitter (page 2–239)*.

   The icon appears as a plane with a perpendicular line pointing in the direction of emission.
3. Adjust the various parameters on the command panel.

Interface

This section describes the Display Icon group in the Basic Parameters rollout, the Particle Motion group in the Particle Generation rollout, and the Material Mapping and Source group in the Particle Type rollout. These are the only controls unique to Blizzard. The other Blizzard rollouts and their contents are identical with those in Particle Array, except that Object Fragments and associated settings are not available on the Particle Type rollout. See *PArray (page 2–256)* for details or choose from the following list for rollout information.

Particle Generation Rollout (page 2–260)

Particle Type Rollout (page 2–262)

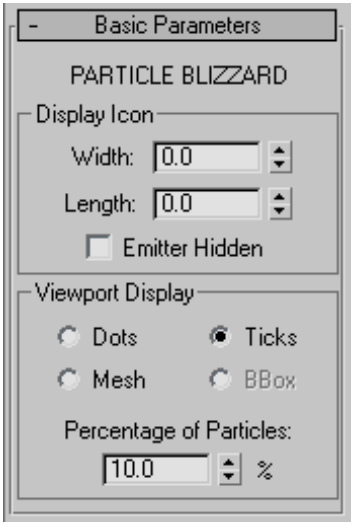
Rotation and Collision Rollout (page 2–268)

Object Motion Inheritance Rollout (page 2–269)

Particle Spawn Rollout (page 2–271)

Load/Save Presets Rollout (page 2–274)

Basic Parameters rollout > Display Icon group



The emitter specifies the location where particles are generated in the scene. It has a geometry you can display in viewports, but it isn't renderable.

The emitter is displayed as a rectangle with a vector pointing out of one side. The vector shows the direction in which the system emits particles.

You set emitter parameters on the particle system's Basic Parameters rollout, in the Display Icon group.

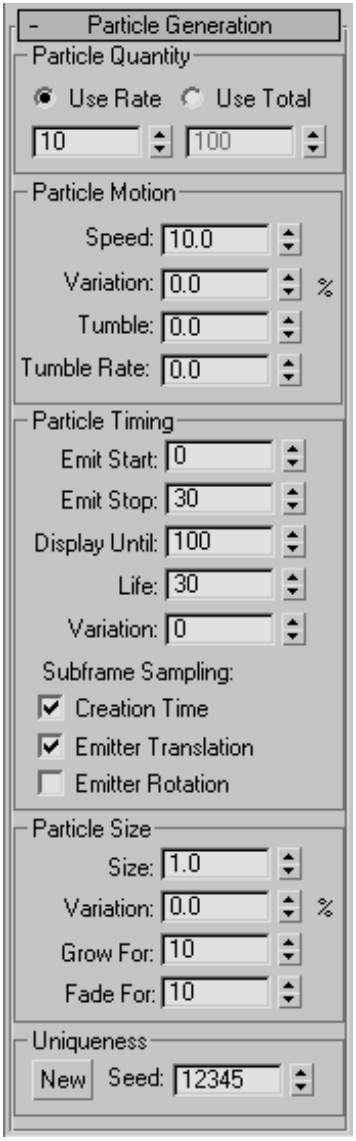
Width and Length—You implicitly set the initial value of these parameters when you drag in a viewport to create the emitter. You can adjust the values on the rollout.

The space occupied by the particle system at any given time is the result of a combination of its initial parameters (such as size of the emitter, and speed and variation of emission) and any space warps that have been applied.

Emitter Hidden—Hides the emitter in viewports. When off, the emitter is displayed in viewports. The emitter is never rendered. Default=off.

Particle Generation rollout > Particle Motion group

Specifies the number, size, and motion of particles.



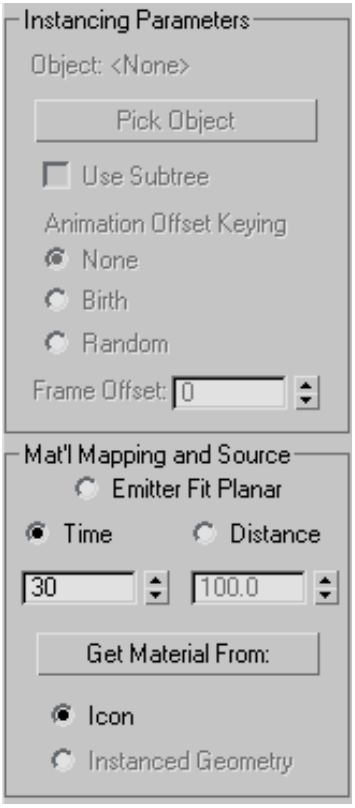
Speed—The speed of the particle at birth, in units per frame.

Variation—Applies a percentage of variation to the speed of emission for each particle.

Tumble—Amount of random rotation of the particles.

Tumble Rate—Speed at which the particles rotate.

Particle Type rollout > Mat'l Mapping and Source group



Emitter Fit Planar—Maps particles at birth, based on their point of emission from the rectangular Blizzard emitter icon. The UV range of the mapped material runs from 0 to 1 over the width and length of the emitter.

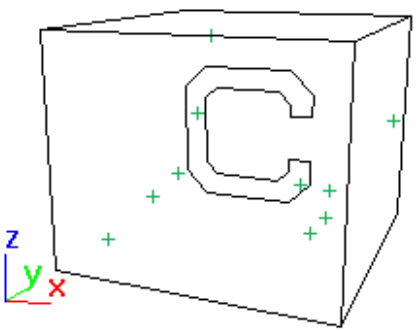
PCloud Particle System

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > PCloud

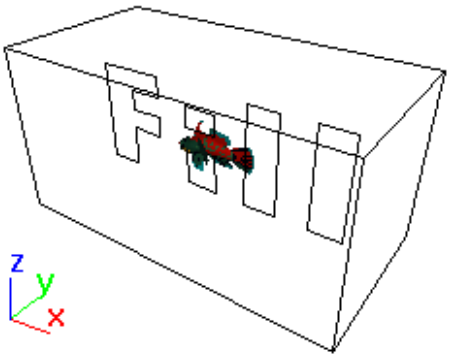
Create menu > Particles > PCloud

Use the PCloud (or Particle Cloud) particle system when you want a "cloud" of particles that fill a specific volume. PCloud can provide a flock of birds, a starfield, or a troupe of soldiers marching over terrain. You can confine the particles using basic supplied volumes of a box, sphere, or cylinder, or you can use any renderable object in the scene as a volume as long as that object has depth. Two-dimensional objects do not work with PCloud.

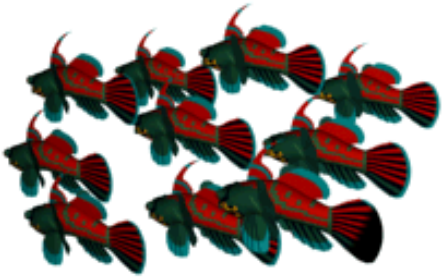
Tip: There is no automatic way to hide the object chosen as the object-based emitter. Hide it by using Hide Selected on the Display panel, or by applying a Hide key in Track View.



PCloud viewport icon (default emitter)



PCloud viewport icon (object-based emitter)



PCloud used to form a school of fish (each fish is a particle)

Procedure

To create a particle cloud:

1.  On the Create panel, make sure the Geometry button is active and Particle Systems is chosen in the object category list, then click PCloud.
2. Drag in a viewport to create the PCloud emitter; see *Creating a Particle Emitter (page 2–239)*.

   Adding a PCloud emitter works the same way as creating a box primitive: First drag out the length and width, then release the mouse button and move the mouse vertically to set the height, and then click to finish.

   The emitter appears with a letter "C" representing the particle cloud.
3. Adjust the various parameters on the command panel.

Interface

This section describes the Object-Based Emitter, Particle Formation, and Display Icon groups in the Basic Parameters rollout, and the Particle Motion group in the Particle Generation rollout. These are the only controls unique to PCloud. The other PCloud rollouts and their contents are identical with those in Particle Array, except that Object Fragments and associated settings are not available on the Particle Type rollout. See *PArray (page 2–256)* for details or choose from the following list for rollout information.

Particle Generation Rollout (page 2–260)

Particle Type Rollout (page 2–262)

Rotation and Collision Rollout (page 2–268)

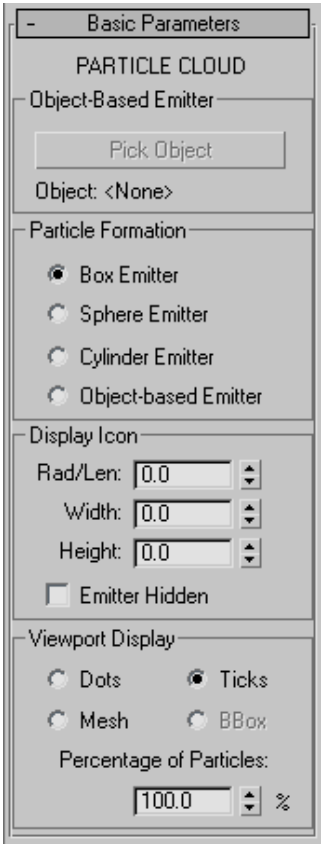
Object Motion Inheritance Rollout (page 2–269)

Bubble Motion Rollout (page 2–270)

Particle Spawn Rollout (page 2–271)

Load/Save Presets Rollout (page 2–274)

Basic Parameters rollout



Object-Based Emitter group

This button lets you select a renderable mesh object to use as a particle emitter. This object is used only when the Object-Based Emitter option is chosen in the Particle Formation group.

Pick Object—Click this, and then select a renderable mesh object to be used as a custom emitter.

Object—Displays the name of the picked object.

Particle Formation group

These options let you specify the shape of the emitter.

Box Emitter—Chooses a box-shaped emitter.

Sphere Emitter—Chooses a sphere-shaped emitter.

Cylinder Emitter—Chooses a cylindrical emitter.

Object-Based Emitter—Chooses the object selected in the Object-Based Emitter group.

Note: With regard to animation of the object-based emitter, the particles will properly fill a deformed object at frame 0, but they can't stay with the emitter while it's moving. If the emitter moves in a straight line, this can give the appearance of the cloud moving with the emitter.

Display Icon group

These options let you adjust the dimensions of the emitter icon when a custom object is not used as an emitter. When a custom object is used you can still resize the "Fill" icon using these options.

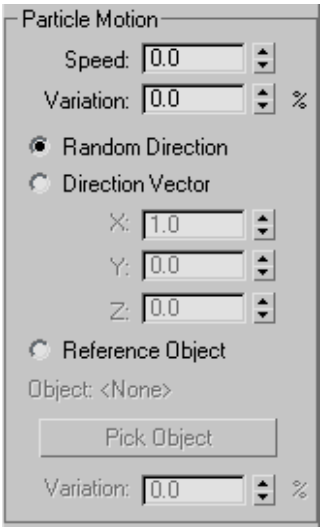
Rad/Len—Adjusts the radius of a spherical or cylindrical icon, and the length of a box icon.

Width—Sets the width of a box emitter.

Height—Sets the height of a box or cylindrical emitter.

Emitter Hidden—Hides the emitter.

Particle Generation rollout > Particle Motion group



Speed—The speed of the particle at birth along the normal, in units per frame.

Note: For the correct volume effect, Speed should be set to 0.

Variation—Applies a percentage of variation to the speed of emission for each particle.

Random Direction—One of three options that affect the direction of the particles. This option emits particles in random directions.

Direction Vector—Specifies the direction of the particles by a vector defined by the three X, Y, and Z spinners.

X/Y/Z—Displays the particle direction vectors.

Reference Object—Emits particles in the direction of the local Z axis of a specified object.

Object—Displays the name of the picked object.

Pick Object—Click this, and then select an object in the scene to use as a reference object. This button is available only when you choose Reference Object.

Variation—Applies a percentage of variation to the direction when you choose either the Direction Vector or Reference Object option. This spinner is unavailable and has no effect when you choose Random Direction.

PArray Particle System

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > PArray

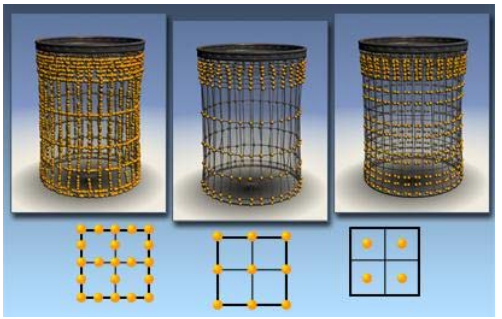
Create menu > Particles > PArray



Basket used as a distribution object for random particles over its surface

The PArray (Particle Array) particle system provides two types of particle effects:

- You can use it to emit particles using a selected geometric object as the emitter template (or pattern) for the emission. This object is referred to here as the distribution object.



How particles can be distributed on an object:
Left: Edges
Center: Vertices
Right: Faces

- You can use it to create sophisticated object explosions.

  Tip: A good way to create explosions is to set the particle type to Object Fragments and then apply a *PBomb space warp (page 2–68)*.

See also

Particle System Usage (page 2–238)

Using Materials with Particle Array (page 2–239)

Achieving Particle Motion Blur (page 2–240)

Using Mapped Materials with Particle Systems (page 2–240)

Using Multi/Sub-Object Materials with Particle Systems (page 2–242)

Using Spawned Particles (page 2–242)

Using Interparticle Collision (page 2–243)

Procedures

To set up a particle array:

1. Create an object that will become the distribution object, providing the emitter pattern (or exploded object) for the particle array.
2.  On the Create panel, make sure the Geometry button is active and Particle Systems is chosen in the object category list, then click PArray.
3. Drag anywhere in a viewport to create the particle-system object; see *Creating a Particle Emitter (page 2–239)*.
4. On the Basic Parameters rollout, click Pick Object, and then click the object to use as the distribution object.
5. Adjust the various options in the Particle Array rollouts to achieve the effect you want.

   The distribution object can be any object containing renderable faces.

   The particle-system object does not appear in the rendered scene. Its placement, orientation or size in the scene has no influence on the particle effect. To access the Particle Array parameters after creation, open the Modify panel and select the particle-system object (or click any visible particles in the viewport). To transform or otherwise modify the distribution object, select the distribution object itself.

   Also, you can share a single distribution object among more than one particle-system object. The distribution object merely provides the template for the particles, which are actually generated by the particle system.

Example: To assign three different mapped materials to fragments:

1. Apply mapping coordinates to the distribution object, either by turning on Generate Mapping Coords, if necessary, or by applying a *UVW Map modifier (page 1–922)*.
2. Assign a *Multi/Sub-Object material (page 2–1594)* to the distribution object.
3. Set up the first three sub-materials to be mapped materials.

4. Select the PArray icon.
5. On the Particle Type rollout in the Mat'l Mapping And Source group, choose Picked Emitter.
6. In the Particle Type rollout > Particle Types group, turn on Object Fragments.
7. Make sure the three spinners in the Fragment Materials group are set to 1, 2, and 3, respectively (or match the numbers with the sub-materials you've assigned in your multi/sub-object material).

   Tip: Using a complex distribution object for object-fragment particles can really slow down viewport interaction. You can use a simple stand-in object as the distribution object, and then later use the File/Replace command to replace the distribution object with a more complex object of the same name.

Interface

Basic Parameters Rollout (PArray) (page 2–258)

Particle Generation Rollout (page 2–260)

Particle Type Rollout (page 2–262)

Rotation and Collision Rollout (page 2–268)

Object Motion Inheritance Rollout (page 2–269)

Bubble Motion Rollout (page 2–270)

Particle Spawn Rollout (page 2–271)

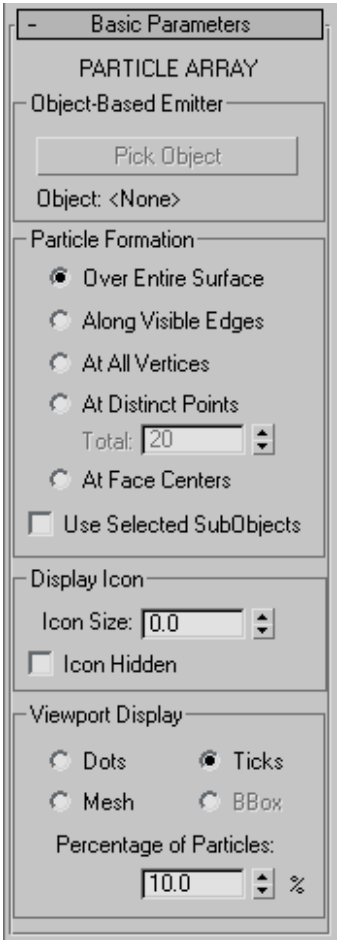
Load/Save Presets Rollout (page 2–274)

Basic Parameters Rollout (PArray)

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > Parray > Basic Parameters rollout

The items on the Basic Parameters rollout let you create and adjust the size of the particle system, and pick the distribution object. In addition, they let you specify the initial distribution of the particles in relation to the geometry of the distribution object, and the initial velocity of the particles from the distribution object. From here, you can also specify how the particles appear in the viewport.

Interface



Object-Based Emitter group

Pick Object—After you create the particle-system object, the Pick Object button becomes available. Click this button, and then click to select an object in your scene. The selected object becomes the

object-based emitter, and is used either as the source geometry over which particles form, or the source geometry used to create particles that appear to be fragments of the object.

Object text field—Displays the name of the picked object.

Particle Formation group

These options determine how standard particles are initially distributed over the surface of the object-based emitter. These controls are available only when the picked object is used as a distribution grid for standard particles, MetaParticles, or instanced geometry; see *Particle Type rollout (page 2–262)*. When Object Fragments is chosen in the Particle Type rollout, these controls are unavailable.

Over Entire Surface—Emits particles randomly over the entire surface of the object-based emitter. This is the default choice.

Along Visible Edges—Emits particles randomly from the visible edges of the object.

At All Vertices—Emits particles from the vertices of the object.

At Distinct Points—Places a specified number of emitter points randomly over the surface of the object.

Total—Specifies the number of emitter points used when At Distinct Points is chosen.

At Face Centers—Emits particles from the center of each triangular face.

Use Selected SubObjects—With mesh-based emitters, and to a limited extent with patch-based emitters, limits the source of the particle stream to the sub-object selection passed up the modifier stack in the object-based emitter. For example, if your emitter object is a cylinder converted to an editable mesh, and the top cap of that cylinder is selected at the Face or Polygon sub-object level, if Use Selected SubObjects is on *and* Particle Formation group > At Face Centers is on, the particles will stream only from the top cap of the cylinder. Default=off.

The type of particle formation you specify determines the type of sub-object geometry used, as follows:

- **Over Entire Surface—**Faces
- **Along Visible Edges—**Edges
- **At All Vertices—**Vertices
- **At Distinct Points—**Faces
- **At Face Centers—**Faces

If you've converted your object to an editable mesh, and selected different sub-object sections of it with vertex, edge, and face selection, as you switch particle formation options, you'll see the particles emit from different areas of the object.

Note: Use Selected SubObjects is applicable to patch object emitters only at the patch and element sub-object levels, and is not applicable NURBS objects used as emitters.

Tip: You can best see the emission patterns by first setting Speed on the Particle Generation rollout > Particle Motion group to 0. Move to a frame in which the particles appear, and then choose the various particle formation options.

Display Icon group

Adjusts the display of the particle-system icon in the viewports. (The particle-system icon is usually called the "emitter." In this case, however, it doesn't actually emit particles, so we're avoiding the term.)

Icon Size—Sets the overall size of the icon, in units.

Icon Hidden—When on, the PArray icon is hidden in the viewports. Note that the icon does not render, in any case. Default=off.

Viewport Display group

Specifies how the particles are displayed in the viewports.

Dots—Displays the particles as dots.

Ticks—Displays the particles as crosses.

Mesh—Displays the particles as mesh objects. This results in slower viewport redraws.

BBox—For instanced geometry only, this displays each instanced particle, whether a single object, a hierarchy, or a group, as a bounding box.

Percentage of Particles—This spinner specifies the number particles displayed in the viewports as a percentage of the number of rendered particles. Default=10 percent.

Set the display percentage to 100 percent if you want to see the same number of particles as will be rendered in your scene. However, this can considerably slow viewport display.

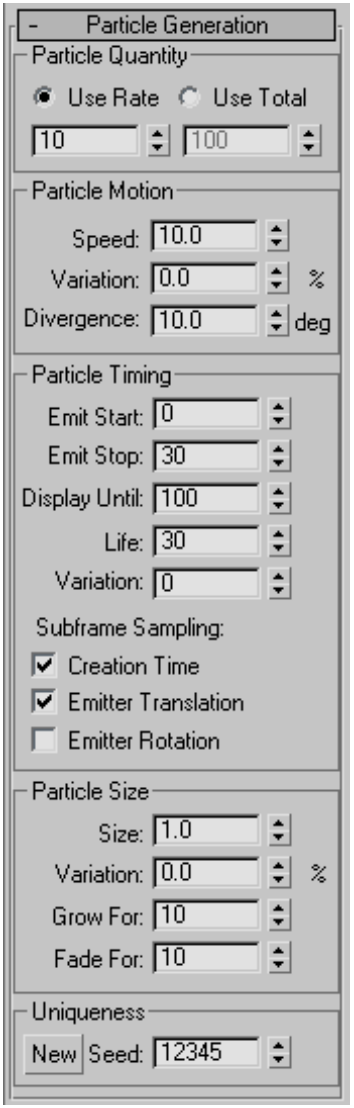
Particle Generation Rollout

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > SuperSpray/Blizzard/Parray/PCloud > Particle Generation rollout

Select a SuperSpray/Blizzard/Parray/PCloud emitter. > Modify panel > Particle Generation rollout

Items on this rollout control when and how quickly particles form, how particles move, and the size of the particles over time.

Interface



Particle Quantity group

In this group, you can choose one of two methods by which the number of particles is determined over time. These settings are unavailable if you set Particle Type (in the *Particle Type rollout (page 2–262)*) to Object Fragments.

Use Rate—Specifies a fixed number of particles emitted per frame. Use the spinner to set the number of particles formed per frame.

Use Total—Specifies a total number of particles formed over the life of the system. Use the spinner to set the number of particles formed per frame.

The life of the system, in frames, is specified by the Life spinner in the Particle Timing group, described later in this topic.

Tip: Generally, Use Rate is best for a continuous flow of particles, like a trail of pixie dust, while Use Total is better for bursts of particles over a short period of time.

Particle Motion group

These spinners control the initial particle velocity, which is directed along the surface, edge, or vertex normals (interpolated for each emitter point).

Speed—The velocity of the particle at birth, along the normal, in units traveled per frame.

Variation—Applies a percentage of variation to the speed of emission for each particle.

Divergence—Applies an angular degree of variation by which each particle's velocity can vary from the emitter normal.

Note: The initial direction for a fragment cluster is the normal of the cluster's seed face. Clusters are created by choosing a single face (the seed face), and then creating a cluster outward from that face, depending on the method chosen in the Object Fragment Controls group on the Particle Type rollout.

Particle Timing group

These options specify when particle emission starts and stops, and the lifespan of the individual particles.

Emit Start—Sets the frame at which particles begin to exist in the scene.

Emit Stop—Sets the last frame at which particles are emitted. This setting has no effect if you choose the Object Fragments particle type.

Display Until—Specifies the frame at which all particles will disappear, regardless of other settings.

Life—Sets the lifespan in number of frames of each particle from the frame of creation.

Variation—Specifies the number of frames by which the life of each particle can vary from the norm.

Subframe Sampling—Turning on any of the three check boxes below helps avoid particle "puffing" by sampling particles at a much higher subframe resolution, instead of the relatively coarse frame resolution. Depending on your needs, you can do this over time, over motion, or over rotation. "Puffing" is the effect of emitting separate "puffs" or clusters of particles, rather than a continuous stream. This effect is especially noticeable when the emitter is animated.

- **Creation Time—**Enables the addition of a time offset to the equations of motion that prevents puffing in time. This setting has no effect with the Object Fragments particle type. Default=on.
- **Emitter Translation—**If the object-based emitter is moving in space, particles are created at integral times at positions along the geometry's path between renderable positions. This prevents puffing in space. This setting has no effect if Object Fragment particle type is on. Default=on.
- **Emitter Rotation—**If the emitter is rotating, turn this on to avoid puffing and produce smooth spiral effects. Default=off.

  Important: Each additional subframe sampling check box that you turn on progressively increases the necessary computation. In addition, the methods are listed in order of least amount of computation

to most. Thus, Emitter Rotation is more costly than Emitter Translation, which is more costly than Creation Time.

Particle Size group

These spinners specify the size of the particles.

Size—This animatable parameter specifies the target size for all particles in the system, depending on the type of particle:

- **Standard Particles—**The major dimension of the particle.
- **Constant—**The dimension, in rendered pixels, of a Constant type of particle.
- **Object Fragments—**No effect.

Variation—The percentage by which the size of each particle may vary from the norm. This is applied to the Size value. Use this parameter to get a realistic mix of large and small particles.

Grow For—The number of frames over which the particle grows from being very small to the Size value. The result is subject to the Size/Variation value, since Grow For is applied after Variation. Use this parameter to simulate natural effects such as bubbles growing as they reach the surface.

Fade For—The number of frames over which the particle will shrink to 1/10th its Size setting prior to its death. This is also applied after Variation. Use this parameter to simulate natural effects such as sparks fading to ash.

Uniqueness group

By changing the Seed value in this spinner, you achieve different results using otherwise identical particle settings.

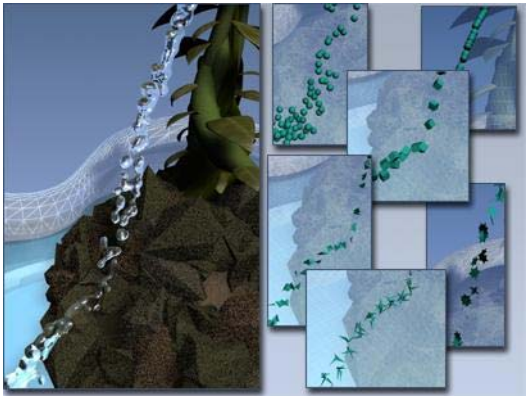
New—Randomly generates a new seed value.

Seed—Sets a specific seed value.

Particle Type Rollout

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > SuperSpray/Blizzard/Parray/PCloud > Particle Type rollout

Select a SuperSpray/Blizzard/Parray/PCloud emitter. > Modify panel > Particle Type rollout

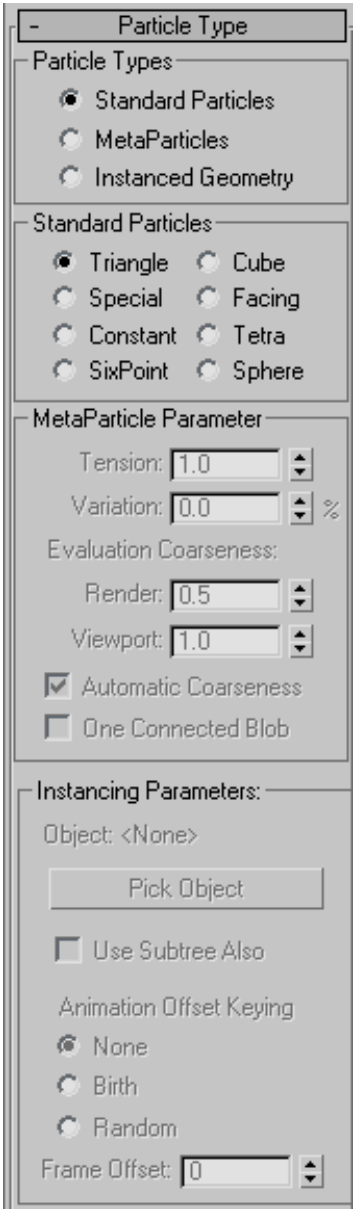


Left: A stream from a fountain

Right: The same stream with various types of particles

The controls on this rollout let you specify the type of particle used and the type of mapping performed on the particles.

Interface



Particle Types group

These options specify one of four categories of particle type. Depending on which option you choose, different controls become available in the lower portion of the Particle Type rollout.

Standard Particles—Uses one of several standard particle types, such as triangle, cube, tetra, and so on.

MetaParticles—Uses Metaball particles. These are particle systems in which the individual particles blend together in blobs or streams.

Object Fragments—Creates particles out of fragments of an object.

Object Fragments is available only with Particle Array. Choose it when you want to fracture the particle emitter object and use the pieces as particles. This option is useful for animating explosions and shattering collisions.

The fragments are created at the Emit Start frame. The Use Rate, Use Total, Emit Stop, and Particle Size parameters are unavailable.

Instanced Geometry—Generates particles that are instances of either an object, a linked hierarchy of objects, or a group. The object is selected in the Instancing Parameters group on the Particle Type rollout.

Choose Instanced Geometry when you want particles to be identical instances of another object in your scene. Instanced geometry particles are extremely effective for creating crowds, flocks, or flows of very detailed objects. Here are a few examples:

- Instance a red blood cell and use Super Spray to animate blood flowing in an artery.
- Instance a bird and use Particle Cloud to animate a flock of birds flying.
- Instance a rock and use Particle Cloud to animate an asteroid field.

  Note: Only one type of particle can be used for the particle system. However, you can have more than one particle array bound to a

single object, and each particle array can emit a different type of particle.

Tip: Image motion blur, described in *Object Properties (page 1–117)* > Motion Blur group, is known not to work properly with instanced particles. Use object motion blur with instanced particles, or use image motion blur with standard particles.

Standard Particles group

When you choose Standard Particles in the Particle Types group, the options in this group become available. Choose one of the following options to specify the particle type:

Triangle—Renders each particle as a triangle. Use Triangle particles with noise opacity for steam or smoke.

Cube—Renders each particle as a cube.

Special—Each particle consists of three intersecting 2D squares. These are effective when you use a face-map material, described in *Shader Basic Parameters Rollout (page 2–1466)*, optionally along with an opacity map, to create the effect of a three-dimensional particle.

Facing—Renders each particle as a square that always faces the view. Use with an appropriate opacity map for bubbles or snowflakes.

Constant—Provides a particle that remains the same size, in pixels, specified in the Size spinner. This size never changes, regardless of its distance from the camera.

Important: You must render either a camera or a perspective view for Constant particles to render correctly.

Tetra—Renders each particle as a mapped tetrahedron. Use Tetra particles for raindrops or sparks.

The default alignment of the tetra particles depends on the particle system type and emitter setup. To specify an alignment, use controls in the *Rotation and Collision rollout (page 2–268)*.

SixPoint—Renders each particle as a six-pointed, two-dimensional star.

Sphere—Renders each particle as a sphere.

MetaParticle Parameters group

When you choose the MetaParticles option in the Particle Types group, the options in this group become available, and metaballs are used as particles. Metaparticles take extra time to render but are very effective for spraying and flowing liquid effects.

Tension—Determines the tightness of the particles, with regard to their tendency to blend with other particles. The higher the Tension, the harder the blobs, and the harder it is for them to merge.

Variation—Specifies the percent of variation of the Tension effect.

Evaluation Coarseness—Specifies how accurately the metaparticle solution is calculated. The higher the coarseness values, the less calculation. However, if the coarseness is too high, there may be little or no metaparticle effect at all. Conversely, if the coarseness is set too low, the time for calculation can become extremely long.

Render—Sets the coarseness for metaparticles in the rendered scene. This option is unavailable when Automatic Coarseness is on.

Viewport—Sets the coarseness for the viewport display. This option is unavailable when Automatic Coarseness is on.

Automatic Coarseness—A general rule is to set the Coarseness value between 1/4 and 1/2 the size of the particles. When this item is on, the rendering coarseness is automatically set, based on the size

of the particles, and the viewport coarseness is set to about twice that of the rendering coarseness.

One Connected Blob—When off (the default), all particles are calculated; when on, a shortcut algorithm is used that calculates and displays only those particles that are adjoining or contiguous to each other.

Note: One Connected Blob mode speeds particle calculations, but you should use it only when your metaparticles form one connected blob, as the label indicates. That is, all particles' blobs must be touching. For example, if you were to use One Connected Blob on a stream of particles containing a mass of 10 contiguous particles followed by a space, then 12 contiguous particles followed by another space, and finally 20 contiguous particles, one of the particles will be chosen, and only the mass connected to that particle will be displayed and rendered.

Tip: When in doubt, leave this option off. If you think all your particles are contiguous and want to save time, turn on One Connected Blob, and then display various frames to see if everything appears.

Object Fragment Controls group

With a particle array, when you choose the Object Fragments particle type, the items in this group become available, and the object-based emitter is exploded into fragments, rather than being used to distribute particles.

Note: To see the fragments in the viewports, choose Mesh in the Viewport Display group near the bottom of the *Basic Parameters rollout (page 2–258)*.

The items in this group include a Thickness spinner, along with three option buttons that determine how the fragments are formed.

Tip: There is no automatic way to hide the distribution object that explodes into fragments. To create the illusion that an object is exploding, you must either set the original object to be invisible at the start of the explosion, as described in *Add Visibility Track (page 2–556)*, or move or scale the original object so it doesn't remain in view.

Thickness—Sets the thickness of the fragments. At 0, the fragments are single-sided with no thickness. When greater than 0, the fragments are extruded, at fragmentation time, by the amount specified. The outer and inner surfaces of the fragment use identical smoothing, which is picked up from the object-based emitter. The edges of the fragments are not smoothed.

The three options that follow specify how the object fragments.

All Faces—Each face of the object becomes a particle. This results in triangular particles.

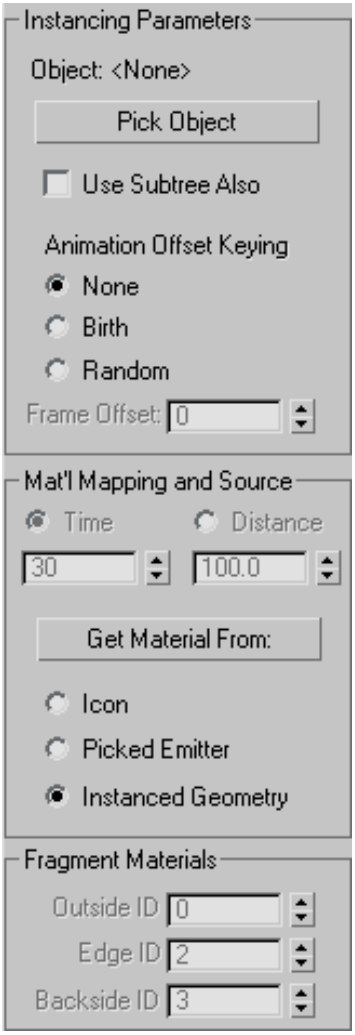
Number of Chunks—The object breaks into irregular fragments. The Minimum spinner, below, specifies the minimum number of fragments that will appear. The method of calculating the chunks may result in more fragments than specified.

Minimum—Determines a number of "seed" faces in the geometry. Each seed face collects connecting faces surrounding it until all available faces are exhausted. Any leftover faces become unique particles, thus increasing the minimum number of fragments.

Smoothing Angle—The fragments are broken based on the angles between face normals, as specified in the Angle spinner. Generally, the higher the Angle value, the fewer the number of fragments.

Angle—Sets the amount of smoothing angle.

Instancing Parameters group



These options are used when you specify Instanced Geometry in the Particle Types group. They let you generate each particle as an instance of either an object, a linked hierarchy of objects, or a group.

Note: Instanced objects can be animated, providing the animation incorporates one or more of the following types:

- Animation of object geometry parameters, such as a sphere's Radius setting.
- Animation of object-space modifiers, such as the Angle setting of a *Bend modifier (page 1–560)*.
- Transform animation of a hierarchical object's children. Transform animation of the top-level parent and non-hierarchical objects is not supported. For example, if you use the toolbar *Select and Rotate function (page 1–439)* to animate a box rotating, and then use the box as instanced geometry with a particle system, the system will not use instanced box's keyframed animation.

Object—Displays the name of the picked object.

Pick Object—Click this, and then select an object in the viewport to be used as a particle. If you select an object that's part of a hierarchy and Use Subtree Also is on, then the picked object and its children become a particle.

If you pick a group, all objects in the group are used as a particle.

Use Subtree Also—Turn this on when you want to include the linked children of the picked object in the particle. If the picked object is a group, all children of the group are included. Note that you can turn this on or off at any time to alter the particles.

Animation Offset Keying—Because the instanced objects can be animated, the options here let you specify the timing of the animations for the particles.

None—Each particle duplicates the timing of the original object. As a result, the animation of all particles will be identically timed.

Birth—The firstborn particle is an instance of the current animation of the source object at the moment of that particle's birth. Each subsequent particle then uses the same start time for the animation. For example, if the animation of the source object is a bend from 0 to 180 degrees,

and the first particle is born at frame 30, when the object is at 45 degrees, then that particle, and all subsequent particles will be born starting at a bend of 45 degrees.

Random—This option is the same as None when Frame Offset is set to 0. Otherwise, each particle is born using the same animation as the source object at the time of birth, but with a random offset of frames, based on the value in the Frame Offset spinner.

Frame Offset—Specifies an offset value from the current timing of the source object.

Mat'l Mapping and Source group

Specifies how a mapped material affects the particles, and lets you specify the source of the material assigned to the particles. A detailed description of how materials affect particles is in the topic *Using Mapped Materials with Particle Systems (page 2–240)*.

Time—Specifies the number of frames from the birth of the particle that it takes to complete one mapping of the particle.

Distance—Specifies the distance, in units, from the birth of the particle that it takes to complete one mapping of the particle.

Note: Tetra particles are an exception. They always have their own local mapping from head to tail, as described in the following section.

Get Material From—Updates the material carried by the particle system, using the source specified by the option buttons below this button.

Important: Remember to click Get Material From whenever you choose a different source option button, or whenever you assign a new material to the specified source. Only a single material (or multi/sub-object material) is carried by the particle-system object at any time. Thus, when you change sources, you actually overwrite the currently assigned material with an instance of the source material.

Icon—The particles use the material currently assigned to the particle system icon.

Note: The Time and Distance options are available only when you choose this option.

Picked Emitter—The particles use the material assigned to the distribution object.

Instanced Geometry—The particles use the material assigned to the instanced geometry. This option is available only when you choose Instanced Geometry in the Particle Types group.

Important: When you turn on either Picked Emitter or Instanced Geometry, an instance of the material from the chosen source is copied to the emitter icon, overwriting the material originally assigned to the icon. Thus, if you've assigned a material to the particle emitter, and then switch to Picked Emitter, the material originally assigned to the icon is replaced by an instance of the material carried by the picked object. If you then return to the Icon option, the particle system does not revert to the material that was assigned the icon, but retains the material taken from the picked object.

Fragment Materials group

These spinners let you assign different material ID numbers to outside surfaces, the edges and the back sides of fragment particles. You can then assign different materials to the front, edge and back of the fragments by using a multi/sub-object material.

Outside ID—Specifies which face ID number is assigned to the outside faces of the fragments. This spinner defaults to 0, which is not a valid ID number, forcing the outside of the particle fragments to use whatever material is currently assigned to the associated faces. Thus, if your distribution object already has several submaterials assigned to its outer faces, these materials are retained by using ID 0. If you want a single,

specific submaterial, you can assign it by changing the Outside ID number.

Edge ID—Specifies which submaterial ID number is assigned to the edges of the fragments.

Backside ID—Specifies which submaterial ID number is assigned to the back sides of the fragments.

Rotation and Collision Rollout

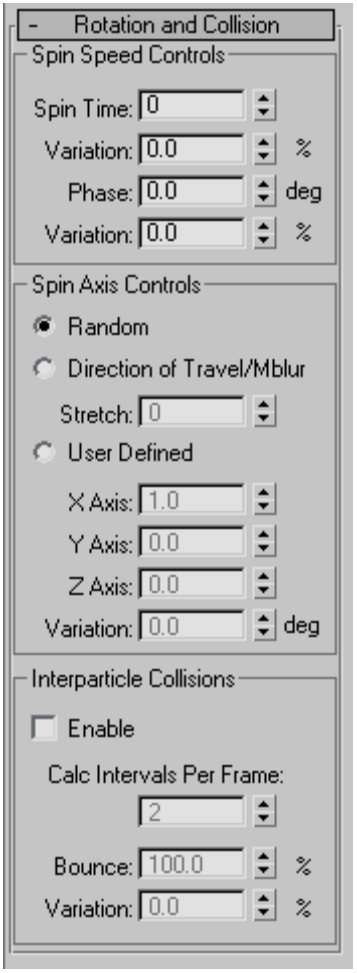
Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > SuperSpray/Blizzard/Parray/PCloud > Rotation and Collision rollout

Select a SuperSpray/Blizzard/Parray/PCloud emitter. > Modify panel > Rotation and Collision rollout

Particles often move at high rates of speed. In such cases, you might want to add motion blur to the particles to enhance their motion. Also, real-world particles typically rotate as they move, and collide with each other.

The options on this rollout affect the rotation of the particles, provide motion blur effects, and control inter-particle collisions.

Interface



Spin Speed Controls group

Spin Time—The number of frames for one rotation of a particle. If set to 0, no rotation takes place.

Variation—The percent of variation of the Spin Time.

Phase—Sets the initial particle rotation, in degrees. This has no meaning for fragments, which always begin with zero rotation.

Variation—The percent of variation of the Phase.

Spin Axis Controls group

These options determine the spin axis for the particles, and provide a partial method of applying motion blur to the particles.

Random—The spin axis for each of the particles is random.

Direction of Travel/Mblur—(Not available with the Blizzard particle system.) Rotates the particles about a vector formed by the direction in which they're moving. This option also lets you apply a form of motion blur to the particles by using the Stretch spinner. For further information, see *Achieving Particle Motion Blur (page 2–240)*.

Stretch—When greater than 0, the particles stretch along the travel axis, depending on their speed. Specifically, the Stretch value specifies the percent of their length per each unit of the Speed setting (in the Particle Motion group). Thus, if you set Stretch to 2 while Speed is set at 10, the particles are stretched 20 percent longer than their original size along the axis of their travel. This spinner is available only when you choose Direction of Travel/Mblur.

Note: For best results when using Stretch, you should also assign the Particle MBlur map as an opacity map of the material assigned to the particle system.

User Defined—Uses a vector defined in the three X, Y, and Z axis spinners.

X/Y/Z Axis—Specifies the spin vector of the X, Y, or Z axis, respectively. These spinners are available only when User Defined is chosen.

Variation—The amount, in degrees, by which the spin axis of each particle may vary from the specified X Axis, Y Axis, and Z Axis settings. This spinner is available only when you choose User Defined.

Interparticle Collisions group

These options enable collisions between particles, and control how the collisions occur. Note that this involves intensive calculation, particularly when large numbers of particles are involved.

Enable—Enables inter-particle collisions in the calculation of the particle movements.

Calc Intervals Per Frame—The number of intervals per rendering interval, during which an inter-particle collision test is conducted. The higher the value, the more accurate the simulation, but the slower the simulation will run.

Bounce—The degree to which speed is recovered after a collision.

Variation—The percentage of random variation of the Bounce value, applied to the particles.

Object Motion Inheritance Rollout

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > SuperSpray/Blizzard/Parray/PCloud > Object Motion Inheritance rollout

Select a SuperSpray/Blizzard/Parray/PCloud emitter. > Modify panel > Object Motion Inheritance rollout

Each particle's position and direction of movement are determined by the position and orientation of the emitter at the time the particle is created. If the emitter is moving through the scene, particles are scattered along the emitter's path.

Use these options to affect the motion of the particles by the motion of the emitter.

Procedure

Example: To observe object motion inheritance:

1. Create a Super Spray particle system.
2. Animate the emitter moving sideways (perpendicular to the direction of the particle stream) between frames 1 and 15.

3. Play the animation while observing it from the Top viewport.

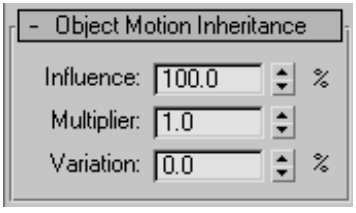
   The emitter stops at frame 15, while the particles it has emitted up to that point continue moving along the diagonal between the emitter's path and that of the particles. The remaining particles move straight out from the emitter.

4. On the Object Motion Inheritance rollout, set Influence to 50.

5. Play the animation again.

   This time, only some of the particles inherit the emitter's motion, while the rest move straight out from the emitter. Note that the latter form a diagonal stream because each particle emerges at a subsequent point on the emitter path.

Interface



Influence—The percent of particles that inherit the motion of the object-based emitter at the moment of particle formation. For example, when this is set at 100 (the default), all particles travel along with a moving object; when it's set at 0, none of the particles are affected by the translation of the object and fall behind its movement.

Multiplier—Modifies the amount by which the emitter motion affects the particle motion. This can be a positive or negative number.

Variation—Provides a percentage of variation of the Multiplier value.

Bubble Motion Rollout

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > SuperSpray/Parray/PCloud > Bubble Motion rollout

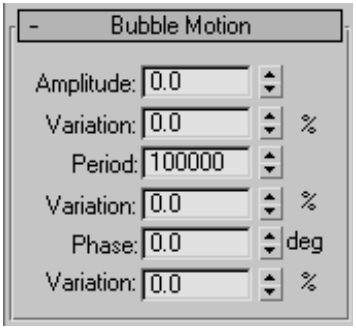
Select a SuperSpray/Parray/PCloud emitter. > Modify panel > Bubble Motion rollout

Bubble motion provides the wobbling effect you see in bubbles rising underwater. Typically, it's used when the particles are set to rise in thin streams. Bubble motion is similar to a waveform, and the Bubble Motion parameters let you adjust the amplitude, period, and phase of the bubble "wave."

Note: The bubble motion is not affected by space warps, so you can use a space warp to control the direction of the particle flow without altering the local, wobbling bubble effect.

Tip: InterParticle Collisions, Deflector Binding, and Bubble Noise do not get along well together. Particles may leak through the deflector when these three are used together. Instead of bubble motion, use animated mapping. Use facing particles with an animated map of a bubble, where the bubble is smaller than the map size. The bubble is animated moving around the map. This simulates bubble motion at the map level.

Interface



Amplitude—The distance the particle moves off its usual velocity vector as it travels.

Variation—The percent of Amplitude variation applied to each particle.

Period—The cycle time for one complete oscillation of a particle through the bubble "wave." A recommended value might be 20 to 30 intervals.

Note: Bubble motion is measured in time, not in rate, so a very large Period value means the motion takes a long time to complete. Thus, there is no motion, effectively. Period is therefore set to a very large default value so that the default motion of this type is none.

Variation—The percent of Period variation for each particle.

Phase—The initial displacement of the bubble pattern along the vector.

Variation—The percent of Phase variation for each particle.

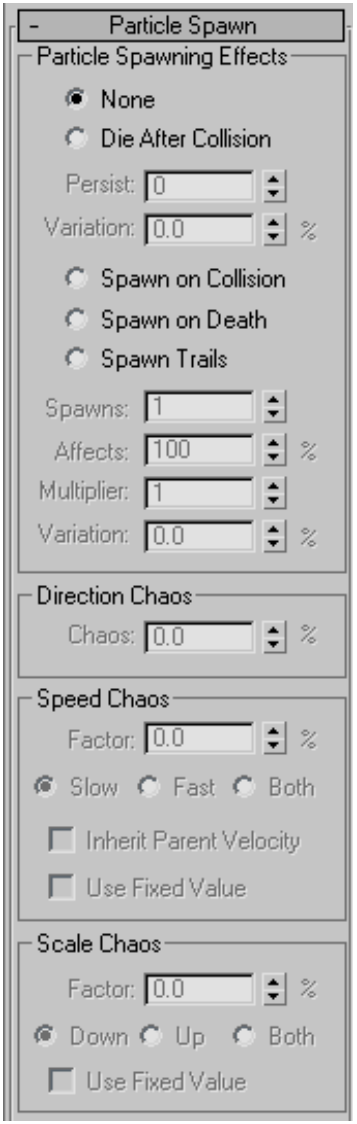
Particle Spawn Rollout

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > SuperSpray/Blizzard/Parray/PCloud > Particle Spawn rollout

Select a SuperSpray/Blizzard/Parray/PCloud emitter. > Modify panel > Particle Spawn rollout

The options on the Particle Spawn rollout let you specify what happens to particles when they die, or when they collide with a particle deflector. Using the options on this rollout, you can have particles generate additional generations of particles upon collision or death.

Interface



Particle Spawning Effects group

Choose one of these options to determine what happens to the particles at either collision or death.

None—Uses none of the spawning controls, and the particles act as they normally would. That is, upon collision, they either bounce or stick,

depending on Particle Bounce settings in the deflector, and on death they disappear.

Die After Collision—Particles disappear when they strike a deflector to which they're bound, such as the SDeflector.

Persist—The life, in frames, that the particle will persist after the collision. Setting this to 0 (the default) causes particles to vanish immediately after the collision.

Variation—Varies the Persist value of each particle, when Persist is greater than 0. This lets you "feather" the dying off of particle density.

Spawn on Collision—Spawn effects take place upon collision with a bound deflector.

Spawn on Death—Spawn effects take place at the end of each particle's life.

Spawn Trails—Particles are spawned from existing particles at each frame of that particle's life. The Multiplier spinner specifies the number of particles that are spawned from each particle. The base direction of the spawned particles is opposite that of the parent's velocity. The Scale, Direction, and Speed Chaos factors are applied to that base direction.

Note: If the Multiplier is greater than 1, at least one of the three Chaos factors must be greater than 0 in order to see the additional spawned particles. Otherwise, the multiples will occupy the same space.

Warning: This option can produce *many* particles. For best results, begin by setting Particle Quantity on the Particle Generation rollout to Use Rate and to 1.

Spawns—The number of spawns beyond the original particle generation. For example, if this is set to 1, and you're spawning at death, one spawning will occur beyond the original lifespan of each particle.

Affects—Specifies the percentage of particles that will spawn. Reducing this reduces the number of particles that produce spawned particles.

Multiplier—Multiplies the number of particles spawned at each spawning event.

Variation—Specifies a percentage range by which the Multiplier value will vary, frame by frame.

Direction Chaos group

Chaos—Specifies the amount by which the direction of the spawned particle can vary from the direction of the parent particle. A setting of 0 means no variance. A setting of 100 causes the spawned particle to travel in any random direction. A setting of 50 causes the spawned particle to deviate from its parent's path by up to 90 degrees.

Speed Chaos group

These options let you vary the spawned particles' speed randomly in relation to their parents' speed.

Factor—This is the range of a percentage of change in the speed of the spawned particle relative to its parent. A value of 0 means no change.

Slow—Applies the speed factor randomly to slow the speed of the spawned particles.

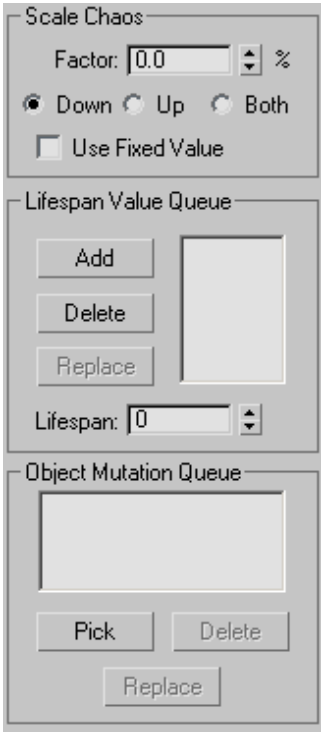
Fast—Randomly speeds up particles based on the speed factor.

Both—Some particles speed up, while others slow down, based on the speed factor.

Inherit Parent Velocity—Spawned particles inherit the speed of their parents, in addition to the effect of the speed factor.

Use Fixed Value—Uses the Factor value as a set value, rather than as a range applied randomly to each particle.

Scale Chaos group



These options apply random scaling to the particles.

Factor—Determines a random percentage range of scaling of the spawned particles relative to their parents, and dependent on the options below.

Down—Randomly scales down spawned particles to be smaller than their parents, based on the Factor value.

Up—Randomly scales up spawned particles to be larger than their parents.

Both—Scales spawned particles both larger and smaller than their parents.

Use Fixed Value—Uses the Factor value as a fixed value, rather than a range of values.

Lifespan Value Queue group

These options let you specify a list of alternative lifespan values for each spawned generation of particles. The spawned particles use these lifespans rather than the lifespan specified for the original particles in the Life spinner on the *Particle Generation rollout (page 2–260)*.

List window—Displays a list of lifespan values. The first value on the list is used for the first generation of spawned particles, the next value is used for the next generation, and so on. If there are fewer values on the list than there are spawnings, then the last value is used repeatedly for all remaining spawnings.

Add—Adds the value in the Lifespan spinner to the list window.

Delete—Deletes the currently highlighted value in the list window.

Replace—Lets you replace a value in the queue with a value in the Lifespan spinner. To use, first place a new value in the Lifespan spinner, then select the value in the queue you want to replace, and click the Replace button.

Lifespan—Use this to set a value, and then click the Add button to add the value to the list window.

Object Mutation Queue group

The options in this group let you switch between instanced-object particles with each spawning (as set with the Spawns spinner). These options are available only if Instanced Geometry is the current particle type.

List window—Displays a list of objects to be instanced as particles. The first object on the list is used for the first spawning, the second for the second spawning, and so on. If there are fewer objects on the list than there are spawnings, then the last object on the list is used for all remaining spawnings.

Pick—Click this, and then select an object in the viewport to add to the list. Note that the type of object you use is based on the settings in the Instancing Parameters group of the Particle Type rollout. For example, if you've turned on Subtree in that group, you can pick object hierarchies. Likewise, if you've picked a group, you can use groups as your spawned particles.

Delete—Deletes the currently highlighted object in the list window.

Replace—Replaces an object in the queue with a different object. Select an object in the queue to enable the Replace button. Click Replace, and then pick an object in the scene to replace the highlighted item in the queue.

Load/Save Presets Rollout

Create panel > Geometry button > Choose Particle Systems from the drop-down list. > Object Type rollout > SuperSpray/Blizzard/Parray/PCloud > Load/Save Presets rollout

Select a SuperSpray/Blizzard/Parray/PCloud emitter. > Modify panel > Load/Save Presets rollout

These options let you store preset values that can be used in other, related particle systems. For example, after setting parameters for a PArray and saving it under a specific name, you can then select another PArray system, and load the preset values into the new system.

Interface



Preset Name—An editable field in which you can define a name for your settings. Click the Save button to save the preset name.

Saved Presets—Contains all the saved preset names. A number of presets are included with 3ds Max; to see what they do, create a particle system, load a preset, and play back the animation. Some of the presets, such as Particle Array's Shimmer Trail, are most effective with moving particle systems.

Load—Loads the preset currently highlighted in the Saved Presets list. Alternatively, double-click the preset name in the list to load it.

Save—Saves the current name in the Preset Name field and places it in the Saved Presets window.

Delete—Deletes the selected items in the Save Presets window.

Note: Animated parameter values subsequent to frame 0 are not stored.

12 Animation

3ds Max gives you a variety of different ways to create animation, and a wealth of tools for managing and editing animation. See the following links for documentation of some of the animation methods in 3ds Max:

- *Animation Concepts and Methods (page 2–275)* provides information about basic animation with 3ds Max.
- Animation in 3ds Max is accomplished by means of controllers; *Working with Controllers (page 2–289)* gives you an overview of how they work. For specifics about controllers, see *Animation Controllers (page 2–307)*.
- A special type of animation controller is the constraint; for details, see *Animation Constraints (page 2–392)*.
- A powerful method for creating sophisticated, procedural animation is covered in *Wire Parameters (page 2–411)*.
- For animating characters both humanoid and other, as well as complex mechanical assemblies, it's important to know about *Hierarchies and Kinematics (page 2–416)*.
- *Track View (page 2–501)* gives you access to powerful tools for analyzing and editing animation.
- *Animation Utilities (page 2–653)* can help ease tedious animation tasks.
- When animating a number of different objects the same way, or transferring animation from one object to another, it's useful to know the information covered in *Saving and Loading Animation (page 3–472)*.
- The *Motion Mixer (page 2–604)* is a powerful tool for combining individual movements into complex animations.
- Another way of combining animation segments is the *Animation Layers (page 2–326)* system.

Animation Concepts and Methods

With 3ds Max, you can create 3D computer animation for a variety of applications. You can animate characters or vehicles for computer games, or you can animate special effects for film or broadcast . You can create animation for serious purposes such as medical illustration or forensic presentation in the courtroom. Whatever reason you have to animate, you'll find 3ds Max a capable environment for achieving your goals.

The basic way to animate is quite simple. You animate the transform parameters of any object to change its position, rotation, and scale over time. Turning on the *Auto Key (page 3–717)* button, then moving the time slider places you in a state in which any changes you make will create animation for the selected objects in the viewport.

Animation is used throughout 3ds Max. You can animate the position, rotation, and scale of an object, and almost any parameter setting that affects an object's shape and surface. You can link objects for hierarchical animation, using both forward and inverse kinematics, and you can edit your animation in *Track View (page 2–501)*.

This section discusses the basics of creating animation. It looks briefly at a comparison between computer animation and classic hand-drawn animation, and then describes the creation of keyframed animation. The following topics are covered here:

Animation Concepts (page 2–276)

Using Auto Key Mode (page 2–278)

Using Set Key Mode (page 2–280)

Viewing and Copying Transform Keys (page 2–283)

Controlling Time (page 2–285)

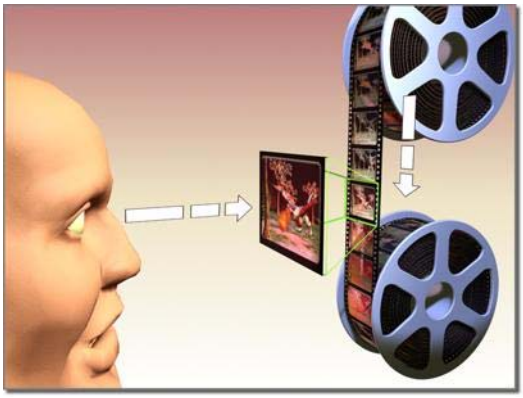
Setting Time Segments (page 2–286)

Moving Through Time (page 2–287)

Choosing a Frame Rate and Playback Speed (page 2–288)

Animation Concepts

Animation is based on a principle of human vision. If you view a series of related still images in quick succession, you perceive them as continuous motion. Each individual image is referred to as a *frame*.



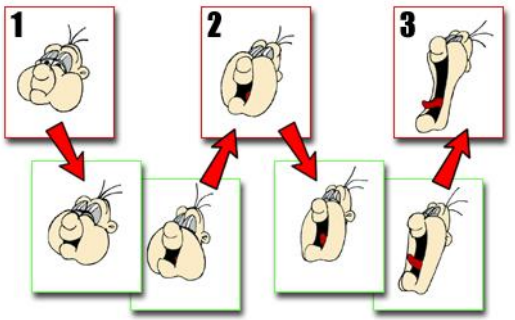
A frame is a single image in an animated film.

Traditional Animation Method

Historically, the main difficulty in creating animation has been the effort required of the animator to produce a large number of frames. One minute of animation might require between 720 and 1800 separate images, depending on the quality of the animation. Creating images by hand is a big job. That's where the technique of *keyframing* comes in.

Most of the frames in an animation are routine, incremental changes from the previous frame directed toward some goal. Traditional animation studios realized they could increase the productivity of their master artists by having them draw only the important frames, called *keyframes*. Assistants could then figure out what belonged on the frames in between the keyframes. The in-between frames were called *tweens*.

Once all of the keyframes and tweens were drawn, the images had to be inked or rendered to produce the final images. Even today, production of a traditional animation usually requires hundreds of artists to generate the thousands of images needed.



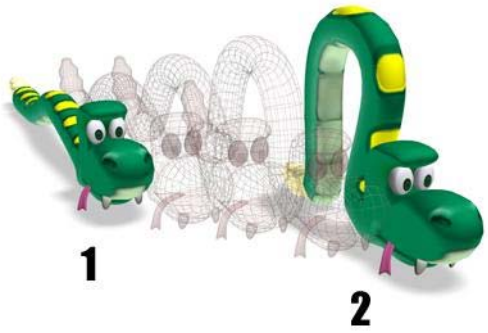
The frames marked 1,2, and 3 are key frames. The other frames are tweens.

The 3ds Max Method

This program is your animation assistant. As the master animator, you create the keyframes that record the beginning and end of each animated sequence. The values at these keyframes are called *keys*. The software calculates the interpolated values between each key to produce the completed animation.

3ds Max can animate just about any parameter in your scene. You can animate modifier parameters, such as a Bend angle or a Taper amount, material parameters, such as the color or transparency of an object, and much more.

Once you have specified your animation parameters, the renderer takes over the job of shading and rendering each frame. The result is a high-quality animation.

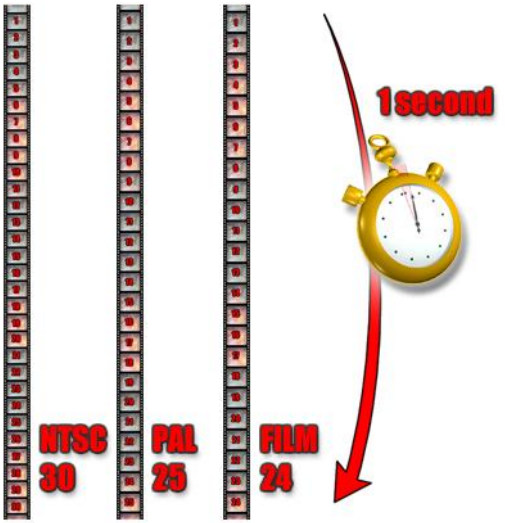


The object's position at 1 and 2 are the keyframed models at different times. The computer generates the in-between frames.

Comparing Frames and Time

Traditional animation methods, and early computer animation programs, are rigidly locked to the concept of producing animation frame by frame. This is okay if you always work in a single format or do not need to specify an animated effect at a precise time.

Unfortunately, animation comes in many formats. Two of the more common formats are film at 24 frames per second (FPS) and *NTSC video (page 3–980)* at 30 FPS. Also, the need for accurate time-based animation versus frame-based animation is critical as animation becomes more common for scientific and legal presentation.



Different animation formats have differing frame rates.

3ds Max is a time-based animation program. It measures time, and stores your animation values, with an internal precision of 1/4800 of a second. You can configure the program to display time in a format best suited for your work, including traditional frames format. See *Time Configuration (page 3–725)*.

Many of the examples in the following sections describe time using the frames method for the sake of tradition and familiarity. Keep in mind that you are really animating using a very precise time-based method, and frames aren't created until you instruct the software to render your animation.

Identifying Animation Tools

You can locate the basic animation tools in the following areas of the interface:

Track View—Provides detailed animation editing capabilities in several floating or dockable windows. See *Track View (page 2–501)*.

Track Bar—Provides quick access to keyframes and interpolation controls. Can be expanded for function curve editing. See *Track Bar (page 3–703)*.

Motion Panel—Use this panel to adjust transform controllers that affect all position, rotation, and scale animation. See *Working with Controllers (page 2–289)*.

Hierarchy Panel—Use this panel to adjust all parameters governing the linkage of two or more objects. These include inverse kinematics parameters and pivot point adjustments. See *Hierarchy Panel (page 3–773)*.

Time Controls—Use these controls to move through time in the viewport displays. You can move to any point in time, and play animations in the viewports. The time controls, documented in *Animation and Time Controls (page 3–716)*, include:

- *Time slider (page 3–701)*
- *Auto Key button (page 3–717)*
- *Set Key buttons (page 3–718)*
- *Animation playback buttons (page 3–723)*
- *Current frame field (page 3–724)*
- *Key Mode toggle (page 3–724)*
- *Time Configuration button (page 3–725)*

To get started with animation, see *Using Auto Key Mode (page 2–278)*.

Using Auto Key Mode



Start creating an animation by turning on the Auto Key button, setting a current time, and then changing something in your scene. You can change the position, rotation, or scale of an object, or change almost any setting or parameter.

When you make a change, a key storing the new value for the changed parameter is created at the current time. If that key was the first animation key created for the parameter, a second animation

key is also created at time 0 to hold the parameter's original value.

Keys are not created at time 0 until you create at least one key at another time. After that, you can move, delete, and re-create keys at time 0.

Turning Auto Key on has the following effect:

- The Auto Key button, the time slider and the border of the active viewport turn red to indicate you are in animation mode.
- Keys are created whenever you transform an object or change an animatable parameter.
- The time slider sets the time where keys are created.

To begin animating an object:

1. 

Click Auto Key to turn it on.
2. Drag the time slider to a time other than 0.
3. Do one of the following:
   - Transform an object.
   - Change an animatable parameter.

For example, if you have a cylinder that has not been animated yet, it has no keys. If you turn on Auto Key, and at frame 20 you rotate the cylinder 90 degrees about its Y axis, rotation keys are created at frames 0 and 20. The key at frame 0 stores the original orientation of the cylinder, while the key at frame 20 stores the animated orientation of 90 degrees. When you play the animation, the cylinder rotates 90 degrees about its Y axis over 20 frames.

Modeling Without Animating

Just as you can animate at any time by turning Auto Key on, you can also model at any time in your animation without creating animation keys.

The results of changing an object or any other parameters with Auto Key off varies according to whether or not the object or parameters have been animated yet.

- If you create a new object, or change an object parameter that has not been animated yet, you can work at any time with Auto Key off. The changes you make are constant through the entire animation.

  For example, you might animate an object bouncing around your scene and then decide to create pads for the object to land on. To do that, you drag the time slider to a time when the bouncing object hits the ground, and make sure Auto Key is off before you proceed. You can then create a pad under the bouncing object and repeat the process at the next time where it hits the ground. Because Auto Key was off, it does not matter at what time the pad objects were created. They remain inanimate through the entire animation.
- If you change an object or parameter that is already animated, while Auto Key is off, the amount of change is applied equally across all the animation keys.

  For example, you might animate a sphere's radius to be 15 at frame 0, 30 at frame 10, and 50 at frame 20. If you drag the time slider to frame 10, turn Auto Key off, and increase the sphere's radius from 30 to 40, the change in the radius is applied to the other two keys as well. Because you increased the radius by 10 units with Auto Key off, all radius keys are increased by 10 units. The sphere's radius is now 25 at frame 0, 40 at frame 10, and 60 at frame 20.

  If Auto Key had been on when you changed the radius, it would have been an animated change applied only to the key at frame 10.

Identifying What Can Be Animated

Because most parameters in 3ds Max can be animated, the easiest way to find out if something

can be animated is just to try it. Usually, if you want to animate a parameter, it can be animated.

Sometimes you need to know in advance if you can animate a parameter. If so, you can use Track View. The Track View Hierarchy list displays every parameter that can be animated. You might also need to add a controller to a track before it can be animated. See *Track View (page 2–501)* and *Animation Controllers (page 2–307)*.

Using Set Key Mode 

The Set Key animation system is designed for the professional character animator who wants be able to try out poses and then commit those poses to keyframes intentionally. It can also be used by any animator to set keys on specific tracks of objects. This way of working allows more control than the Auto Key method, because it gives you the chance to try out ideas and discard them quickly without having to undo work. It lets you pose your character, and selectively key certain tracks on certain objects through the use of Key Filters and Keyable tracks in Track View.

Note: Although the following information uses the example of animating a character, it also applies to animating a complex mechanical assembly.

Straight-Ahead and Pose-to-Pose

Traditional animation is created one of two ways, either straight-ahead animation or pose-to-pose animation. Straight-ahead animation is drawn starting from the beginning and then additional frames are drawn sequentially thereafter, moving straight ahead in time. Pose-to-pose animation is created by drawing the important frames first, (extremes and breakdowns), and then the intervening frames are filled in later.

Once a character has been correctly drawn for a specific frame, pose-to-pose animation requires that all the keyable tracks needs to be keyframed. This creates a pose of the character that will not be affected if animation for the character is edited at other points in time. If all the animatable tracks are keyed in the extremes, the in-betweening work will not destroy any of those poses.

Set Key and Pose-to-Pose Animation

The number of objects and tracks that require keying for a character, even a simple character, is not something that can be easily handled manually. Set Key makes this process easier by listing all the tracks that are parts of the character that have to be keyed when you want to fix a pose and create a snapshot in time. Keyable tracks let you determine which tracks can be keyed, then Key filters let you work on them selectively, placing keys on only the tracks you want.

Differences Between Set Key and Auto Key Modes

Set Key mode differs from Auto Key in a number of ways.



In Auto Key mode, the workflow is to turn on Auto Key, move to a point in time, then transform objects or change their parameters. All changes register as keyframes. When you turn off Auto Key mode, you no longer are creating keys. Changes made to objects when Auto Key mode is off are applied globally to the animation. This is referred to as *Layout mode*.



In Set Key mode, the workflow is similar, but the behavior is fundamentally different. Turn on Set Key mode, then move to a point in time. Before you transform or change object parameters, you determine the tracks you want to set keys on using Keyable icons in Track View and Filters. Once you know what you plan to key, you try

out poses in the viewport (transform the objects, change parameters, and so forth).

 When you like what you see, click the large Set Keys button or press K on the keyboard to set a key. If you don't do this, no keys are set.

If you move to another point in time, your changes are lost and have no effect on your animation. For example, if you find that you have a posed character, but *at the wrong frame* in time, you can hold down Shift and the right-mouse button and drag the time slider to the correct frame without losing your pose.

Using Set Key with Inverse Kinematics

Choosing IK Parameters in Key Filters allows you to use Set Key to keyframe inverse kinematics. This lets you set keys for IK goals and end effectors using Set Key as well as other IK parameters such as Swivel Angle or Twist.

As always, when using Set Key, you can selectively keyframe tracks by combining Keyable icons in Track View with Key Filters.

Set Key doesn't currently support IK/FK Enabling, so don't try to keyframe the Enable button using the Set Key button or the keyboard shortcut. Use the Auto Key method when you want to work with IK/FK blending.

Using Set Key with Materials

If you select Materials in Key Filters, you can use Set Key to create keys for materials. Be forewarned that you need to use Keyable Icons to limit the tracks which get keyed. If you simply turn on Materials and set a key, you will place keys on every Material track, something you probably don't want to have happen.

Using Set Key with Modifiers and Object Parameters

When you want to set a key on an object's parameters, and you have the Object Parameters Key Filter selected, every parameter will receive a key, unless you have turned off the parameter track in the Controller window of Track View using Keyable icons. It might be easier to simply Shift+right-click the parameter spinner to set the key.

Also make sure both Modifiers and Object Parameters are turned on in the *Filters dialog (page 2–542)* when you are keyframing a modifier gizmo.

Additional Set Key Tools

Additional Set Key Tools are found in the *Customize User Interface dialog (page 3–792)*. On the Keyboard panel, choose Set Key Tools from the Category field. Here, you can set keyboard shortcuts to clear the set key buffer, as well as create keys on just a single axis for a transform.

Using Set Key with Sub-Object Animation

When using Set Key with sub-object animation, you must first assign a controller before creating a key. Sub-objects do not have a default controller assigned upon creation. The controller is assigned by animating at the sub-object level.

Other Methods to Set Keys

You can also set position, rotation, and scale keys by right-clicking the Frame Indicator of the time slider. To set keys on parameters that have spinners, hold down the Shift key and right-click to set a key using the existing parameter value.

Procedures

To use Set Key animation:

1. 

Turn on the Set Key button.



When the button is red you are now in Set Key mode.

This is a mode where you can try out ideas before you commit to them.

2. Open Track View (either Curve Editor or Dope Sheet).
3.  Click the Show Keyable button on the Track View toolbar.
4. Turn off all the other tracks you don't want to keyframe.

   The red key means the track will be keyed. If you click the red key, it turns to a gray key, which means that track will not be keyed.

   Tip: You can toggle multiple tracks as keyable by using the Controller menu > Keyable command.

   When you are finished, minimize or close Track View.
5. 

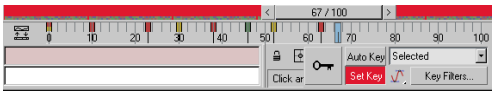
Click the Key Filters button, and turn on the Filters to choose the tracks you want keyframed.

   Position, Rotation, and Scale are on by default.

   You can use the Key Filters button to work on individual tracks selectively. For example, if you are in Track View and the Rotation and Position tracks of a character's arm are keyable, you can use the key filters to turn off the Position filter and only work on the Rotation tracks.
6. Move the time slider to another point in time, transform your objects or adjust parameters in the command panels to create animation.

   This does not yet create keyframes.
7.  Click the Set Keys button or press K on the keyboard to set a key.

 When the button turns red, it sets a key which appears on the time ruler. The keys are color coded to reflect which tracks are being keyed.



If you don't click Set Keys and you move to another point in time, the pose is lost.

Tip: To move the pose to another point in time, use the right mouse button to press and drag the time slider. This lets you move to another frame number without losing the pose.

To animate a vertex using set key:

1. Create an editable spline.
2. Select a vertex.
3. Turn on Set Key mode.
4. Move the selected vertex.
5. Click the Set Keys button.

   Now a controller has been assigned to the vertex. From this point forward, you can animate.

To set a key on every keyable track of an object:

1. Select the object for which you need to set keys.
2. On the Key Filters dialog, click All.
3. Click the Set Keys button or press K on the keyboard.

Spinner Right-Click Menu

To open the spinner right-click menu, right-click the editable field of an animatable parameter. This menu provides options for cutting, copying, and pasting values, wires, and animation tracks

between various object parameters. It also allows you to show the parameter in *Track View (page 2–501)* or in the *Parameter Wiring dialog (page 2–412)*.

Interface

With the exceptions of Undo and Select All, each of the following commands takes effect on the parameter you right-click. It's not necessary to first click in the value field.



Undo—Reverses the effect of the last action.

Cut—Removes highlighted text to be pasted elsewhere.

Copy—Copies highlighted text to be pasted elsewhere.

Paste—Inserts cut or copied text into the editable field.

Delete—Removes highlighted text.

Select All—Highlights all text within the active editable field.

Note: In order for Select All to work, the text cursor must be active in the field that you right-click to open the menu. For best results, first click in the field, and then right-click the field and choose Select All.

Copy Animation—Copies the animation controller, including all animation keys, to a buffer, which you can then assign elsewhere by pasting.

Paste Animation - Copy—Assigns a copy of animation controller in the buffer to the right-clicked field.

Paste Animation - Instance—Assigns an *instance (page 3–957)* of the copied animation controller to the right-clicked field.

Paste Animation - Wire—Opens a *Parameter Wiring dialog (page 2–412)* with the copied parameter and the paste destination parameter selected. This is equivalent to creating a *wire parameter (page 2–411)* in the viewport between two parameters.

Edit Wire—Opens the *Parameter Wiring dialog (page 2–412)* expanded and positioned to show the wire controller assigned to the selected parameter.

Note: Edit Wire is available only if a *two-way connection* is currently applied on the right-clicked parameter.

Show in Track View—Opens a Track View - Curve Editor window, titled "Selected," expanded and positioned to show the selected parameter. This is very useful for fast access to a particular curve or track.

Show in Parameter Wire dialog—Opens a Parameter Wiring dialog expanded and positioned to show the selected parameter.

Viewing and Copying Transform Keys

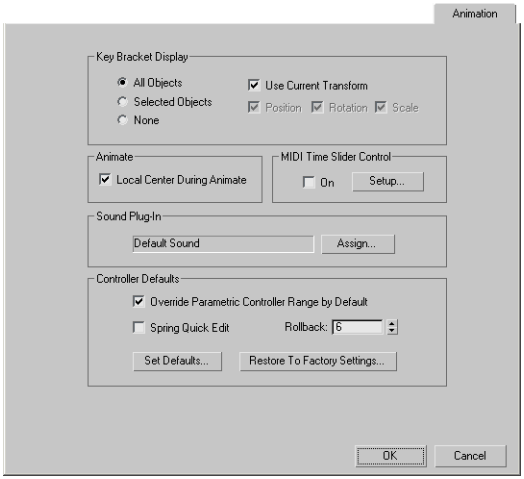
The viewports display white brackets around objects that have transform keys at the current time. These key brackets only appear in viewports using the wireframe shading method.

Use the Track View to view all key types. You can also see all keys for the current selection in the *track bar (page 3–703)*.

For example, suppose you animate a sphere by moving it at frame 20, and scale and rotate it at frame 50. When you drag the time slider, white brackets appear around the sphere at frames 50, 20, and 0, and keys appear at the same frames in the track bar.

If you then apply a modifier such as Bend, and animate its Angle setting at frame 40, you won't see a white bracket around the sphere at frame 40, but track bar displays a key for the Bend animation.

Controlling Key Bracket Display

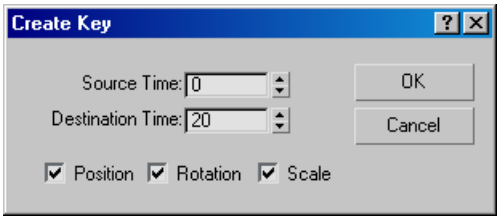


Animation panel of the Preference Settings dialog

You can control the display of key brackets using options in the Preference Settings dialog > *Animation panel (page 3–828)*.

Creating Transform Keys with the Time Slider

You can use the *time slider (page 3–701)* to create transform keys by copying transform values from one time to another. To specify the type of key to create and the source and destination time for the key values, right-click the time slider to display the Create Key dialog.



You set parameters in the Create Key dialog:

Source Time—Specifies the time from which transform values will be copied.

Destination Time—Specifies the time where the key will be created.

Position, Rotation, Scale—Determine which transform key values will be copied to the destination time.

When you click OK, new keys for the specified transforms are created at the destination time, using values from the source time. Keys do not have to exist at the source frame, because the interpolated values at the frame are used.

You might find it easier to create and manipulate keys with the *track bar (page 3–703)*.

When Auto Key mode is on, you can right-click and drag the time slider at the same time. When you do this, the Source time uses the frame number that you were on when you pressed the mouse button, and the Destination time accepts the frame number that you move the time slider to.

When Set Key mode is on, you can right-click and drag the time slider to move to another frame in time, without losing your character pose. If you find you have posed your character on the wrong frame, simply right-click and drag the time slider, and the pose will be copied to the new frame. Click Set keys to set keys for the pose on the new frame.

Creating Position Lock Keys and Rotation Lock Keys

Creating a lock key creates a key with Linear interpolation. If you create the lock key while an existing key is selected, it changes that key's interpolation from Smooth to Linear. (Different types of interpolation are described in *Bezier Controllers (page 2–310)*.)

You can create a lock key for position or for rotation.

Lock keys are useful when you want an object to be stationary, but smooth interpolation is causing it to "wobble" on its stationary spot.

To create a lock key:

1. From the Customize menu, choose Customize User Interface.
2. On the Customize User Interface dialog, click the Quads tab, and then, from the drop-down list at the upper right of the dialog, choose Animation.
3. Click the upper left quadrant of the four gray squares that define the quad menus.

   The quad turns yellow.
4. In the Action list to the left, find Create Position Lock Key. Drag it to the window on the right below any menu item. Choose Save and click OK to apply this and close the dialog.
5. (Optional.) Select a key.
6. Alt+right-click the object you're animating.

   Note: The Auto Key button need not be on.
7. From the quad menu that appears, choose Create Position Lock Key .

   You can also create keyboard shortcuts for these two commands. Create Position Lock Key and Create Rotation Lock Key are main user interface shortcuts. See *Keyboard Shortcuts (page 3–871)*.

Controlling Time

You create animation by changing your scene over time. You can exercise great control over time, as follows:

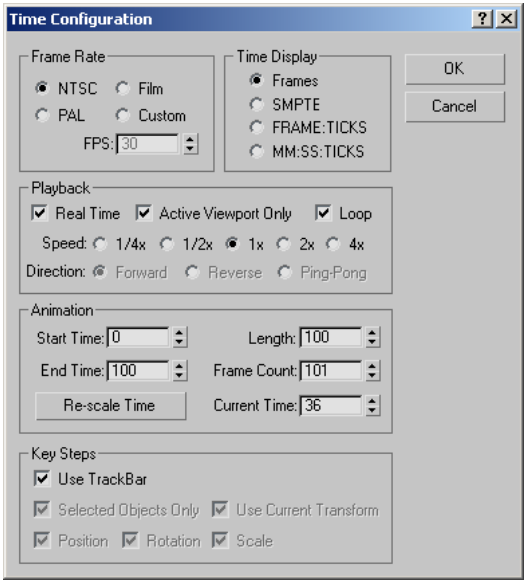
- How time is measured and displayed.
- The length of the active time segment (the part of the animation in which you're currently working).
- How much time is covered by each rendered frame of your animation.

Other issues in the topics which follow describe how to move through time and how to view animation in the viewports.

Choosing the Time Display Format

When you start 3ds Max, the default time display is in frames, but you can use alternative time-display formats. For example, you might want to see time in seconds and minutes.

You can specify different time-display formats using the *Time Configuration dialog (page 3–725)* > Time Display group settings. When you change the time display format, you not only change the way that time is shown in all parts of the software, but you also change the method with which you access time.



You can use these time display formats:

Frames—Displays time in whole frames.

This is the default display mode. The amount of time covered by a single frame depends on your choice for the current frame rate. For example, in NTSC video each frame represents 1/30th of a second.

SMPTE—Displays time using the Society of Motion Picture and Television Engineers format.

This is the standard time-display format for most professional animation work. From left to right, the SMPTE format displays minutes, seconds, and frames, delineated by colons. For example, 2:16:14 represents 2 minutes, 16 seconds, and 14 frames.

FRAME:TICKS—Displays time using frames and the program's internal time increment, called "ticks."

There are 4,800 ticks per second, so you can actually access time intervals as small as 1/4800 of a second.

MM:SS:TICKS—Displays time in minutes (MM), seconds (SS), and ticks, delineated by colons.

For example, 2:16:2240 represents 2 minutes, 16 seconds, and 2,240 ticks.

Setting Time Segments

The active time segment specifies a block of working time. You might think of it as a window in time that you use to focus on a specific part of your animation.

Specifying an Active Time Segment

You specify the active time segment by setting the Start Time and End Time for the segment on the *Time Configuration dialog (page 3–725)*.

You can change the active time segment whenever you want without affecting the keys you've already created. For example, if you have keys scattered over a range of 1000 frames, you can narrow your active time segment to work only on frames 150-300.

Changing the active time segment has the following effect:

- Restricts the range of time you can access using the *time slider (page 3–701)*.
- Restricts the range of time displayed when using the animation playback buttons.

The default setting for the active time segment runs from frames 0 to 100, but you can set it to any range.

You can also increase the active time segment using the *track bar (page 3–703)*.

Rescaling the Active Time Segment

You use the Re-scale Time button on the *Time Configuration dialog (page 3–725)* to change your entire animation based on the active time segment.

Use Re-scale Time as follows:

- Scale all animation in the active time segment to fit within a new time range.
- Move the entire animation to a new time.

Click Re-scale Time and then enter new Start Time and End Time values on the Re-scale Time dialog. When you click OK all the animation in the active time segment is moved and scaled to fit the new Start Time and End Time settings. Any animation outside the active time segment is moved to match the new active time segment boundaries.

For example, say you have an animation from frame 0 to frame 300 and an active time segment starting at frame 100 and ending at frame 200. Use Re-scale Time to set the new Start Time to frame 200 and the new End Time to frame 250. Clicking OK gives you the following result:

- Animation in the active time segment moves forward 100 frames and shrinks to a length of 50 frames. The new active time segment is from frames 200 to 250.
- Animation in the frames preceding the original active time segment moves forward 100 frames to connect to the start of the new active time segment.
- Animation in the 100 frames after the original active time segment moves forward 50 frames to connect to the end of the new active time segment.

You can also rescale the active time segment using the *track bar (page 3–703)*.

Moving Through Time

You can move to any time in your active time segment by using either the *time slider (page 3–701)*, or the *Current Frame field (page 3–724)* in the time controls area. You can also move through time using the playback control buttons.

Using the Time Slider

The time slider shows you the current time, and lets you move to any time in your active time segment.

To change the current time using the time slider, do one of the following:

- Drag the time slider.
- Click in the empty track to either side of the time slider.
- Click the increment arrows at either end of the time slider.

When you click in the slider track, the time slider jumps to the time where you clicked. This is a faster method of moving through time than dragging the time slider.

The time slider displays the current time, followed by a slash (/), followed by the total time in the active time segment. For example 25/100 means frame 25 of 100 frames. The current time also appears in the current time field. If your scene has been animated, it's played back in all viewports as you drag the time slider.

You can use the < and > keys on the keyboard to advance the time slider a frame at a time. When Key mode is on, this will advance to the next keyframe.

Moving to an Exact Time

The Current Time field always displays the current time. You can also enter a time value and press Enter to move to that time.

Using the Time Control Buttons

You use the Time Control buttons to move forward and backward in time and to play your animation in one or more viewports. They work like VCR controls that you use to move through frames and to start and stop animation playback. The Time Control buttons include four buttons for moving

through time and a center flyout for controlling animation playback:

Go To Start (page 3–722)

Previous Frame/Key (page 3–723)

Play/Stop (page 3–723)

Next Frame/Key (page 3–724)

Go To End (page 3–724)

They work like VCR controls that you use to move through frames and to start and stop animation playback.

Choosing a Frame Rate and Playback Speed

The frame rate of an animation is expressed in frames per second (FPS). This is the number of frames the software displays and renders for every second of real time.

Because the program stores your animation keys using real time using an internal precision of 1/4800 of a second, you can change the frame rate for your animation at any time without affecting your animation timing.

For example, if you create three seconds of animation using the NTSC video frame rate of 30 FPS, you will have a 90-frame animation. If you later discover you need to output to PAL video, at 25 frames per second, you can switch to that frame rate, and your animation is now set to 75 frames of output. No change in the timing of your animation has occurred. Only the number of frames that 3ds Max will display and render has changed.

Setting the Frame Rate

You use the settings in the *Time Configuration dialog (page 3–725)* > Frame Rate group to switch back and forth between frame rates at any time.

NTSC: U.S. and Japanese video standard of about 30 frames per second.

PAL: European video standard of 25 frames per second.

Film: Movie standard of 24 frames per second.

Custom: Frame rate set in the FPS parameter.

Configuring Animation Playback

You use settings in the *Time Configuration dialog (page 3–725)* > Playback group to specify the playback speed, and the number of viewports that play the animation.

Real Time: Animation plays at the selected playback speed, skipping frames, if necessary, to maintain the correct speed. Turn this off and the animation will play every frame without trying to maintain the correct speed. The different playback speeds are also useful when using the *Motion Capture utility (page 2–655)*.

Active Viewport Only: Animation plays only in the active viewport. Turn this off and the animation will play in all four viewports at once.

Speed: Choose one of these options to multiply the frame rate by the selected speed.

Direction: When Real Time is off you have the option to change the direction of the animation playback. Choosing Reverse will play the animation from end to start. Choosing Pingpong will play the animation from start to end, and then end to start.

Loop: When Loop is turned off, the animation will play once and stop.

Viewport Playback Speed

The ability of the program to play your animation at a specified rate depends on many things, including the complexity of the scene, the number of objects moving in the scene, the geometry

display mode, and so on. The worst case is a camera move in shaded mode, in which the viewport is filled with detailed geometry. In such cases, it's best to simplify the viewport display, using either wireframe display or, in extreme cases, box display mode.

Naturally, it takes more computing power to display your animation in four viewports, and playback smoothness is reduced. When Active Viewport Only is on, you can switch active viewports during playback either by clicking the label of an inactive viewport, or by right-clicking in an inactive viewport.

Working with Controllers

Everything you animate in 3ds Max is handled by a controller. A controller is a plug-in that handles the storage and interpolation of all animated values.

The default controllers are:

- **Position**: Position XYZ
- **Rotation:** Euler XYZ
- **Scale:** Bezier Scale

Note: New in 3ds Max 9 Extension 1 is the ability to double-click a controller track name on the Motion panel or Track View hierarchy window for fast access to key info or controller parameters. This productivity enhancement can help speed your work when fine-tuning animation.

Although 3ds Max has many different types of controllers, much of the animation is handled by the *Bezier controller (page 2–310)*. Bezier controllers interpolate between keyframes in a smooth curve. You can adjust the key interpolation of these interpolations through the keys on the track bar or in Track View. This is how you can control acceleration, hesitation and other types of motion.

The default controller for Rotation is Euler XYZ, which breaks the rotation down into three individual Bezier Float tracks. The default controller for Position is Position X,Y,Z. The Scale controller default is Bezier.

Note: If you load in files made in earlier versions of the software, their existing controllers will be maintained.

Warning: Be aware that Euler rotations behave differently than TCB rotations. If you are used to using TCB controllers, you can reassign TCB Rotation as the controller and you will get the same behavior you are used to.

3ds Max has a specialized type of controller, called a *constraint*, that is commonly used to help automate the animation process. A constraint can be used to control an object's position, rotation, or scale through a binding relationship with another object.

You apply constraints and controllers using commands on the Animation menu. When you assign a controller from this menu, a weighted list controller is automatically applied, with the controller you've selected appearing first in the list. The weighted list controller gives you the ability to blend controllers, similar to a nonlinear animation system. If you assign a controller through the Motion panel or Track View, it replaces the existing controller, rather than creating a list controller. You can do that manually if you are working in the Motion panel or Track View.

This section explains techniques for working with controllers. For details about the parameters and use of each controller type, see *Animation Controllers (page 2–307)*.

Understanding Controllers

Controllers are plug-ins that handle all of the animation tasks in 3ds Max, including:

- Storing animation key values.
- Storing procedural animation settings.
- Interpolating between animation key values.

Most animatable parameters don't receive a controller until you animate them. As soon as you change an animatable parameter at any frame other than 0 with the Auto Key button on, or click in its track using Curve Editor > Add Keys, a default controller is assigned to the parameter.

Accessing Controllers

There are two places where you work directly with controllers:

-  **Track View:** Controllers are indicated in the Hierarchy list by the various controller icons. Each controller has its own individual icon. Using Track View, in either Curve Editor or Dope Sheet mode, you can view and work with the controllers for all objects and all parameters. See *Track View (page 2–501)*.
-  **Motion panel:** Contains special tools for working with transform controllers. The Motion panel contains many of the same controller functions as the Curve Editor, plus controls necessary for working with special controllers such as IK Solvers. Using the Motion panel you can view and work with the transform controllers of a single selected object. See *Motion Panel (page 3–774)*.

Categories of Controllers

There are two main categories of controllers. These categories are easy to identify when looking at the Track View - Curve Editor Hierarchy list.

Single-Parameter Controllers: Control animation values of a single parameter. Regardless of whether the parameter has a single component, such as the number of sides of a cylinder, or multiple components, such as the RGB values of a color, the controller is handling a single parameter.

Compound Controllers: Combine or manage multiple controllers. Compound controllers include high-level Transform controllers, such as *PRS (page 2–357)*, the *Euler XYZ Rotation controller (page 2–318)*, the *Transform Script controller (page 2–379)*, and the *List controller (page 2–342)*.

A compound controller appears in the Hierarchy list as a controller icon with subordinate-level branches of other controllers.

Controllers and Constraints

In addition to controllers, the software can animate using constraints. These items are located in the Animation > Constraints menu. The constraints include the following: Attachment, Surface, Path, Link, Position, Orientation, and LookAt.

When you assign a controller in the Motion panel or in Track View, you will see these constraints appearing in the list of available controllers. You can assign them the same as the other controllers, but they are also assignable from the Animation > Constraints submenu. For more information, see *Animation Constraints (page 2–392)*.

Note: If you assign a controller using the Animation > Controller submenu, a list controller is automatically applied to the object, with the selected controller applied beneath the list controller. This is different than controllers that are applied via the Motion panel.

Viewing Controller Types

You can view the controller type assigned to a parameter in both the Curve Editor and in the Motion panel. Before you can view the controller types in Track View, you must do the following:

1.  On the Curve Editor toolbar, click the Filters icon. Then in the Filters dialog > Show group, turn on Controller Types.

   You can then see the name of the Controller type in the Hierarchy view.

2.  Assign controllers to parameters. You can do this on the Assign Controller rollout of the Motion panel, or through the right-click menu in the Hierarchy list in Track View.

The Parameters mode of the Motion panel always displays the transform controller types for the selected object.

Reading Controller Types

You can tell a lot about how a parameter is animated by looking at the controller type.

Parameter Name: Is always visible and is always to the right of the controller icon. It tells you what is animated.

Interpolation Type: Usually follows the parameter name. It tells how animation values are calculated. The label Diffuse: Bezier Color indicates a Diffuse color parameter using Bezier interpolation with Color data.

Data Type: Usually follows the interpolation type. It tells what type of data is used. The label Height: Linear Float indicates a Height parameter using Linear interpolation with a floating point value.

Changing Controller Properties

Certain controllers, including procedural ones like *Noise (page 2–353)*, do not use keyframes. For this type of controller, you can analyze and change your animation by editing controller parameters by means of a Properties dialog. The controller type determines whether or not the controller displays a properties dialog and the type of information displayed.

Using Curve Editor you can view controller properties dialogs for multiple tracks simultaneously. The following rules govern viewing multiple controller properties dialogs:

- Each track displays only one properties dialog in each Track View window.
- When properties dialogs for multiple tracks are visible, only one dialog can be active.
- Properties dialogs for tracks that use keys are disabled unless keys are selected.

Changing Controller Properties

Some controllers do not use keys, using instead a properties dialog that affects the entire animation. Such controllers are usually parametric controllers like Noise, or compound controllers like List.

To view controller properties in Curve Editor:

Do one of the following:

- Highlight the label for a parametric or compound controller, right-click the label, and then choose Properties from the shortcut menu to display the properties dialog.

  The dialog title identifies the controller type, the item or object name, and the parameter. For example: Noise Controller:-Box01\Position

- Double-click the label.

You can also view global properties for some transform controllers in the Motion panel. The same controls described above apply.

Controller properties can also be viewed from the *track bar (page 3–703)*. Right-click any key and choose Controller Properties.

Changing Controller Key Information

Some controllers interpolate between keys that you set on specific frames. Such controllers are always

single parameter controllers such as a Bezier Float controller for Height or TCB for Rotation. These controllers use a Key Info dialog that contains settings for one of more selected keys.

To view key information in the Curve Editor:

- Right-click a key to display the Key Info dialog.

  If more than one key is highlighted, Key Info displays common information for all of the selected keys. Settings that contain values indicate values common to all of the selected keys. Settings that are blank indicate values that change from key to key.

Viewing key information in the Motion panel always displays the settings for a single transform key.

To view key information for transform controllers in the Motion panel:

1. Select an object.
2. On the Motion panel, click Parameters, if it's not already active.
3. Click Position, Rotation, or Scale on the Parameters rollout.

   If the transform controller uses keys, Key Info rollouts appear below the Parameters rollout.

To view key information for transform controllers in the track bar:

1. Select an animated object.
2. Right-click any key in the track bar.
3. From the right-click menu, choose the property to inspect (e.g., Sphere01: X Position), or, if available, choose Controller Properties.

Changing Controllers That Do Not Display Properties

Some controller types do not display any properties in Track View or the Motion panel. You change the animation values of these controllers using the Auto Key button in the viewports and using the tools in the Curve Editor Keys and Curves menus, and in the track bar.

Assigning Controllers

Every parameter has a default controller type that gets assigned the moment the parameter is animated. You can choose from multiple controller types for any parameter and change controllers after the parameter is animated.

Assigning Controllers in Track View

You can assign controllers to any animatable parameter in the Curve Editor by selecting controller items and then choosing Assign Controller on the Controller menu.

You can also assign controllers to any animatable parameter in the Track View — Curve Editor by selecting the parameter in the Hierarchy List then right-clicking and choose Assign Controller from the quad menu.

You can also assign the same controller type to a selection of multiple parameters as long as all the selected parameters can use the same type of controller. For example, you could select the Length, Width, and Height parameters for multiple Box objects and assign the same controller type to all of them. This is because they all use controllers that work on floating point data.

If a parameter has already been animated, then assigning a new controller has one of the following effects:

- The existing animation values are recalculated to produce a similar animation with the new controller. For example, replacing TCB Position with Bezier Position closely preserves the animation.
- The existing animation values are discarded. For example, replacing Smooth Rotation with

Noise Rotation discards the Smooth Rotation animation values.

Assigning Controllers in the Motion Panel

You can assign controllers on the Motion panel by selecting an object and then, on the Assign Controller rollout, choosing a controller and clicking the Assign Controller button. You can change the controllers of only one selected object.

Assigning Controllers Using the Animation Menu

You can assign controllers using the Animation menu. All the controllers and constraints are available by going to the menu bar and choosing Animation, and then selecting the type of Controller you want and picking the controller from the submenu that displays.

When you assign a controller in this method, a list controller is automatically assigned, and the controller you have selected appears as the first entry in the list. This automatic list controller assignment does not occur if you assign controllers using the Motion panel or Track View. Weighted list controllers allow you to blend between various tracks by animating the weights.

Tip: You can use the Animation menu to assign the same type of controller or constraint to several different objects at once. Just select the objects and then choose the controller or constraint from the Animation menu.

Note: If an assigned constraint requires designating an additional subject, such as a Path constraint, a rubber-band line extends from the first selected object to the mouse cursor after choosing the constraint. To complete the assignment, position the cursor over a qualifying target object and click.

Copying and Pasting Controllers

Right-click any blank area on the Track View toolbars, then choose Show Toolbars > Controllers: Track View. This displays the Controller toolbar with buttons for quick access to controller tools.


Click Copy and Paste on the Track View toolbar to copy and paste controllers. For a general discussion of using Copy and Paste, see *Copying and Pasting Items (page 2–575)*.

To copy and paste controllers in the Curve Editor, select the controller track to copy, then right-click and choose Copy from the Track View quad menu. Navigate the Hierarchy list to the target object, select the track, right-click and choose Paste.

Rules for using Copy and Paste are as follows:

- You can copy only single controllers. Compound controllers like List or PRS Transform controllers are considered single controllers for Copy and Paste operations.
- You can paste a copied controller into one or more controllers using the same data type.
- You can choose to make an instance or a copy of the pasted controller.
- You can choose to convert other controller instances in the scene automatically.

Clicking Paste displays the Paste dialog, with three controls for determining how the Paste operation is carried out.

Copy—Pastes the controller as a copy.

Instance—Pastes the controller as an instance of the source controller. Any change you make to either controller will affect the other.

For example, you can paste a box's Length controller as an instance into its Height and Width parameters. This makes the box a cube. Changing either of the Length, Width, or Height parameters changes the other two.

Paste Target: Replace All Instances—When selected, all instances of the target controller receive the paste controller, whether or not they are selected. This keeps all instances of the target controller as instances. When off, the target controller is made unique and the remaining instances are unchanged.

Specifying Default Controllers

You can specify the permanent defaults for controller types and controller settings to match the way you prefer to work. The following defaults are written to the *3dsmax.ini* file:

- The default controller used for each data type
- The default controller settings

Specifying Default Controllers

You specify default controllers by choosing Controller > Assign from on the Track View menu (see *Assigning Controllers (page 2–292)*). When you choose a controller in the Assign Controller dialog you have the option to click Make Default before clicking OK.

Clicking Make Default assigns the chosen controller as the default for all parameters using that data type. It has the following effects:

- The default controller is listed at the bottom of the Assign Controller dialog.
- Many different parameters might share the same data type. For example, selecting the Length parameter of a Box and specifying Linear Float as the default sets the default controller for all parameters that use the Float data type. This includes, Width, Camera FOV, and Scale Deformation Curves.
- Default controller choices are written to the *3dsmax.ini* file and become the default for all new scenes.
- Previously assigned controllers are not affected.

Specifying Default Controller Values

You can specify the default settings for many controller types, or reset controllers to their factory settings on the Preferences dialog.

To specify default controller settings:

1. Choose Customize menu > Preferences.
2. In the Preferences dialog > Animation panel > Controller Defaults group, click Set Defaults to display the Set Controller Defaults dialog.
3. From the list of available controllers, choose a controller type and click the Set button to display the default settings supported by the selected controller. For example, with a Bezier controller you can set the In and Out tangents.

   Note: The default Rotation controller is Euler XYZ, not TCB (Quaternion) as in earlier versions of the software. The default Position controller is now Position XYZ, instead of TCB Position. Only Scale is still Bezier as a default.

   Euler XYZ Rotation behaves quite differently from TCB. It gives you three function curves to manipulate, but does not allow rotations of greater than 180 degrees between keys. If you are used to working with TCB controllers, you can change the default rotation controller back to TCB.

Once you click OK, the controller defaults are changed. Changes to the controller default settings are written to your *3dsmax.ini* file and become the defaults for all newly assigned controllers and all new scenes.

You can also revert to the original program defaults for all controllers by clicking Preference Settings dialog > Animation panel > Controller Defaults group > Restore To Factory Settings.

General-Purpose Controllers

The following controllers are general purpose in that you can apply them to parameters of different data types, yet they behave in essentially the same way for those different parameters.

Within certain general-purpose controllers there might be variations according to the data type used by a parameter.

See *Animation Controllers (page 2–307)* for detailed descriptions of controller properties.

Bezier Controllers

Bezier controllers (page 2–310) interpolate between keys using an adjustable spline curve; they are the default controller for most parameters.

Use Bezier controllers to provide fully adjustable interpolation between keys. Bezier controllers support the following options:

- Adjustable tangent handles.
- Step tangents for abrupt changes from one key to the next.
- Constant velocity control.

You can adjust the key interpolation by choosing among different tangent types in the Key Info dialog. .

TCB Controllers

The *TCB controller (page 2–377)* produces curve-based animation like Bezier controllers. However, TCB controllers use fields to adjust the Tension, Continuity, and Bias of the animation.

Use TCB controllers when you want adjustable, curved interpolation between keys, and you want to use TCB style controls. TCB rotation controllers were the default controller type in versions 4 and earlier.

Linear Controllers

The *Linear (page 2–341)* controller interpolates between animation keys by evenly dividing the change from one key value to the next by the amount of time between the keys.

Use Linear controllers when you want motion to have a rigid, mechanical look.

Noise Controllers

The *Noise (page 2–353)* controller produces random, fractal-based animation over a range of frames. Noise controllers are parametric; they do not use keys.

Noise controllers have many possible uses; as in the following examples:

- Use Noise whenever you need completely random animation around a given value. For example, use a Noise Rotation controller when you want an object to wobble in place. A common use for a noise controller is the creation of camera shake.
- Use Noise in a List controller to apply variations to the result of another controller. For example, use a List controller to combine Noise Position with Bezier Position. The Bezier controller moves the object while the Noise controller makes the object shake and stray a little from the trajectory.

XYZ Controllers

The XYZ controllers such as *Euler XYZ (page 2–318)* and *Position XYZ (page 2–356)* are specifically designed so that you have three separate curves, one for each axis. This allows you to independently view and control the curves individually. This has advantages over rotation controllers like TCB that do not display function curves at all. XYZ controllers are now the default for rotation animation.

You can adjust the interpolation between keys using the Key Info dialog.

Audio Controllers

The *Audio controller (page 2–309)* converts the amplitude of a recorded sound file or real-time sound wave into values that can be used by an animated parameter.

Use the Audio controller to synchronize parameter values with a sound file. For example, use an Audio controller for a Multiplier Curve to scale a parameter in sync with a sound.

Special-Purpose Controllers

The controllers described in this topic are also applied to parameters of different data types, but they are used for special purposes.

List Controllers

The *List controller (page 2–342)* combines multiple controllers into a single effect. It is a compound controller with tools for managing the order in which its internal controllers are calculated. List controllers are weighted, by animating the weights of the different layers you can create an effective non-linear animation system.

Use List controllers to combine controllers as in the following examples:

- Combine Noise Rotation and TCB Rotation controllers to introduce random orbital motion as an object rotates.
- Combine Bezier Position and Path Constraint to make an object follow a path with manually keyframed variation away from the path.

  When you apply any controller from the Animation menu, a list controller is automatically placed on the object and the selected controller placed first in the list.

For hands-on experience using weighted list controllers see the bouncing ball exercise in the tutorials.

Expression/Script Controllers

You write custom code for *Expression controllers (page 2–320)* and *Script controllers (page 2–372)* in an Expression or Script controller dialog. You specify parameters using mathematical expressions, functions, and variables. The code can include values based on the controllers of other objects in the animation.

Note: You can replicate some functions of Expression and Script controllers without having to write code by using parameter wiring or Reaction controllers. See *Parameter Wiring Dialog (page 2–412)* and *Reaction Controllers (page 2–358)*.

See also

Expression Controller Techniques (page 2–324)

Trigonometric Functions (page 1–150)

Vectors (page 1–151)

Motion Capture Controllers

A Motion Capture controller controls parameters in real time from the input of external devices. Currently supported devices are mouse, keyboard, MIDI device and joystick. Each device has specific properties that must be set:

- After assigning the motion-capture controller, open the controller Properties dialog and bind the type of peripheral device and set device parameters.
- Use the *Motion Capture utility (page 2–655)* to test and record captured motion.

Explicit Axis Keys

The XYZ controllers assign an individual track to the X, Y, and Z components of the position, rotation, or scale of an object. However, the controllers always assign three keys (one for each axis), by default.

In previous versions of the software, you had to manually edit the keys in order to create explicit axis keys. However, there are now actions available from the *Customize User Interface dialog (page 3–792)* that let you create explicit keys with the touch of a button.

When an explicit key is created on an axis, any keys on the two remaining axes (at the selected time) will be deleted.

An explicit axis key can only be assigned to objects that already have an XYZ controller assigned to their position, scale, or rotation tracks.

You can also use Keyable Icons to turn on or off tracks for editing. This allows you to set keys on only specified PRS X, Y or Z axes as well. See *Keyable Icons (page 2–531)*.

Interface

Note: These commands are not available in the default user interface. You can use the Customize User Interface dialog to assign a keyboard shortcut to them or add them to your toolbars or menus if you want to use them.

Key Position X, Y, or Z—Creates an explicit key on the designated position axis.

Key Rotation X, Y, or Z—Creates an explicit key on the designated rotation axis.

Key Scale X, Y, or Z—Creates an explicit key on the designated scale axis.

See also

Position XYZ Controller (page 2–356)

General-Purpose Controllers (page 2–295)

Euler XYZ Rotation Controller (page 2–318)

Scale XYZ Controller (page 2–371)

Float Controllers

The following controllers are available only for parameters that use the Float data type.

On/Off Controller

An *On/Off controller (page 2–355)* controls parameters using binary on and off values. The On/Off track displays a solid blue color in frames that are on, and the background in frames that are off. The on/off state of the parameter switches every time you add a key.

Use for binary parameters such as the Smooth object parameter, or for Visibility tracks.

Boolean Controller

The Boolean Controller is a variation on the On/Off controller. It provides a similar functionality, but with one significant difference. The on/off state of the parameter does not switch automatically every time you add a key. This allows you to effectively add keys in the middle of a sequence without creating unintended drastic changes.

Waveform Controller

A *Waveform controller (page 2–381)* generates values by combining periodic waveforms. You can specify different types of waveforms and add waveforms together to create complex animation.

Waveform was originally created to control blinking lights. Use it for any value that you want to vary in a regular, oscillating pattern.

Controlling Transforms

Transform controllers are compound controllers. They set the type and behavior of the controllers used for Position, Rotation, and Scale.

You assign Transform controllers using either Track View — Curve Editor or the Motion panel. Compound Transform controllers do not display properties in Track View. You can access the parameters of Transform controllers only in the Motion panel.

Position/Rotation/Scale Controller

The *Position/Rotation/Scale (PRS) controller (page 2–357)* is a simple Transform controller applicable for most objects. When applied the PRS Transform controller sets up default Bezier Position, TCB Rotation, and Bezier Scale controllers.

Use a PRS controller whenever you want a standard transform set up or when you want single function curve control over the Position, Rotation, and Scale controllers.

Note: PRS controller is no longer the default controller applied to all objects. The latest version of the software uses Position XYZ and Euler XYZ as the new default controllers for position and rotation transforms.

Transform Script Controller

The *Transform Script controller (page 2–379)* contains all of the information contained in a PRS Controller in one scripted matrix value. Instead of having three separate tracks for position, rotation, and scale, all three values can be simultaneously accessed from one script controller dialog.

Because the transform values are defined by a script, they are easier to animate.

The value of the controller script must be a matrix3 value. A matrix3 value is a 4x3 3D transformation matrix. For more information, see the Matrix3 Values topic in the MAXScript Reference.

XRef Controller

The *XRef controller (page 2–383)* lets you externally reference any type of Transform controller from another scene file. When you assign this controller to your object, it nests the source controller, making it only accessible for playback. You can use the XRef controller either on its own, or combine it with an *XRef Object (page 3–394)*.

Controlling Position

Below the *Transform controller (page 3–909)* is the Position controller. Position is a data type that can use most of the standard controllers such as Bezier, TCB, and Noise.

The controllers described in this topic can be used only as position controllers.

XYZ Controller

The XYZ controller breaks a Position controller into three separate Bezier Float controllers. Each of the X, Y, and Z components of position receives its own track.

Use the XYZ controller when you want to have separate key patterns or controller types for each position component.

See *Euler XYZ Rotation Controller (page 2–318)*, *Position XYZ Controller (page 2–356)*, and *Scale XYZ Controller (page 2–371)*.

Controlling Rotation

Below the Transform controller is the Rotation controller. Rotation is a data type that can use most of the standard controllers such as TCB, Linear, and Noise.

Rotations in 3D are very complex. Even the standard controller types behave differently when used for rotation. The most common way of calculating rotations in 3D animation uses four components to define rotation about an arbitrary axis. This is the *quaternion* method.

The benefits of using quaternion rotation include:

- It produces a direct one-to-one relationship between the quaternion values and how objects interactively rotate in the scene.
- It produces smoother rotation than other methods.

The drawbacks of using quaternion rotation include:

- Rotation values in key information can be difficult to understand.
- Quaternion rotation controllers do not display function curves in Track View — Curve Editor. For this reason Euler XYZ is now the default rotation controller assigned to all objects.

You can use the following controllers only as rotation controllers.

Smooth Rotation

Use *Smooth Rotation (page 2–374)* when you want rotation to have a smooth and natural look. Smooth rotation uses nonadjustable curved interpolation, and has the following characteristics:

- You can move keys in Track View to change timing.
- You can directly rotate objects in the viewports to change rotation values.
- You cannot display controller or key properties, or function curves.

Euler XYZ Rotation

Use the *Euler XYZ Rotation controller (page 2–318)* when you want individual function curve control for each axis of rotation. Euler XYZ is a compound controller that combines separate, single-value float controllers to specify an angle of rotation about each of the X, Y, and Z axes. Euler XYZ Rotation is the default controller applied to all objects.

Euler XYZ is not as smooth as quaternion rotation, but it is the only rotation type that you can use to edit rotation function curves.

Euler angles are well suited for animating mechanical rotations. They can also be a good choice for an object in an inverse kinematics chain because IK Rotational joints are defined as Euler angles. X, Y, and Z function curves displayed by this controller in Track View represent rotations in "world" or "parent" space which is the rotation system used by most 3D animation packages.

Customers upgrading from earlier versions of the program might find some difficulty in adapting to using Euler XYZ rotation controllers. If you find yourself experiencing unexpected rotation behavior, this might be the reason. Reassign TCB rotation as the default rotation controller and you will be able to regain your customary workflow.

Controlling Colors

You can animate colors in 3ds Max just like most other parameters. There are two data types used for controlling colors.

Point3: A general-purpose, three-component data type that works with RGB color values. It uses most of the standard controllers.

Color: A special data type designed specifically for working with RGB and HSV color values. Color uses the Bezier and RGB controllers.

Color Point3 Controller

You can assign any of a variety of Point3 controllers to a material's color channels, including Point3 Expression, Point3 List, Color RGB (described later in this topic), and so on. When using Point3 controllers for color parameters, an important issue is behavior of the Key Value fields, labeled X, Y, Z. They store color values using only the RGB color model.

- The X value field stores Red color values.
- The Y value field stores Green color values.
- The Z value field stores Blue color values.

The Point3 Key Value fields do not clamp at the valid color range of 0 to 255. Values that drop below 0 or exceed 255 are ignored by the color parameter but are still displayed in Track View.

Bezier Color Controller

The Bezier Color controller (see *Bezier Controllers (page 2–310)*) is a data type that uses Bezier key interpolation. You can use RGB or HSV color models with the Bezier Color controller.

The choice of color model is global for all keys used by the controller. You can switch between color models anytime and color key values are correctly converted.

The Color Value fields are limited to a range of 0 to 255. You can drag the color keys above 255 in Track View Curve Editor, but the values are clamped at 255.

Color RGB Controller

Assign a *Color RGB controller (page 2–317)* to break a Color controller into three separate Bezier Float controllers. Each of the R, G, and B components of color uses its own track.

Use the Color RGB controller when you want to have separate key patterns or controller types for each color component.

Morph Controllers

You can choose from two morph controllers: Cubic and Barycentric; they manage how morph targets blend from one target to another. Morphing can also be achieved by applying a *Morpher modifier (page 1–729)* to an object.

Cubic Morph Controller

A Cubic Morph controller is a TCB-style controller. It uses Tension, Continuity, and Bias controls much like the generic TCB controller.

The Cubic Morph controller manages only the interpolation from one morph target to the next. If you want to add Morph keys or change the morph target, you must use the *Pick Targets and Current Targets rollouts (page 1–314)* in the Modify panel.

Barycentric Morph Controller

The *Barycentric Morph controller (page 2–309)* is also a TCB controller like the Cubic Morph controller, except that instead of each key representing a single target, each key represents a series of weights for all targets. A Barycentric Morph key represents a new object which is a blending of all targets.

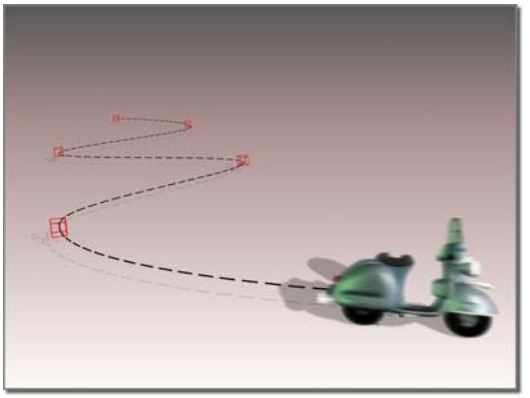
You can adjust each morph key to percentages of the available morph targets, creating subtle adjustments in the animation.

You can also click Add Keys in Track View to create new Barycentric Morph keys. The added keys contain interpolated values for all targets.

Motion Panel Commands

Trajectories

Select an object > Motion panel > Trajectories



Scooter following a trajectory

The Trajectories rollout displays the path an object travels over time. Controls to convert splines into trajectories, trajectories into splines, and to collapse any transform controller into editable keys are also available by choosing Motion panel > Trajectories.

Note: The state of the Trajectory check box under Display in the *Object Properties dialog (page 1–117)* (and on the Display panel > Display Properties rollout) controls whether trajectories will be visible for an object.

Using trajectories, you can do the following:

- Insert and delete keys from the path.
- Move, rotate, and scale keys on the path.
- Convert the path to a spline object.
- Derive a new path from a spline object.
- Collapse transforms.

Values under Sample Range are used in the Spline Conversion and Collapse Transform functions.

Note: You can now assign the four principal Trajectories functions to keyboard shortcuts and other custom UI items. The following new actions are available via *Customize User Interface (page 3–792)*:

- Trajectory Toggle: If one or more objects are selected, this enters and exits Trajectories mode on the Motion panel.
- Trajectory Key Mode Toggle: Enters and exits the Keys sub-object level for direct editing of animation keys via transformation of the trajectory keys.
- Trajectory Add Key Mode Toggle: Enters and exits *Add Key mode*.
- Trajectory Add Key Mode Toggle: Enters and exits *Delete Key mode*.

Procedures

To display an object's trajectory:

1. Select an animated object that moves over time.
2. Right-click the object and choose Properties.

   Trajectory is not available.
3. In the Display Properties group, click By Layer to change this button to By Object. Skip this step if By Object is already displayed.

   Trajectory becomes available.
4. Turn on Trajectory, then click OK.

   The Trajectory is displayed as a red line with white squares and dots. The white squares are keys, the white dots are in-betweens.

To add a key to a trajectory:

1. Select an object.
2. Display the trajectory by following the previous steps.
3. In the Motion panel, click Trajectories.
4. Click Sub-Object to activate Keys and enable editing.
5. Click Add Key.

   The button highlights.
6. Click the trajectory.

   Wherever you click the trajectory a key will be added.
7. Right-click the key and select Key Info from the right-click menu.

   In the Key Info dialog, you can change the In and Out values for the in-betweens on either side of the key. This allows you to use the trajectory key to create the illusion of hesitation or acceleration.

   You can click repeatedly to add many keys in this mode.

To delete a key from a trajectory:

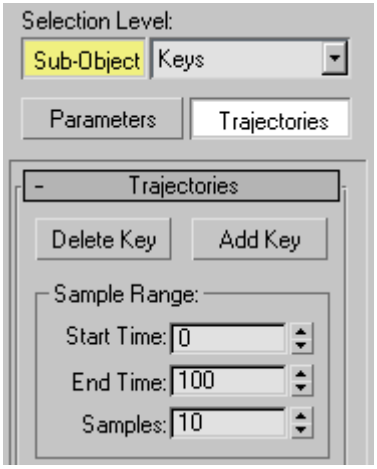
1. Select a key on a trajectory. You can click it, or use window selection.
2. Be sure Add Key is turned off in the Trajectories rollout of the Motion panel.
3. Click Delete Key in the Trajectories rollout of the Motion panel.

   The key is deleted.

To transform keys on a trajectory:

1. Select one or more objects.
2. On the Motion panel, click Trajectories.
3. Click Sub-Object to activate Keys and enable editing.
4. Select one or more keys and use the transform tools on the Main toolbar to move, rotate, or scale the selected keys.

Interface



Sub-Object—Enables key editing. Use the Move, Rotate, and Scale transforms to change the location of a key(s) displayed on a trajectory.

Trajectories rollout

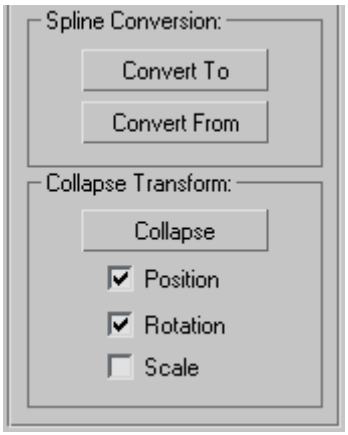
Delete Key—Deletes the selected key(s) from the trajectory.

Add Key—Adds key(s) to the trajectory. This is a modeless tool. When you click this button once, you can then add any number of keys by clicking the trajectory line in the viewport one or more times in succession. To exit Add Key mode, click the button again.

Sample Range group

Start Time/End Time—Specifies the interval for the conversion. If you're converting from position keyframes to a spline object, this is the time interval for which the trajectory will be sampled. If you're converting from a spline object to position keys, this is the interval over which the new keys will be placed.

Samples—Sets the number of samples for the conversion. When converting in either direction, the source is sampled at regular intervals, and keys or control points are created on the target object.



Spline Conversion group

Convert To/Convert From—Converts keyframe position tracks to and from spline objects. This enables you to create a spline trajectory for an object and then convert that spline to keyframes for that object's position track in order to do various keyframe-specific functions (such as applying constant velocity to the keys and normalizing the time). Or, you can convert an object's position keyframes into a spline object.

Tip: Convert To and Convert From can be used to move an object along a path without using a path constraint.

Collapse Transform group

Generates keys based on the current transformation of the selected object. You can apply this for any type of transform controller assigned to an object, but the main purpose of this function is to "collapse" a parametric transform effect, such as one generated by a Path constraint, into standard, editable keys.

Collapse—Collapses the transform of the selected object.

Position, Rotation, Scale—Specifies which transforms you want to collapse.

You must choose at least one check box to activate the Collapse button.

PRS Parameters Rollout

Select an object > Motion panel > Parameters > PRS Parameters rollout

The PRS Parameters rollout provides tools for creating and deleting keys. PRS stands for the three basic transform controllers: Position, Rotation, and Scale.

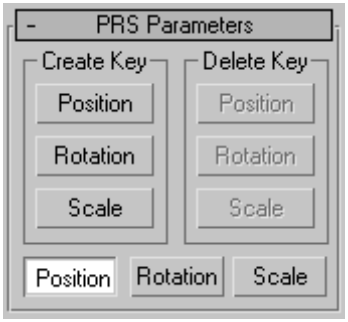
Procedure

To create a PRS transform key:

1. Select an object in the viewports.
2. Drag the time slider to the frame where you want to place a key.
3. On the Motion panel, choose Parameters > PRS Parameters rollout.
4. Click one of the following buttons under Create Key PRS Parameters rollout:
   - Click Position to create a Position key.
   - Click Rotation to create a Rotation key.
   - Click Scale to create a Scale key.

   If a particular Position, Rotation, or Scale controller doesn't use keys, then that button is not available under Create Key. For example, you can't create Position keys if you're using a Noise Position controller.

Interface



Create Key/Delete Key—Creates or deletes a move, rotate, or scale key at the current frame. These buttons become active or inactive depending on the existence of a key type at the current frame.

For example, if you're on a frame containing a Scale key, the Scale button is inactive in the Create column because a key already exists. At the same time, the Position and Rotation buttons are inactive in the Delete column because there are no keys of that type to delete.

Position/Rotation/Scale—Determines the contents of the Key Info rollouts that appear below the PRS Parameters rollout on the Motion panel.

Key Info (Basic) Rollout/Dialog

Select an animated object > Motion panel > Parameters > Key Info (Basic) rollout
Select an animated object. > Right-click a key on the track bar. > Choose a controller track.
Curve Editor > Right-click a key.
Parameter Collector (page 1–138) > Click a Properties button.

The Key Info (Basic) rollout or dialog changes the animation value, time, and interpolation methods of one or more selected keys.

The following parameters are available with the Bezier controller, which is the default position controller. The Key Info controls for a TCB position controller are different.

See also

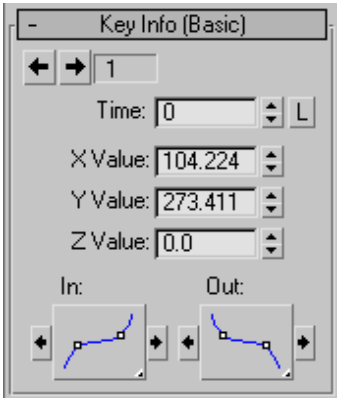
Key Info (Advanced) Rollout/Dialog (page 2–306)

Procedure

To set the tangent type for a key:

1. Select an object and open the Key Info (Basic) rollout.
2. Use the arrows in the upper-left corner of the dialog to select a key.
3. Choose a tangent type from the In or Out tangent flyouts.

Interface



Key number—Shows the current key number. Click the right or left arrows to go to the next or previous key.

Time—Specifies at what time the key occurs.

Time Lock—Prevents horizontal key movement in Track View Edit modes.

Value—Adjusts the position of the selected object at the current key.



Note: In the Key Info dialog *for the Bezier Scale controller*, a Lock button is displayed beside the X Scale spinner. If you click Lock X, only the X value affects all three axes of scale. The Y and Z values are ignored and their function curves are not displayed. When X is locked, the Y and Z values aren't affected by changes in the X value. If you click Lock X when all three axes are at identical values, alter the X value, and then unlock X. The Y and Z values remain where they were while X retains its new value.

Key Tangent Flyouts—With Bezier controller types, sets the interpolation properties of the In tangent and Out tangent of the key. For details, see *Tangent Types (page 2–305)*.

Tangent Copy buttons—Use the arrow buttons at either side of the Key Tangent flyouts to copy the tangent type between the tangents of the current key or between the tangents of the previous and next key.

- The left arrow of the In tangent copies to the Out tangent of the previous key.
- The right arrow of the In tangent copies to the Out tangent of the current key.
- The left arrow of the Out tangent copies to the In tangent of the current key.
- The right arrow of the Out tangent copies to the In tangent of the next key.

Tangent Types

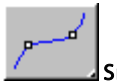
Key Info (Basic) dialog/rollout (page 2–304) > In/Out Tangent Type flyouts

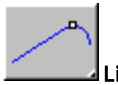
Status bar > Animation controls > Default In/Out Tangents For New Keys flyout

Select an animated object > Motion panel > Parameters > Key Info (Basic) rollout

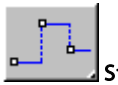
You can adjust the tangent type of the Bezier transforms. This means you can adjust the interpolation between keyframes to create particular motion effects. By manipulating tangent types, you can make something hesitate, speed up, slow down, or even stand still. Each key has two tangents: one to define the interpolation before the key, and a second to define interpolation following the key.

Tip: You can also set the default tangent type through the *Default In-Out Tangents For New Keys (page 3–721)* flyout. By doing so, each new key created with Set Key Mode or Auto Key Mode follows the curve interpolation set by the default tangent type.


Smooth—Creates smooth interpolation through the key.


Linear—Creates linear interpolation at the key.

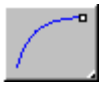
A linear tangent affects the curve near the key only. Full linear interpolation between two keys occurs only when the Out tangent of the first key and the In tangent of the next key both use a linear tangent.

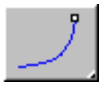

Step—Creates binary interpolation from one key to the next. Step tangents require a matched set between the Out tangent of one key and the In tangent of the next key.

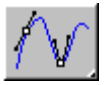
Choosing Step for the In tangent of the current key also sets the Out tangent of the previous key to Step. Likewise, choosing Step for the Out tangent of the current key also changes the In tangent of the next key to Step.

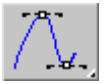
Using Step tangents, the outgoing value of a key is held constant until the time of the next key is reached. The value then abruptly jumps to the value of the next key. Use this tangent

when you want to animate On/Off switching or instantaneous changes from one value to the next.

 **Slow—**Causes the interpolated rate of change to slow down around the key. A slow In tangent decelerates as it approaches the key. A slow Out tangent begins slow and accelerates as it leaves the key.

 **Fast—**Causes the interpolated rate of change to speed up around the key. The effect is the opposite of using slow. A fast In tangent accelerates as it approaches the key. A fast Out tangent begins fast and decelerates as it leaves the key.

 **Custom—**Displays adjustable tangent handles at the key in Function Curves mode in 3ds Max.

 **Flat Tangent—** Displays a smooth interpolation type designed to eliminate overshoot with no editable handles. Tangent slopes automatically take the most direct route to the next key value.

See also

Default In/Out Tangents For New Keys (page 3–721)

Key Info (Advanced) Rollout/Dialog

Select an object > Motion panel > Parameters > Key Info (Advanced) rollout

Select an animated object in the viewport. > Right-click a key in track bar. > Choose the name of the controller track such as Sphere01: Z Position. > Key Info dialog > Click Advanced.

Parameter Collector (page 1–138) > Click a Properties button. > Key Info dialog > Click Advanced.

The Key Info (Advanced) rollout or dialog lets you control velocity in three ways:

- You can specify the absolute velocity at a key using the In/Out fields.
- You can average velocity over a time period using Normalize Time.
- With certain controller types, you can force constant velocity from one component key to the next using Constant Velocity.

If you are creating Path animation, you can force constant velocity by placing a *Normalize Spline modifier (page 1–747)* on the spline used as a path.

See also

Key Info (Basic) Rollout/Dialog (page 2–304)

Procedure

To set normalize time for a key:

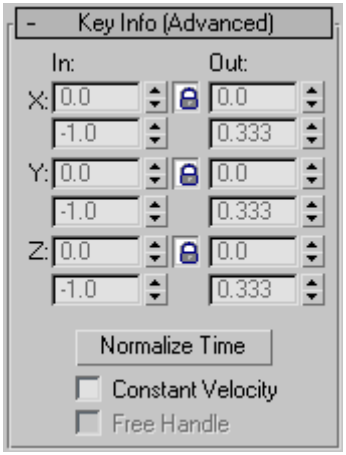
1. Select an object and choose a key to Normalize.

   You can use the arrows in the Key Info (Basic) rollout to step through Keyframes.

2. In the Key Info (Advanced) rollout, click Normalize Time.

   The key is moved in time to average the velocity through the key.

Interface



In/Out—The In field is the rate of change as the parameter approaches the key. The Out field is the rate of change as the parameter leaves the key.

- These fields are only active for keys using the Custom tangent type.
- The number in the field is the rate of change expressed as parameter units per tick. By changing the two values for X, Y, and Z you are changing the length and angle of the tangent handle.

Lock button—Changes one Custom tangent by changing the other an equal but opposite amount. For example, if you click the Lock button and the In value is 0.85, then the Out value is -0.85.

Normalize Time—Averages the position of the keys in time and applies them to any consecutive blocks of selected keys. Useful if you have an object that speeds up and slows down repeatedly, and you want to smooth out the motion.

Constant Velocity—When on, interpolates values between the key and the next one so that the object moves at a constant velocity across that curve segment. Available only with certain controller types, such as Bezier.

Free Handle—Used for automatically updating the length of the tangent handle. When this is turned off, the length of the tangent is at a fixed percentage from its neighboring key. As you move a key, the handles adjust to stay the same percentage away from the neighboring keys. When this is turned on, the handle lengths are based on time lengths.

Animation Controllers

Track View > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Choose a controller in the dialog.

Animation menu > Constraints/Transform Controllers/Position Controllers/Rotation Controllers/Scale Controllers> Choose a controller.

Select an object. > Motion panel > Assign Controller rollout > Select a track. > Assign Controller > Choose a controller in the dialog.

Controllers, like *constraints (page 2–392)*, handle the animation tasks in a scene, including:

- Storing animation key values
- Storing procedural animation settings
- Interpolating between animation key values

An object or parameter doesn't receive a controller until you animate it. As soon as you change an animatable parameter with Auto Key on, or *add a key (page 2–560)* in *Track View - Dope Sheet (page 2–501)*, the software assigns a controller to the parameter. The software chooses a default type for the controller, depending on the animation. You can change the default controller to another type.

The animation controllers are organized in the following categories:

- **Float controllers:** for animating floating-point values

- **Point3 controllers:** for animating three-component values such as colors or 3D points
- **Position controllers:** for animating positions of objects and selection sets
- **Rotation controllers:** for animating rotation of objects and selection sets
- **Scale controllers:** for animating the scale of objects and selection sets
- **Transform controllers:** for animating general transforms (position, rotation, and scale) of objects and selection sets

To change an assigned controller, use *Controller > Assign (page 2–546)* in Track View or in the *Motion panel (page 3–774)*.

The controllers and constraints are:

Attachment Constraint (page 2–393)

Audio Controller (page 2–309)

Bezier Controllers (page 2–310)

Barycentric Morph Controller (page 2–309)

Block Controller (page 2–313)

Boolean Controller (page 2–316)

Color RGB Controller (Point3 XYZ Controller) (page 2–317)

Euler XYZ Rotation Controller (page 2–318)

Expression Controller (page 2–320)

Inverse Kinematics (IK) (page 2–435)

Layer Controller (page 2–325)

Limit Controller (page 2–335)

Linear Controller (page 2–341)

List Controller (page 2–342)

Link Constraint (page 2–403)

LookAt Constraint (page 2–406)

Motion Capture Controller (page 2–347)

Master Point Controller (page 2–346)

Noise Controller (page 2–353)

On/Off Controller (page 2–355)

Orientation Constraint (page 2–409)

Path Constraint (page 2–398)

Position Constraint (page 2–401)

Position XYZ Controller (page 2–356)

PRS Controller (page 2–357)

Reaction Controllers (page 2–358)

Scale XYZ Controller (page 2–371)

Script Controller (page 2–372)

Slave Parameters Dialog (Block Controller) (page 2–391)

Smooth Rotation Controller (page 2–374)

Spring Controller (page 2–375)

Surface Constraint (page 2–396)

TCB Controllers (page 2–377)

Transform Script Controller (page 2–379)

Waveform Controller (page 2–381)

See also

Working with Controllers (page 2–289)

Animation Constraints (page 2–392)

Audio Controller

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Audio

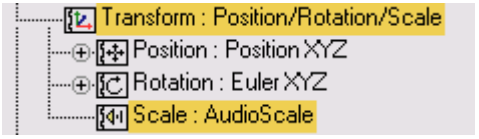
Graph Editors > Track View – Curve Editor > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Audio

Use the Audio controller to drive the animation of almost any parameter in the software. The Audio controller converts the amplitude of a recorded sound file or real-time sound wave into values that can animate an object or parameter.

With the Audio controller, you have full control over sound channel selection, base threshold, oversampling, and parameter range.

The Audio controller works with most parameters in Track View, including:

- Transforms
- Float values
- Point3 values (color)



Audio controller assigned to the scale track of this object

Procedure

To animate the Z axis scale of a box using a sound file:

1. Create a box.
2. In the Track View hierarchy, select the Scale track.
3. On the Track View menu, click Controller > Assign, and choose the AudioScale controller.

   The Audio Controller dialog appears.
4. Select Choose Sound, and select a *.wav* file.
5. In the Base Scale group, set the Z field to **0**.
6. In the Target Scale group, set the Z field to **600**.
7. Close the Audio Controller dialog and play the animation.

   The box scale is animated along the Z axis.

   Tip: If you want to hear the sound as the animation plays, include the same audio file in the *Sound Track (page 2–520)*. (Sound Track in the Track View hierarchy)

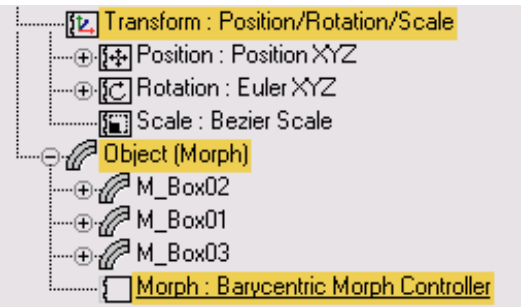
   In the Time Configuration dialog, turn on Real Time.

Interface

See *Audio Controller Dialog (page 2–386)*

Barycentric Morph Controller

Select an object. > Create panel > Geometry > Compound Objects > Morph



The *Barycentric (page 3–914)* Morph controller is automatically applied when a morph object is created in Create > Geometry > Compound Objects > Morph. Morph targets are selected and keys are created at different times to morph the original object into the shape of the Morph Targets.

Note: The Morpher modifier provides an alternative to the Morph controller to morph objects.

See also

Morph Compound Object (page 1–314)

Morpher Modifier (page 1–729)

The Barycentric Morph controller represents each key as a series of weights for all targets. One barycentric key represents a new object which is a blending of all targets.

You can adjust each morph key to various percentages of the available morph targets. This lets you create subtle adjustments in the animation.

You can add keys between existing morph keys. The added keys contain interpolated values for all targets.

To access the key properties dialog, select the Morph track, select one of its keys, and click Properties to display the Key Info dialog.

Procedures

See *Morph Compound Object (page 1–314)* and *Barycentric Morph Controller Key Info Dialog (page 2–388)*

Interface

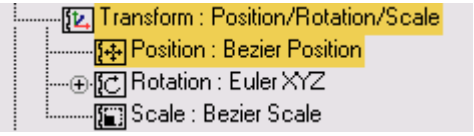
After assigning the Barycentric Morph controller in Create panel > Compound Objects > Morph, then morph parameters for the controller display in the Modify panel and in the Barycentric Controller Key Info dialog, which is displayed by right-clicking over a morph key in Track View — Dope Sheet or the track bar.

See *Morph Compound Object (page 1–314)* for Morph parameters.

Bezier Controllers

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Bezier

Graph Editors > Track View – Curve Editor > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Bezier



The Bezier controller is the most versatile controller available in the program. Bezier controllers interpolate between keys using an adjustable spline curve. They are the default controller for most parameters.

Use Bezier controllers whenever you want fully adjustable interpolation between keys. Bezier is the only controller that supports the following:

- Dragging tangent handles
- Step tangents for abrupt changes from one key to the next
- Constant velocity controls

Procedures

To set the tangent type for a key:

1. Select an object that has some animation keys and open the Key Info (Basic) rollout in the Motion panel.
2. Use the arrows to select a key.
3. Choose a tangent type from the In or Out tangent flyouts.

To normalize time for a key:

1. Select an object, and then choose a key to normalize.

   You can use the arrows in the Key Info (Basic) rollout to step through keyframes.

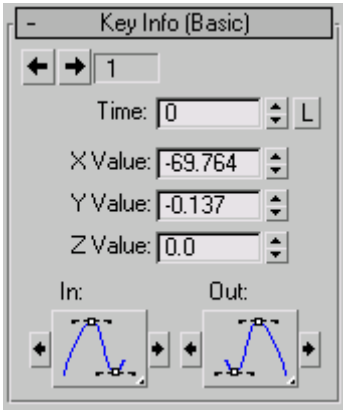
2. On the Motion panel, in the Key Info (Advanced) rollout, click Normalize Time.

   The key is moved in time to average the velocity through the key.

Interface

Motion panel > Parameters > Key Info (Basic) rollout

A Key Info dialog with identical parameters to this rollout is available by right-clicking a key in Track View or the track bar.



Key number—Shows the current key number. Click the right or left arrows to go to the next or previous key.

Time—Specifies when in time the key occurs.

Time Lock—Controls dragging the key horizontally in Track View edit modes.

X/Y/Z—Adjusts the position of the selected object at the current key.

Key Tangent—Two flyouts set the interpolation properties of the in tangent and out tangent of the key.

Choosing Bezier Tangent Types—Sets the tangent types for one or more keys in the same track using the In and Out tangent flyouts.

See *Tangent Types (page 2–305)* for detailed information on each of the available tangent choices.

Tangent Copy—Copies the tangent type between the tangents of the current key or between the tangents of the previous and next key. Use the arrow buttons on either side of the Key Tangent flyouts.

The left arrow of the In tangent copies to the Out tangent of the previous key.

The right arrow of the In tangent copies to the Out tangent of the current key.

The left arrow of the Out tangent copies to the In tangent of the current key.

The right arrow of the Out tangent copies to the In tangent of the next key.

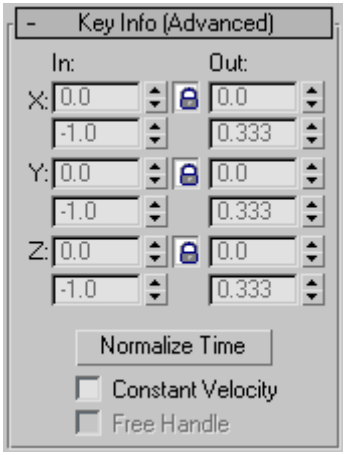
Choosing Bezier Tangent Types—Sets the tangent types for one or more keys in the same track using the In and Out tangent flyouts.

Bezier Scale Controller (Lock X)—Causes the X value to affect all three axes of scale. The Y and Z values are ignored and their function curves are not displayed.



When X is locked, the Y and Z values are not affected by changes in the X value. If you lock X when all three axes are at identical values, alter the X value, and then unlock X, the Y and Z values remain where they were while X retains its new value.

Motion panel > Parameters > Key Info (Advanced) rollout



Controls in the rollout affect velocity in three ways:

- Control the absolute velocity at a key using the In/Out values
- Average velocity over a specified amount of time using Normalize Time
- Force constant velocity from one component key to the next using Constant Velocity

In/Out—The In field displays the rate of change as the parameter approaches the key. The Out field displays the rate of change as the parameter leaves the key.

These fields are active only for keys using the Custom tangent type.

The number in the field is the rate of change expressed as parameter units per Tick. By changing the two values for X, Y, and Z you are changing the length and angle of the tangent handle.

Lock button—When on, changing one Custom tangent changes the other by equal but opposite amount. For example, if the Lock button is on and the In value is 0.85, then the Out value is -0.85.

Normalize Time—Averages the position of the keys in time and is applicable to any consecutive blocks of selected keys. Useful if you want to smooth out the motion and have an object that speeds up, slows down, speeds up, and slows down.

Constant Velocity—Interpolates values between a key and the next one in a way that makes the object move at a constant velocity across that curve segment.

Free Handle—Used for automatically updating the length of the tangent handle. When this is turned off, the length of the tangent is at a fixed percentage from its neighboring key. As you move a key, the handles adjust to stay the same percentage away from the neighboring keys. When this is turned on, the handle lengths are based on time lengths.

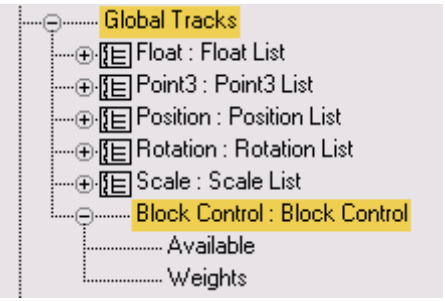
Create Position Lock Key / Create Rotation Lock Key—Makes the incoming handle of the current key linear, and the out going handle of the previous key linear. This is to prevent overshoot from the spline interpolation.

To use these two features, you must first create a keyboard shortcut in Customize > Customize User Interface. Look for Create Position Lock Key and Create Rotation Lock Key in the keyboard shortcut list and assign a key. Or you can also add these commands to the quad menu.

Block Controller

Main toolbar > Curve Editor (Open) > Expand Global Tracks in the Track View hierarchy. > Block Control

Graph Editors > Track View – Curve Editor > Expand Global Tracks in the Track View hierarchy. > Block Control



A Block controller is a global *List controller (page 2–342)* that allows you to combine several tracks from multiple objects over a range of time, and group them as "Blocks." These Blocks are then used to re-create the animation anywhere in time. Blocks can be added, removed, scaled, moved graphically in Track View, and saved. Blocks can represent either absolute or relative animation.

For example to animate a hand forming a guitar chord, all the rotations of the fingers and hand can be saved as a block. This block can then be used to recreate the hand and finger position, rotation, and animation whenever the chord is played in an animation.

Essentially, block controllers allow you to build up libraries of animations and apply them to objects as you choose.

Tip: Euler Rotations will give you better results than TCB Rotation controllers when you are working with blocks. Block controllers blend clips in a relative repeat manner. If your motions do not loop exactly, your rotations start to drift. Also note that block controllers only work with Keyframe controllers (non-procedural).

Master Block Parameters Dialog

After assigning a Master Block, right-click the MasterBlock track to display the *Master Block Parameters dialog (page 2–390)*. This is the first step in creating a block.

With this dialog, you can save blocks and then load them later. Blocks are saved as BLK files.

Track View Pick Dialog

The *Track View Pick dialog (page 2–392)* is displayed when track selection is necessary; for example, when you click Add in the Master Block Parameters dialog. Select tracks in the dialog to include in a block. Valid tracks are shown as darker.

Block Parameters Dialog

After you create a block by clicking Add in the Master Block Parameters dialog and then selecting tracks in the Track View Pick dialog, clicking OK in the Track View Pick dialog displays the *Block Parameters dialog (page 2–389)*.

Attach Controls Dialog (Loading Blocks)

The *Attach Controls dialog (page 2–388)* is displayed when you click Load in the Master Block Parameters dialog. Tracks in blocks saved previously can be mapped to tracks in the current scene.

Slave Parameters Dialog (Slave Controller)

Every time a block is created, all tracks within the block are assigned a slave controller, which allows the MasterBlock to transfer key data. The slave controller tracks appear with the original tracks you used to create the block. See *Slave Parameters dialog (page 2–391)*.

Procedures

Example: To create a block:

1. Animate two objects moving in the viewports. Make the final keyframe for the objects frame 10.
2. Open Track View — Dope Sheet.
3. On the Track View hierarchy, expand Global Tracks, expand Block Control, and then select Available.

   Available is the track below the Block Control track.
4. From the Track View toolbar, choose Controller menu > Assign, select Master Block in the dialog, and then click OK.

   The Master Block Parameters dialog automatically displays.
5. Click Add on the Master Block Parameters dialog.
6. On the Track View Pick dialog, expand the tracks for the two animated objects.
7. While holding down Ctrl, click the X,Y,Z Position tracks for both objects, and then click OK.
8. On the Block Parameters dialog, type a name in the name field.

   Choose a name that will remind you of the animation in this block.
9. Set the End value to **10**, and then click OK.

   The Block will contain animation between frame 0 and frame 10.
10. Click OK to close the Master Block Parameters dialog.

    The block has been created, and you can now use it.

    - Once a block is added, a slave controller is added to the controllers on the original tracks. This enables communication between the track and the Block controller.
    - With the Blend track (below the MasterBlock track in the Track View hierarchy) you can animate how much of the block animation will be in effect. Negative values reverse the animation.
    - Controllers that each block uses are listed under the block name. This allows you to adjust the data for a particular block.

Example continued: To use a block:

1. In the Dope Sheet Key window, right-click the MasterBlock track.

   A pop-up menu displays the name of blocks already created.
2. Select the name of the block you created earlier.

   The Block is displayed in the Key window.
3. Drag the Block to start at frame 20, and click Play.

   The animation repeats at frame 20.

   You can create a number of blocks for different periods of an object's animation, and use them at different locations in the MasterBlock track.

To move an inserted block:

- In the Key window, select the block and then drag it left or right.

To scale (resize) an inserted block:

- Select the key at the lower left or right corner of the block, and then drag left or right.

To create a copy of a block:

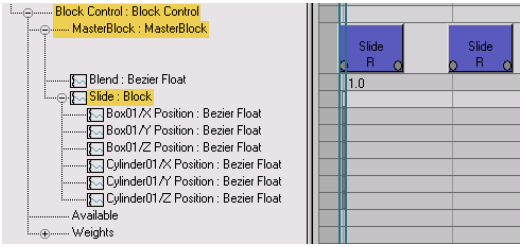
- Hold down Shift, and draft the block.

  This creates a new instance of the block, which you can place at a different time.

To remove an inserted block:

- Click to select the block, then press Delete.

Interface



Inserted blocks appear in the Key window to the right of the MasterBlock controller in the hierarchy under Global Tracks.

Block Control track—Displays in Track View under Global Tracks.

To create a MasterBlock Control, expand Block Control, select Available, and choose Controller > Assign.

MasterBlock Track—Displays in Track View after you assign a Master Block controller to an available track.

You can right-click the track in the Key window. This displays a pop-up menu that shows the name of blocks that have been created. Choosing a block inserts the block into the MasterBlock track. The dialog also has an entry, Properties, that displays the Master Block Parameters dialog.

In the MasterBlock track, inserted blocks display as colored rectangles. The name of the block appears at the center. In the lower left and right corners are keys that indicate the beginning and end of the block's animation. Below the name of the block is the letter "R" for relative, or "A" for absolute. Click and drag the center of a block to move it in time. Select a key in the corner of a Block to move one edge of the block (scale time).

Right-click an inserted lock to display the Block Key Properties dialog (described below).

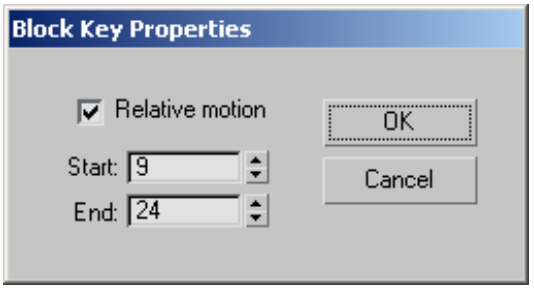
MasterBlock Subtracks

Below the main MasterBlock track are subtracks. The first of these is always Blend. The remaining subtracks are initially copies of the tracks used to create the block.

Blend track—Lets you animate the influence of the block. You can reduce the block's influence by creating Blend keys with values less than 1.0. Default=1.0.

Block-specific subtracks—Display the block name and its associated tracks. Initially these are copies of the tracks used to create the block. You can edit keys on these tracks to change the block's behavior.

Block Key Properties Dialog



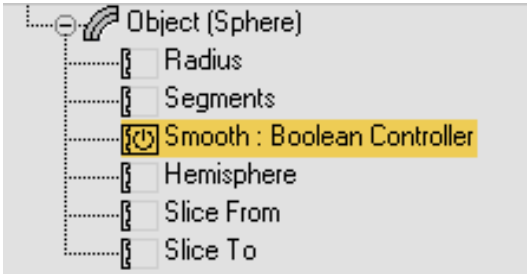
Relative Motion—Toggles between relative and absolute motion.

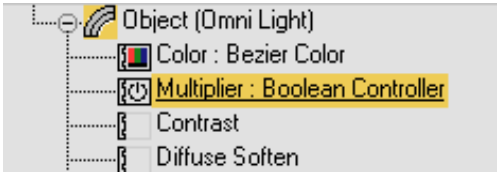
Start, End—Set the first and last frames for this instance of the block (this scales the block instance).

Boolean Controller

Main toolbar > Curve Editor (Open) > Select a track containing a float value in the Track View hierarchy > Track View menu bar > Controller menu > Assign > Boolean controller

Graph Editors > Track View – Dope Sheet > Select a track containing a float value in the Track View hierarchy > Track View menu bar > Controller menu > Assign > Boolean controller





The Boolean controller is similar to the *On/Off controller (page 2–355)*. By default, it is assigned to tracks (such as an object's Visibility track) that provide only binary on and off control.

Note: The Boolean controller is useful for controlling the *Enabled* state in the *History Independent (HI) IK system (page 2–453)*.

It differs from the On/Off controller in two ways: first, each key has its own float value, either *1.0 or 0.0*, signifying its on or off state. This enhancement leads to the second component that differentiates the Boolean controller from its On/Off cousin: the ability to lay down sequential keys without inadvertently changing the on/off state of any keys downstream.

The Boolean controller can be assigned to any parameter that would normally be controlled by a float or boolean class controller, such as sphere's *Hemisphere* or *Smooth* track.

Boolean controller key values can be changed in on of two ways: through Track View's Dope Sheet editor or through MAXScript.

Note: Although the Boolean controller displays a function curve in Track View's Curve Editor, it does not display keys. Key time and value must be changed in the Dope Sheet editor.

Procedures

To assign a Boolean controller and add keys:

1. Open Track View's Dope Sheet editor and select any track that would normally be assigned a float or boolean class controller.
2. Right-click and select Assign Controller. Pick the Boolean controller from the Assign Controller dialog, and click OK
3.  From the main Track View toolbar, select the Add Keys button and click anywhere on the selected track.

   New keys with a value of 0.0 are added.

To change Boolean controller key values in Track View:

1. Open the Track View Dope Sheet editor and select a track that contains a Boolean controller.



2. Select any key and type **1.0** or **0.0** into the key value field at the bottom of the Track View dialog.



Note: Values greater than 0.0 will automatically be clamped to 1.0

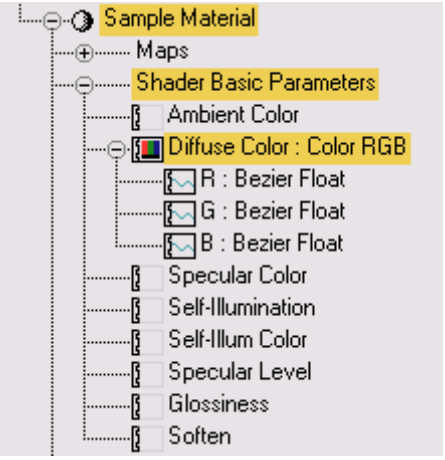
To change Boolean controller key values via MAXScript:

- Enter the following into either the MAXScript Listener or Mini Listener:**<node>.<animatable_property>.keys[<index_integer>].value = <float>** where *<node>* is the object containing the Boolean-controlled parameter (for example, *$Sphere01*, *<animatable_property>* is the parameter itself (for example, *Smooth*, *<index_integer>* is the actual sequential number of the key in the key array, and *<float>* is either 1.0 or 0.0.

Color RGB Controller (Point3 XYZ Controller)

Main toolbar > Curve Editor (Open) > Select a material color track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Color RGB

Graph Editors > Track View – Curve Editor > Select a material color track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Color RGB



The Color RGB controller splits the R, G, and B components into three separate tracks. You can use this controller with color tracks. By default, each track is assigned a Bezier Float controller.

The Bezier Float controller is a single parameter controller.

Note: The Color RGB and the Point3 XYZ controllers are identical in function, except that the labels of their tracks are RGB and XYZ respectively. Point3 XYZ controllers can be used for functions such as Noise Strength, and the FFD modifier.

Procedure

Example: To animate the background color of a rendered animation:

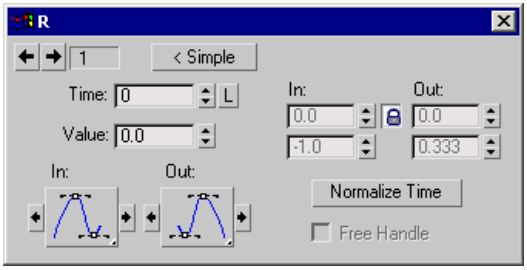
1. Open Track View – Curve Editor.
2. In the Track View hierarchy, expand the Environment track and then select Background Color.
3. On the Track View menu, click Controller > Assign and select the Color RGB controller.
4. Expand the Background Color track.
5.  On the Track View toolbar, click Add Keys.
6. Add three keys along the red track at frames 0, 50, 100.
7. Right-click one of the keys to display the Bezier Float dialog. Change the key values to **0**, **400**, and **0**, respectively.

   To see the background color change, move the time slider to the same frame as the key that is being adjusted and render the scene. To see the color change in an animated fashion, you must render the animation.

   Although you can set values over 255 in the value field, the actual color value stops at 255.

Interface

Bezier Float controller properties.



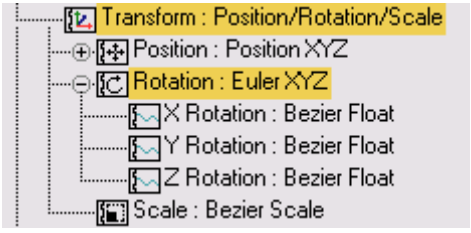
See *Bezier Controller (page 2–310)* for a description of these parameters.

Euler XYZ Rotation Controller

Main toolbar > Curve Editor (Open) > Select a rotation track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Euler XYZ

Graph Editors > Track View – Curve Editor > Select a rotation track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Euler XYZ

Select an object in the viewport > Motion Panel > Assign Controller rollout > Select the rotation track. Euler XYZ is assigned as the default controller.



The Euler XYZ Rotation controller is a compound controller that combines separate, single-value float controllers to specify an angle of rotation about each of the X, Y, and Z axes. Euler XYZ is not as smooth as quaternion rotation (used by the TCB Rotation controller), but it is the only rotation type that you can use to edit rotation function curves.

The Euler XYZ Rotation controller assigns an individual track to the X, Y, and Z components of the position, rotation, or scale of an object. However, the controller assigns three keys (one for each axis), by default.

There are actions available from the *Customize User Interface dialog (page 3–792)* that let you create explicit axis keys.

The Gimbal *reference coordinate system (page 1–443)* is meant for use with this controller. With other coordinate systems, rotating about one axis always changes at least two tracks. With Gimbal rotation, rotating about one axis changes only one track, making it easier to edit function curves.

Euler versus TCB Rotation

Euler rotation offers several advantages over the TCB Rotation controller, which was the default rotation controller in versions 4 and earlier. Euler rotations provide function curves, while TCB do not. Euler rotations allow for 3 separate curves that can be manipulated in the Curve Editor. TCB rotations can only be controlled using tension continuity and bias settings, making it difficult to keyframe or manipulate X, Y or Z rotations independently.

On the other hand, TCB rotations allow for rotations greater than 180 degrees on a given keyframe. If you turn on the Rotation Windup function, you can get rotation values greater than 360 degrees using TCB.

Tip: Users who are used to working with TCB rotations might find difficulty in adapting to Euler workflow. For these users, it is advised that you reassign TCB as the default rotation controller. See *Specifying Default Controllers (page 2–294)*.

Euler Rotation and the Waveform Float Controller

Because the Euler Rotation controller uses radians, unit adjustments should be made when other controllers are applied to Euler axes. For example, the Waveform Float controller has a default amplitude of 100 in the Characteristic Graph. When the Waveform Float controller is applied to an Euler axis, the default amplitude is 100 radians.

This sets the Amplitude setting to 5729.598 (the number of degrees in 100 radians).

Euler Rotation and the Noise Float Controller

When a Noise Float controller is applied to an Euler axis, the default Strength setting is 286.479 or 50 percent of 10 radians in degrees (maximum deflection).

Euler Rotation and the MIDI Motion Capture Controller

When a MIDI Motion Capture controller is applied, the Parameter Scaling is taken in radians so that the Max. default setting of 1.0 results in an upper boundary of 57.2958 degrees.

Procedures

Example: To use Euler XYZ Rotation:

1. Create a box, and then select the rotation track in Motion panel > Parameters > Assign Controller rollout.
2. Click Assign Controller, and then select Euler XYZ in the Assign Controller dialog list. Click OK.
3. Turn on Auto Key.
4. Select Rotation at the bottom of the PRS Parameters rollout.
5. In the Euler Parameters rollout, select the X axis.
6. In the Create Key group of the PRS Parameters rollout, click Rotation.

   A rotational key is created.
7. Move the time slider to frame **50**.
8. Again, in the Create Key group of the PRS Parameters rollout, click Rotation.
9. In the Key Info (Basic) rollout, enter **500** in the value field.
10. Select the Z axis in the Euler Parameters rollout, and then enter **90** in the value field of the Key Info (Basic) rollout.

    Play the animation. The box rotates 500 degrees around the X axis and 90 degrees around the Z axis. In this case the X axis is rotated first, then Z.

To use Euler XYZ with List controllers for local rotation control:

You can mimic the local Euler rotation controller (available in earlier versions of the software) by combining a list controller with an Euler XYZ rotation controller.

1. Select the object for which you want to have local rotational control. For example, create a teapot and a dummy.
2. Animate the dummy.
3. Open the Motion Panel.
4. In the Assign Controller rollout, select the Rotation transform, then assign a LookAt controller to teapot, with the dummy assigned as the LookAt Target.

   You now have an object with rotation controlled through the dummy.
5. With the Rotation: LookAt controller selected in the window, again choose Assign Controller and pick Rotation List.

   Now there is a list controller, with the LookAt Constraint applied as the first rotation controller in the list.
6. Scroll down in the window and select the entry labelled Available, then again click Assign Controller.
7. Choose Euler XYZ.

   You now have a list controller with a LookAt Rotation as the first controller and an Euler XYZ as the second controller.

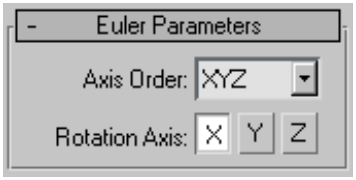
8. In the Rotation List rollout, highlight Euler XYZ in the Layers window, then click Set Active.
9. In the Euler Parameters rollout, change the Axis Order to ZYX.

   Now you can animate the rotation independent of the LookAt rotation control. It should now behave the same as the Local Euler controller that was available in versions 4 and earlier.

Interface

Euler parameters are displayed in the Motion panel.

The Euler parameters rollout is displayed when Rotation is selected in the *PRS Parameters rollout (page 2–357)*.



Axis Order—Selects the order that rotations are calculated. The default is X,Y,Z order, where the X axis is rotated first.

X—Displays controller properties for X axis rotation angle.

Y—Displays controller properties for Y axis rotation angle.

Z—Displays controller properties for Z axis rotation angle.

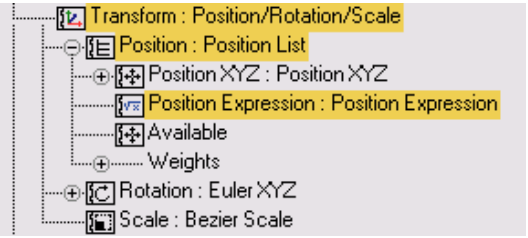
Each axis uses its own independent controller using the float data type. For example, the X and Y Rotation axes could use Bezier Float controllers, while the Z rotation axis uses a Noise Float controller.

Expression Controller

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Expression

Graph Editors > Track View – Curve Editor > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Expression

Animation menu > Position or Scale controller > Expression



Interface (page 2–322) Procedures (page 2–321)

The Expression controller lets you use mathematical expressions to control the following animation aspects:

- Object parameters such as length, width, and height
- Transform and modifier values such as an object's position coordinates

You can constrain values by basing them on the controller values of other objects in the scene.

An expression is a mathematical function that returns a value. 3ds Max evaluates the expression once for each frame of an animation, generating values that can change from frame to frame.

You can create and manage Expression controllers in *Track View (page 2–501)* and on the *Motion panel (page 3–774)*. You can assign expressions to the following kinds of scene elements:

Scene element	Controller
Creation parameters	Any numeric creation parameter
Transforms	Position [X, Y, Z]
	X Rotation
	Y Rotation
	Z Rotation
	Scale [X%, Y%, Z%]
Modifiers	Any numeric modifier parameter (including creation parameters)
Materials	Colors [R, G, B]
	Any numeric material parameter

Note: Expression controllers can work only with the individual XYZ components of Euler rotation. You can't assign an expression to TCB Rotation or other kinds of rotation controllers.

See also

Expression Techniques (page 1–146)

Expression Controller Techniques (page 2–324)

Trigonometric Functions (page 1–150)

Vectors (page 1–151)

Procedures

To assign a constant value to a variable:

1. Highlight the variable name in the Scalars or Vectors list.
2. Click Assign To Constant.

   A new dialog opens.
3. On the dialog, enter the new value (or, in the case of a vector, values) for the constant, and then click OK.

To assign a controller to a variable:

1. Highlight the variable name in the Scalars or Vectors list.
2. Click Assign To Controller.

   The Track View Pick subdialog is displayed, showing the track hierarchy. The dialog display is similar to the Track View hierarchy.
3. Highlight the track for the variable to use, and then click OK.

Example: To create an expression that moves a sphere in a precise circle:

1. Create a sphere with Radius=**15.0**.

   You'll use Track View to create the Expression controller.
2. In the active viewport, right-click the sphere and click Curve Editor.
3. In the Hierarchy list, scroll down to the Objects branch and, if necessary, expand Sphere01 the branch so the sphere's Position track is visible. Click the Position label to highlight it.
4. In the Hierarchy list, right-click the Position label and then choose Assign Controller

   The Assign Controller dialog opens.
5. Choose Position Expression from the list of controller types, and then click OK.

   The Expression Controller dialog opens.

   Note: You can also open the Expression Controller dialog with Track View > Controller menu > Properties.
6. Replace the default expression by typing the following position expression in the Expression field:

   [100*cos(360*NT), 100*sin(360*NT), 0]

   The expression specifies a circular path for the sphere. NT is a variable that means normalized time. Movement based on NT happens exactly once per the active time segment, regardless of how many frames are in the animation.
7. Click Evaluate.

8. Play the animation. The sphere moves in a circle about the world origin (0,0,0). The radius of the circular path is 100 units.

Example continued: To change the radius of the circle:

The two 100s in the position expression from the previous procedure specify the radius. To adjust the radius of the circle's path, create a symbolic variable to represent the radius. The variable has a constant value that is easy to edit.

1. Reopen Track View and the Expression Controller dialog if necessary.
2. In the Name field of the Expression Controller dialog, type **radius** Make sure Scalar is chosen, and then click Create.

   The variable name "radius" appears in the Scalars list of the dialog.
3. Click Assign To Constant.

   A dialog titled "radius" opens.
4. Enter **150** in the Value field, and then click OK.

   The radius variable is now 150.

   Next you'll use the new variable in the expression.

Example continued: To replace the literal value with the variable name:

1. In the Expression field, change 100 to **radius** in both places. The expression should now look like this:

   [radius*cos(360*NT), radius*sin(360*NT), 0]
2. Click Evaluate.
3. Play the animation. The sphere moves in a circle about the world origin (0,0,0). The radius of the circular path is 150 units.

Example continued: To make the sphere rotate about a box:

1. Create a box about 40 units square, and animate its position over three or four keyframes.
2. Select the sphere.
3. In the Name field of the Expression Controller dialog, enter **boxposn**. Choose Vector, and then click Create.

   The name "boxposn" is displayed in the Vectors list in the lower-left area of the dialog.

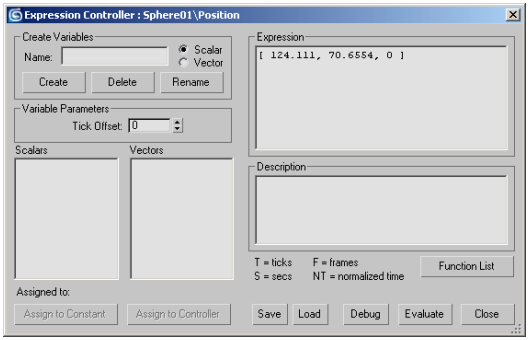
   Variable names are case sensitive; the variable name should be lower case.
4. Click Assign to Controller.

   The Track View Pick dialog is displayed. It shows the object hierarchy as it appears in the left side of Track View-Dope Sheet.
5. In the Hierarchy list, highlight the Position controller for Box01 and then click OK.
6. In the Expression field, add boxposn as an offset:

   [radius * cos(360*NT), radius * sin(360*NT), 0]**+boxposn**.
7. Click Evaluate, and then click Close.

   Play the animation again. The sphere moves in a circle around the box and follows the box wherever it moves.

Interface



Tip: You can resize the dialog by dragging an edge or a corner.

Create Variables group

Name—The variable name.

Scalar/Vector—Choose the type of variable to create.

Create—Creates the variable and adds it to the appropriate list.

You must enter a name and specify a type before clicking Create.

Delete—Deletes the highlighted variable in the Scalar or Vector list.

Rename—Renames the highlighted variable in the Scalar or Vector list.

First highlight the variable in the list; this places the name in the Name field. Edit the name in the Name field, and then click Rename; the new name replaces the previous one in the list.

Variable Parameters group

Tick Offset—Contains an offset value.

A tick is 1/4800 of a second. If a variable has a non-zero tick offset, that value is added to the current time.

Expression group

Type an expression in the Expression box.

The expression must be a valid mathematical expression. The result is either a three-value vector for a vector expression (position, scale, or point3) or a scalar value for a float expression.

Description group

Type text in this group to document an expression. For example, you can describe user-defined variables.

Function List—Displays a list of Expression controller functions.

In the list, p, q, and r represent scalar values or scalar expressions; V and W represent vector values or vector expressions.

Save—Saves an expression. Expressions are saved as files with a *..xpr* file name extension.

Load—Loads an expression.

A saved expression does not include variable definitions or values. After loading the expression, you need to redefine them.

Debug—Displays the Expression Debug window.

This window shows the values of all variables, and the value of the expression. When you change the variables or move the time slider, the Debug window automatically updates so you can interactively view what's happening with the expression. The values for frames (F), normalized time (NT), secs (S), and ticks (T) are also displayed.

Evaluate—Evaluate the expression for each frame in the animation.

There is no explicit assignment (= or := operator) as in a conventional programming language; the assignment is implicit and takes place over time.

If the expression has a syntax error, an error message is displayed. The error message is the first part of the expression itself. The last character in the error message is the point of the error. This is usually where the error actually is, unless the problem is that opening and closing parentheses (or the braces for vectors) don't match. In this case, evaluation can proceed further before the error is detected.

Close—Closes the Expression Controller dialog.

Expression Controller Techniques

This topic summarizes some useful expression techniques. The following topics review trigonometric functions and vector arithmetic. If you're familiar with these subjects you can skip the review topics. Use these expressions when you have applied an Expression controller using the Assign Controller rollout in the Motion panel.

Commonly Used Expressions

This topic lists some expressions that you might find useful in situations when you animate.

Circular Path

```
[ Radius * cos(360*Time),
Radius * sin(360*Time), 0 ]
```

where *Time* is one of the predefined time variables such as NT or S.

If you make the two Radius values unequal, you get an elliptical path.

If you specify a nonzero Z component, the path is no longer planar.

Following Another Object

```
[X, Y, Z] + Position
```

where *Position* is the Position controller of the second object.

The vector [X, Y, Z] can be an offset from the second object. (If it's [0,0,0], the two objects occupy the same position.) It can also be a vector expression that specifies some movement in itself.

Keeping an Object Between Two Objects

```
(Position1 + Position2) / 2
```

where *Position1* and *Position2* are the Position controllers of two objects.

The divisor 2 constrains the object to be halfway between the two other objects. Other values constrain the object to other locations.

Bouncing Between Other Objects

```
(1+sin(360*Time))/2 * (Pos1-Pos2) + Pos2
```

where *Time* is one of the predefined time variables such as NT or S; *Pos1* and *Pos2* are the Position controllers of two other objects.

The subexpression (1+sin(360*Time))/2 is a value that oscillates between 0 and 1 over time. (Pos1-Pos2) is the vector between the two other objects. Multiplying the two and then adding Pos2 as an offset locates the object along this vector.

Changing the Number of an Object's Segments Based on Camera Distance

This expression varies the number of segments in a cylinder based on the distance of a camera. It is assigned to the cylinder's Segments creation parameter.

```
if ( (length(Camera-Myself) > 35),
3 + (50*Height) / length(Camera-Myself),
MaxSegs)
```

where *Camera* is the position controller of the camera; *Myself* is the cylinder's position controller; *Height* (= 70) is the cylinder's height; *MaxSegs* (=100) is the maximum number of segments.

When the camera is closer, more segments make the cylinder smoother; when the camera is distant, the smoothing is less important and fewer segments render more quickly.

The if() function returns its second argument if the first argument is true; otherwise, it returns its third argument. In this example, if the camera is more than 35 units away from the cylinder, the expression calculates the number of segments; if the camera is 35 units away or closer, the number of segments is the MaxSegs constant.

The values in the second argument are chosen so that as the distance decreases toward the threshold of 35, the number of segments increases toward MaxSegs. The addition "3+" ensures that the cylinder always has at least three segments, even when the division rounds to zero (Segments is an integer).

Note: To the expression, it doesn't matter whether the camera is moving, or the cylinder, or both.

Layer Controller

Main toolbar > Curve Editor (open) > Highlight a Layer controller track and right-click > Properties

Select layer-enabled object. > Motion Panel > Parameters > Position/Rotation/Scale

The Layer Controller dialog provides commands and options related to the Layer controllers in your scene, which the system automatically assigns for you when you enable *animation layers (page 2–326)* on an object.

Unlike other controllers, you cannot assign a Layer controller explicitly to a track; you first need to enable layers via the *Animation Layers toolbar (page 3–690)* or the Controller menu of the Curve Editor.

The Layer controller dialog has similarities to the *List controller (page 2–342)* dialog. You can refer to it for complementary informations on some of the options.

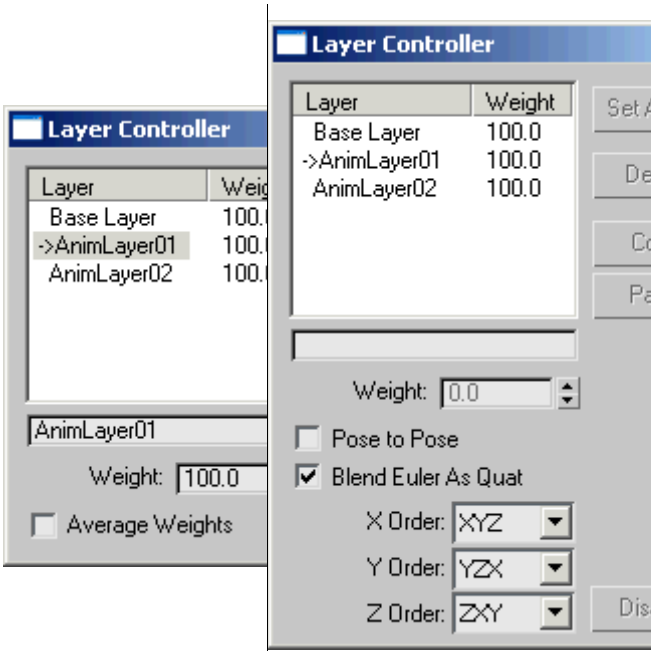
Animation Layers and Autodesk VIZ

Although the Animation Layers feature is not available in Autodesk VIZ, VIZ can load and save objects that use this feature without any data loss. The Layer controllers continue to exist in VIZ, and can even be viewed and edited in Track View, although editing of the animation-layer data is not recommended if the objects are to be returned to 3ds Max.

Tip: For best results with objects that are to be brought into Autodesk VIZ, *collapse* any animation layers before saving the objects.

Interface

There are two types of dialogs based on whether you assign a Layer controller to a position or scale track, or to a rotation track.



List Window—Displays all Layer controllers for the selected object, along with their respective weight value.

Set Active—Determines on which layer your animation keys are set. The active controller is marked with an arrow in the list.

Tip: You can also switch between active layers from the drop-down list on the *Animation Layers (page 2–326)* toolbar.

Delete—Deletes the highlighted controller. A confirmation dialog prompts you before.

Copy—Copies the highlighted controller's data and enables Paste.

Paste—Puts the copied content onto the highlighted controller.

Weight—Sets the effects of the highlighted Layer controller

Average Weights—When on, the weight values of all the controllers in the list. Available only when you assign a Layer controller to a position track. Default=off.

Pose to Pose—Enables blending among controllers in the list. Available only when you assign a Layer controller to a rotation track. Default=off.

Refer to the *List controller (page 2–342)* for more information on this option.

Blend Euler As Quat—When on, exposes the rotation axis order for blending the Euler controllers, which can prove useful for controlling *gimbal (page 1–443)*. Available only when you assign a Layer controller to a rotation track. Default=off.

X/Y/Z Order—Sets the order in which the system calculates each rotation axis. Available only when Blend Euler As Quat is on.

Disable—Removes the Layer controller from the select object and reverts the animation keys on the Base Layer to the original controller.

Note: You have to delete or collapse (available on the *Animation Layers (page 2–326)* toolbar) all layers above the Base Layer before you can disable it.

Animation Layers

Right-click unused area of any toolbar. > Animation Layers

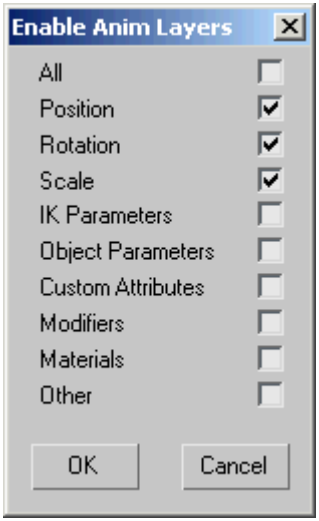
Animation menu > Animation Layers

Animation layers let you combine multiple animation tracks on the same object. You can use layers to store your animation experiments as you try them out, turning them on or off to your liking. To enable, add, and control layers, use commands from the Animation Layers toolbar.

Using animation layers is comparable to using both the *List Controller (page 2–342)* and *biped layer system (page 2–974)*, while it is more flexible and simpler for the animators to use.

Enabling Animation Layers

When you *enable animation layers* for a selected object by turning them on in the Enable Anim Layers dialog, 3ds Max assigns a *Layer controller (page 2–325)* to the animation tracks.

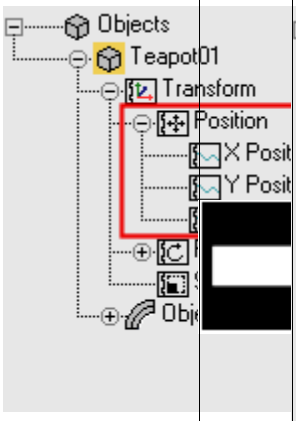


When 3ds Max adds a Layer controller, it copies the original controller into its first layer (called

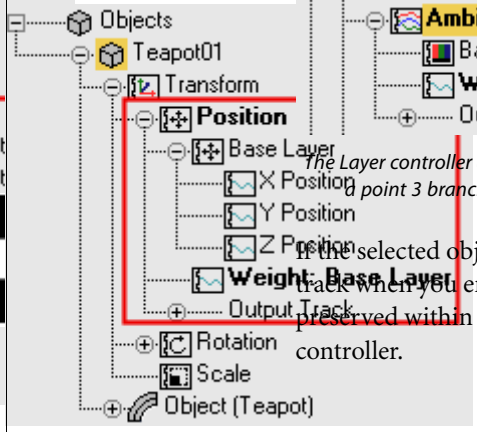
Base Layer), preserving any animation data it may already contain.

Note: The Base Layer is not a new layer; it is on the same level as any other controller track.

Note: If you layer-enable a controller but do not see a Layer controller appear in the Track View hierarchy window, open and close the Enable Anim Layers dialog again.



The original controller before you enable animation layers



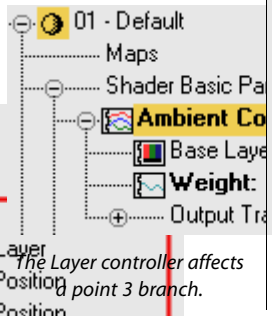
The Layer controller nests the original controller in its Base Layer.

To revert from a Layer controller to the original controller, click *Disable Anim Layer* on the Animation Layers toolbar. However, if your object contains more than one animation layer, you first need to either delete or collapse them before you can disable the Base Layer.

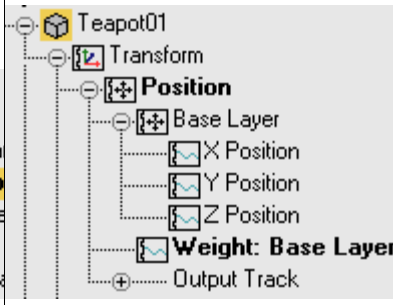
Tip: You can also disable a Layer controller through the *Layer Controller dialog (page 2–325)*.

Note: Turning off a track in the Enable Anim Layers dialog does not disable it.

You can assign a Layer controller either at the leaf level of a controller track (such as a material's Diffuse Color track), or at a branch level (such as a Position track).

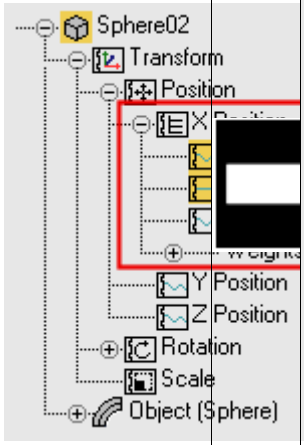


The Layer controller affects a point 3 branch.

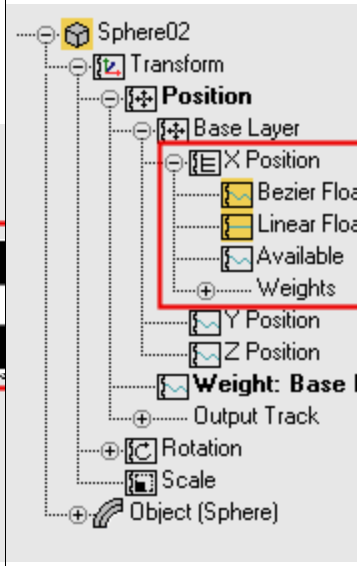


The Layer controller affects a leaf track.

If the selected object already has a list controller track when you enable layers, that list controller is preserved within the Base Layer of the new Layer controller.



The original List controller



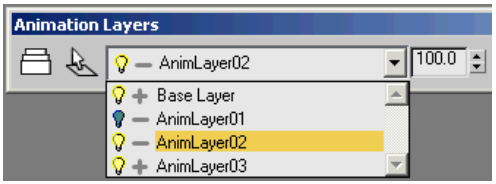
The new Layer controller nests the List Controller.

Note: You cannot nest a Layer controller within a List controller or another Layer controller.

Working with Animation Layers

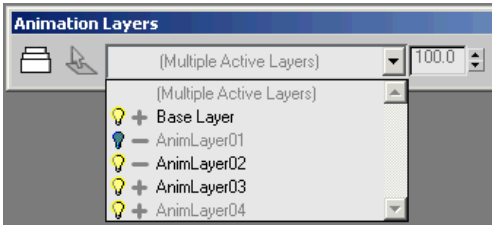
The animation layers list works as follows (Refer to the *procedure section below (page 2–329)* for common workflow examples):

- The list displays all existing layers for the selected object. If the object isn't layer-enabled, "(Enable Layers)" appears instead.
- By default, the Base Layer is the active layer when you first layer-enable an object. This means that it stores all future animation keys.
- Every subsequent layer you create becomes the new active layer and appears after the previous one in the list.
- When you select an object, its active layer is automatically chosen. When expanded, the drop-down list highlights the active layer.



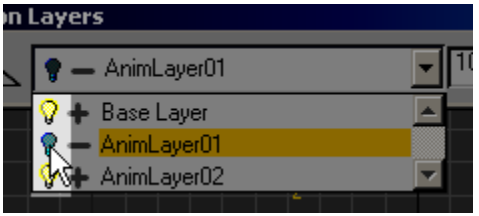
- When you select multiple objects that have different active layers, "(Multiple Active Layers)" appears.

  The drop-down list displays all common layers (that is, with the same name); layers not common to all are unavailable.



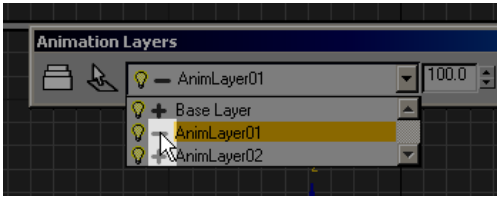
When you select multiple objects, layers not common to all are unavailable on the drop-down list.

- You can turn on or off any layer in the expanded list by clicking its light bulb icon. Turning off a layer hides its animation without deleting it.



To toggle a layer, click its light bulb icon.

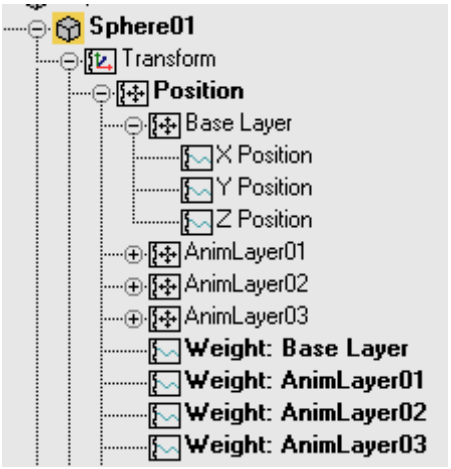
- You can include or exclude any layer from the output track by clicking its plus/minus sign icon. The output track contains the sum of all included layers. You can use this track to control other objects' tracks via MAXScript and parameter wiring.



To toggle a layer's inclusion in the output track, click the plus/minus sign icon.

Tip: For a procedure that demonstrates how to use this option, see *Example: To link two objects with wire parameters using the Layer Controller's output track: (page 2–330)*.

Each animation layer has a global weight value which, when changed, impacts every controller within that active layer. Similar to the *List controller (page 2–342)*, you can animate a layer's weight and toggle the visibility of the resulting keys in the Track Bar using the *Track Bar Filter menu (page 3–706)*.



When you *add a new animation layer (page 2–334)*, you can pick a name already in use by an existing layer; this links both layers' weight track, which is now instanced.

Using Merged or XRef Scenes Containing Layers

When you *merge (page 3–463)* or *externally reference (page 3–407)* a source scene (or object) containing layers into your master scene, the system adds the incoming layers to the master scene's layer list.

Note: All animations within an incoming XRef object or scene is preserved within an *XRef controller (page 2–383)*. You cannot edit them unless you merge the XRef object or XRef controller into your master scene.

If a layer name from the merged or XRef scene matches one from the master scene, both layers become synchronized ; that is, the weight track is instanced to both of them.

Collapsing Layers

Collapsing a layer merges keys if any of these conditions are met:

- the two controllers are of similar types (Bezier, Float, TCB, etc.).
- they have the same tangent type.
- the Blend Euler As Quat option of the *Layer Controller dialog (page 2–325)* is turned off.

However, many factors can cause a collapse per frame (that is, setting a key on every frame):

- If one controller type is TCB (Quaternion) and the other is Bezier (Euler).
- If the Blend Euler As Quat option of the *Layer Controller dialog (page 2–325)* is turned on, and both rotation controllers are Euler-driven.
- If the Blend Euler As Quat option of the Layer controller is turned off, and both rotation controllers are Euler-driven, but both have different tangent types.
- If either controller is non-keyable.

Procedures: Integrating Animation Layers in a Worklow

The following sequence of procedures illustrate how you can use animation layers in your workflow to quickly create and combine different animation tracks together for the same object. You should be familiar with the basics of animation and key creation before going through the following steps.

Example: To enable layers:

1. Create a teapot, and then turn on Auto Key and set position keys on the X axis at frames 0, 15, and 30.
2. Open the *Curve Editor (page 2–507)* and locate the teapot's Position track in the hierarchy window.
3. Choose Animation > Animation Layers to open the Animation Layers toolbar.
4.  Make sure the teapot is selected, then click Enable Anim Layers.

5. On the Enable Anim Layers dialog, turn on only the Position track, and click OK.

   The new Layer controller now nests your teapot's animation.

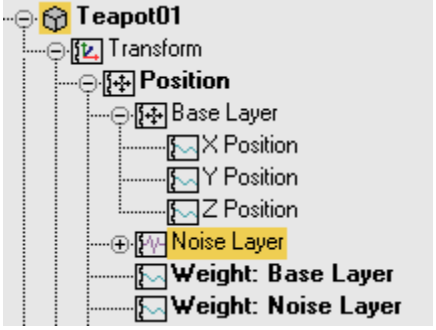
Example: To add a layer and change its controller type:

This continues from the previous procedure, in which you enabled layers for an animated object.

1.  Click Add Anim Layer on the Animation Layers toolbar.
2. On the Create New Animation Layer dialog, enter "Noise Layer" as name, and click OK.

   This adds the new animation layer under the Base Layer in the hierarchy list.
3. In the hierarchy window of the *Curve Editor (page 2–507)*, Highlight *Noise Layer* and right-click it.
4. From the contextual menu, choose Assign Controller and pick a Noise Position controller.
5. This opens the Noise Controller dialog. Leave the options as they are and close the dialog.

   Two different controllers are now driving the teapot's position.



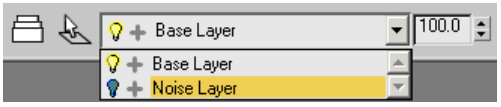
Example: To turn a layer on and off:

This procedure discusses turning existing layers on and off. It continues from the previous procedure.

1. On the Animation Layers toolbar, expand the drop-down list.

   The list currently contains both the original base layer and *Noise Layer* from the previous procedure.
2. Move your cursor across the list and click the light bulb icon next to *Noise Layer*.

   This turns off the layer and hides the Noise controller track from the rest of the tracks.

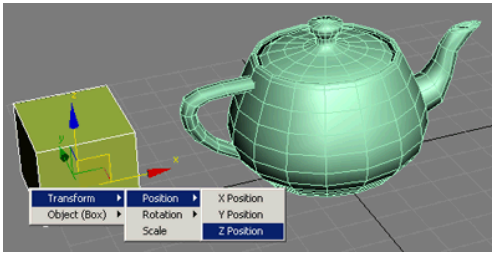


3. Drag the time slider and notice how the teapot does not shake anymore.
4. Try turning on and off both layers in the list to see the results.

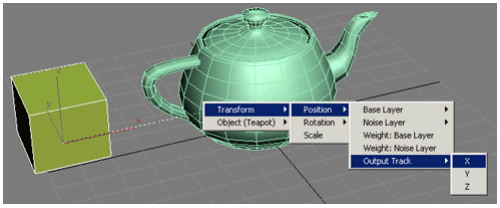
Example: To link two objects with wire parameters using the Layer Controller's output track:

This procedure refers to a simple way to use the output track to control a cube's position through wire parameters. It continues from the previous procedure.

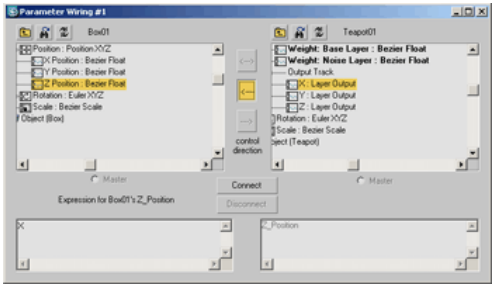
1. Make sure both layers listed on the Animation Layers toolbar are on, as indicated by the light bulb icons.
2. Create a box next to the teapot.
3. Right-click it and choose Wire Parameters from the quad menu.
4. Navigate through the pop-up menu and choose Z Position.



5. A dashed line appears. Click the teapot and choose the X Output Track from the pop-up menu.

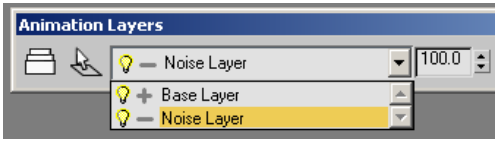


6. The *Parameter Wiring dialog (page 2–412)* opens, with the two highlighted parameters. Click the left directional arrow button to control the cube's position using the teapot's animation.



7. Click Connect and close the window. Drag the time slider and notice how the cube's position in Z matches the teapot's layered animation in X.
8. Select the teapot and expand the layer list from the Animation Layers toolbar.
9. Click the plus sign icon next to the *Noise Layer*.

   This excludes the layer from the output track.



10. Drag the time slider again. The cube doesn't move erratically like the teapot.

    Try including and excluding both layers in the list to see the results.

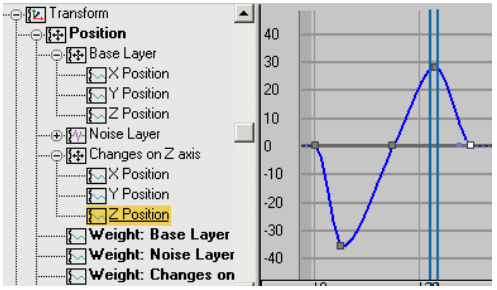
Example: To copy and paste a layer and update an active layer:

This procedure continues from the previous procedure and focuses solely on the teapot.

1. Make sure the teapot is selected and choose Base Layer from the animation layers list to make it active.
2.  Click Copy Anim Layer on the Animation Layers toolbar to buffer this layer and the data it contains.
3.  Click Paste New Layer. On the Rename Anim Layer dialog, enter "Changes on Z axis" as name and click OK.

   This new layer becomes the active layer.
4. While Auto Key is on, move the teapot on the Z axis at frames 5 and 20.

   This animation update only affects the position tracks of the active layer.



Example: To collapse a layer:

This procedure continues from the previous procedure.

1. In the animation layers list, turn off *Noise Layer*. Then, choose *Changes on Z axis* to make it active.

2.  Click Collapse Anim Layer on the Animation Layers toolbar.

   This merges the layer onto the next available one (Base Layer), skipping *Noise Layer*, which is off.

3. Turn *Noise Layer* back on in the list.

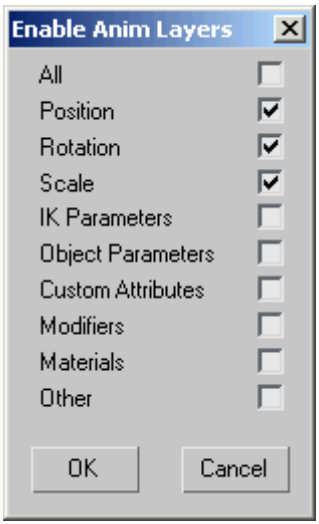
Interface



 **Enable Anim Layers—**Assigns a Layer controller to tracks you specify.

Note: Enabling animation layers does not create a new layer, but rather transfers all chosen controller tracks to the Base Layer.

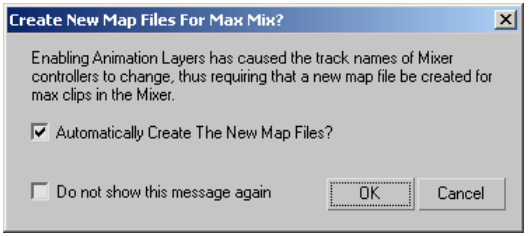
This opens the Enable Anim Layers dialog, which, similar to the *Set Key Filters dialog (page 2–280)*, lets you turn on tracks you want to assign a Layer controller to.



The Position, Rotation, and Scale tracks are on by default.

Note: Once you turn on a track, you can turn it off only by disabling the respective Layer controller.

If you enable animation layers for an object that has animation loaded in the *Motion Mixer (page 2–604)*, the following dialog opens:



Remapping is necessary in this case because enabling Animation Layers causes the full controller names to change. For example, if a sphere's X position track before enabling Animation Layers is Sphere01\Transform\Position\X Position, then after enabling animation layers it might change to Sphere01\Transform\Position\Base Layer\X Position (the layer name is inserted into the controller name).

If you accept the default dialog option, to create the new map files automatically, 3ds Max automatically generates any necessary *XMM files (page 3–478)* with the same names as the XAF files and places them in the same directory as the original map files, or if no map files exist, in the same directory as the XAF files. Thereafter, when you open this scene file the map file is loaded automatically and no additional mapping is required.

However, if you turn off Automatically Create The New Map Files? before clicking OK, no map files are generated, and the next time you open the scene file you're prompted to create map files. Without map files, the animation doesn't appear correctly.

 **Select Active Layer Objects—**Selects all objects in your scene containing the active layer.

[animation layer list]—Displays all existing layers for the selected object. Each layer in the list contains toggle icons to turn it on or off, as well as include or exclude it from the controller's output track. See *Working with Animation Layers (page 2–328)* for more details.

[weight]—Sets a global weight for the active layer, which affects If the active layer is shared by several objects (ex. Base Layer), changing its weight affects them all.

 **Anim Layer Properties—**Opens the *Layer Properties dialog (page 3–662)*, which lets you set global options for layers.

 **Add Anim Layer—**Opens the *Create New Animation Layer dialog (page 2–334)*, which lets you specify settings related to the new layer. This adds a new layer to every track that has a Layer controller.

 **Delete Anim Layer—**Removes the active layer, along with the data it contains. A confirmation dialog prompts you before deletion.

Note: You cannot delete the Base Layer. Alternatively, Click *Disable Anim Layer* to removes the Layer controller entirely.

 **Copy Anim Layer—**Copies the active layer's data and enables Paste Active Anim Layer and Paste New Layer.

 **Paste Active Anim Layer—**Overwrites the active layer's controller type and animation keys with the copied data.

Note: You cannot paste a copied layer from one object onto another'.

 **Paste New Layer—**Creates a new layer with the copied layer's controller type and animation keys. Opens the Rename Anim Layer dialog, in which you can use the default layer name or enter your own.

 **Collapse Anim Layer—**Collapses the active layer to the one below it, as long as it is not turned off. If it is, the collapsed layer cycles through the list until it finds an available layer.

Note: Collapsing a layer does not remove the Layer controller. Click Disable Anim Layer to remove it.

 **Disable Anim Layer—**Removes the Layer controller from the selected object. The animation keys on the Base Layer revert to the original controller. A dialog prompts for confirmation.

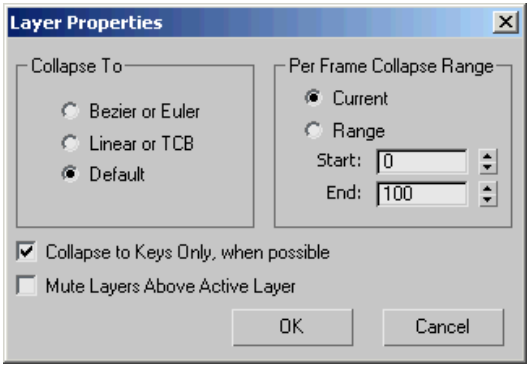
Note: Before you can disable the Base Layer, you have to delete or collapse all layers above.

Layer Properties Dialog

Right-click unused area of any toolbar. > Animation Layers > Create New Animation Layer

This dialog provides global options in regards to collapsing *animation layers (page 2–326)* and isolating the active layer from the rest.

Interface



Collapse To— Sets the controller type for when you collapse a controller track onto a non-keyable controller track, such as a Noise controller.

- **Bezier or Euler—**On collapse, the resulting track is assigned either a Position XYZ controller (for a Position track), a Bezier Scale controller (for a Scale track), or an Euler XYZ controller (for a Rotation track).
- **Linear or TCB—**On collapse, the resulting track is assigned either a Linear controller (for Position and Scale tracks) or a TCB controller (for a Rotation track).
- **Default—**On collapse, the resulting track is assigned a default controller based on the original controller track. Refer to *Specifying Default Controllers (page 2–294)* for more details.

Per Frame Collapse Range— Sets the range to cover when you collapse a layer.

Note: For best results, turn off Collapse to Keys Only, When Possible; otherwise, based on the collapsed layer's controller and tangent type, the system might collapse to keys instead.

- **Current—**When chosen, a key is set on every frame of your scene's animation range when you collapse a layer.
- **Range—**When chosen, a key is set on every frame of the animation range you specify.

Start/End—Sets the collapse range boundaries.

Collapse To Keys Only, When Possible—When you collapse a layer, the keys are merged only when the respective controllers are of the same type, same tangent types, and Blend Eulers As Quats is on. Default=on.

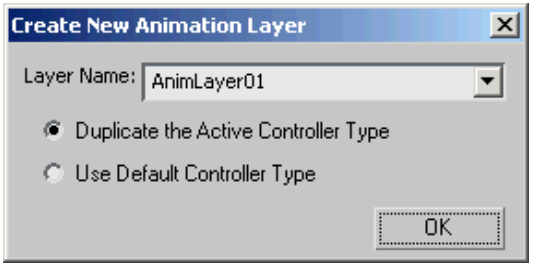
Mute Layers Above Active Layer—When on, you see the effects of the layers only up to the active layer, inclusively. This is similar to the Visible Before/After options of the *biped Layers rollout (page 2–974)*.

Create New Animation Layer Dialog

Right-click unused area of any toolbar. > Animation Layers > Create New Animation Layer

This dialog lets you choose both the name and controller type of the new layer.

Interface



Layer Name—Lists all available layers in your scene. You can use the default name, enter a custom name, or pick one from the list. If you choose an existing layer from the list, it inherits that layer's name and weight properties. Refer to *Working with Animation Layers (page 2–328)* for more details.

Tip: You can rename a layer after it is added from the *Layer Controller dialog (page 2–325)*.

Note: All animation layer names persist in the Layer Name list even after you delete their host object. Resetting the scene clears the list.

[controller type]—Determines the controller type to assign to each track within the animation layer.

- **Duplicate the Active Controller Type—**The new layer's controller type becomes the same as the one from the active layer, on a per-track basis.

  For example, if a track in the active layer has a Noise controller, adding a new layer copies that controller type.

- **Use Default Controller Type—**When chosen, the new layer's controller does not take into account the active layer's controller type and instead uses the original track's type.

Limit Controller

Main toolbar > Curve Editor (Open) > Highlight a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Float Limit

Graph Editors > Track View - Curve Editor > Highlight a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Float Limit

The Limit controller lets you specify upper and lower limits to available controller values, thus restricting the potential value range of the controlled track. For example, in a character rig you could use this to restrict rotation on finger joints so the fingers can't bend backward. Basically, once a track is limited, and the limit is active, you can't set a value for the track beyond the limit. You can apply a Limit controller to any other type of controller; the limited controller (that is, the original controller) then appears as a child of the Limit controller in the Track View hierarchy.

Because the Limit controller doesn't change the original controller, you can easily switch back and forth between the original and the limited animation by toggling the limits. But if you *collapse (page 2–522)* the Limit controller, the result is the limited animation, and the original animation is no longer available.

Use of the Limit controller makes it faster to set up and create animation. It eliminates the need for expressions or scripts to do the same thing, this improving the ease with which you can set up automation in hierarchies and rigs, creating effects such as avoiding collision, FK joint limits, etc.

For example, a Technical Director might want to assign limits to float values that animators will adjust to make posing easier, enforce standards, or prevent inconsistencies and mistakes. A TD could use limits for:

- light intensity
- finger rotation
- position extents for sliding drawers
- camera field-of-view
- limiting patch resolution to that supported by the game engine

Character riggers can use limits to create complex relationships or shortcuts in rigs. Examples include:

- Wire wrist-twist bones to the rotation of the hand, but limit the rotation of the wrist-twist or hand to stop short of any flipping that might occur.
- Prevent unnatural translation of a spline IK helper by causing limit values for its local position tracks to react to the angle or distance of the helper from surrounding helpers.
- React to when IK is disabled, and match the limits of the solver in FK.
- Limit the position of the IK chain swivel-angle target for the knee so that it can never go behind the character. That is, the target is linked to character's root.

Use Cases

Following is a brief list of examples of specific uses for the Limit controller:

- Select the Y Rotation tracks of all finger bones in a character's hands and assign a Limit controller. Use Set Key mode to pose the fingers at the extremes of their ranges of motion and then use the *right-click menu (page 2–340)* > Set Upper/Lower Limit commands to limit the fingers' motion.

  Note: Limits are included for Set Key via the Other category.

- You apply a Noise controller to the X Rotation track of a window shutter object to make it rattle, but the object intersects the building at the lower end of its motion range. Apply a Limit controller to the track and set the Lower Limit value to prevent the intersection.

- Following from the previous example, say the wind tears the window shutter from the building at frame 100. If the Noise controller is in a Float List with a Bezier Float, you can weight the noise to 0 at frame 100, when the shutter is torn from the building by the wind, at which point the animator will key the animation. Limits are unnecessary after frame 100. At frame 99, you'd use Set Lower Limit. Then, at frame 100, you'd edit the Lower Limit value to set the limit to much lower, essentially removing the limit. Finally, you would set keyframes to animate its movement past frame 100.

Limitations of the Limit Controller

It's important to be aware of the following limitations of Limit controllers:

- Limit controllers do not limit IK joints. All IK solvers use their own limits, and ignore Limit controllers. Interactive IK uses the current controllers, but it also has its own limits, so using Limit controllers with interactive IK will produce unpredictable results.

- The Limit controller currently supports only Float controllers. Thus, the Limit controller is currently available only under the name Float Limit.

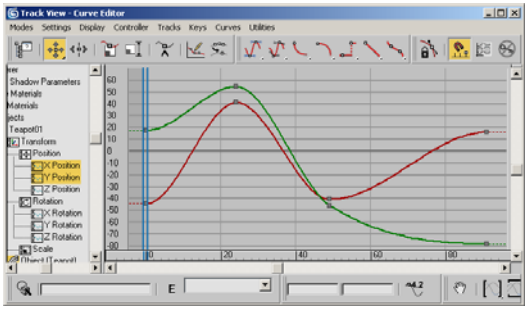
- Upper and lower limits values are expressed in the same units or coordinates as those used by the limited controller. Thus, for example, Float limits on an X Position track are in local coordinates, not in global coordinates.

Procedure

Example: To use the Limit controller:

1. Start with an animated object, and open Curve Editor.

   In this example, the teapot position was animated on the X axis between extents of about -45 to 40, and on the Y axis between extents of about -80 and 54.



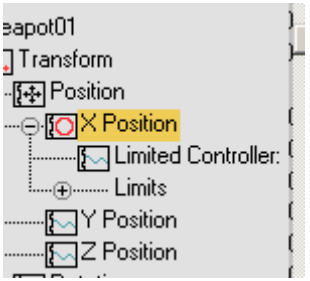
2. Assign the Float Limit controller to a track. In this example, we'll assign it to the X Position track.

   This opens the Float Limit Controller dialog, where you can set limits and other values, but you might find it easier to do this interactively, so for now just accept the defaults by closing the dialog.

   The track retains its original name (X Position, in this example), but the icon has changed to

indicate that it's now a Limit controller, and is expandable (see illustration in following step).

3. Expand the track.

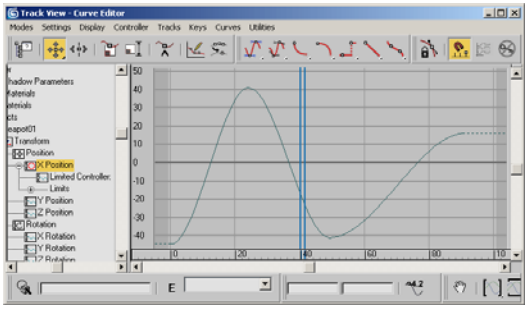


The hierarchy contains a child track, named Limited Controller, and a new, expandable child branch named Limits.

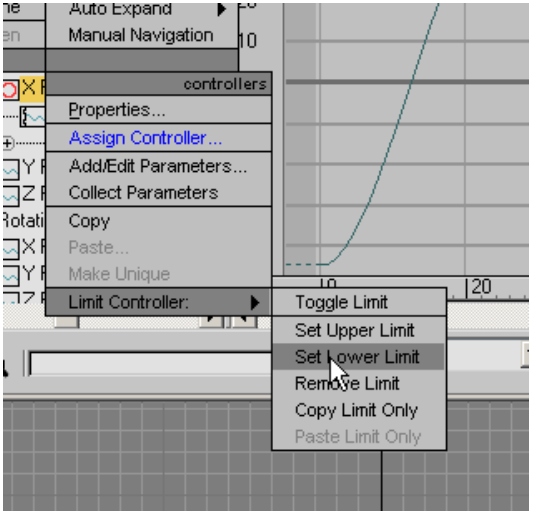
4. Click the X Position track and then the Limited Controller track to compare them.

   Because the original X-axis movement falls within the default limits, the two tracks are identical.

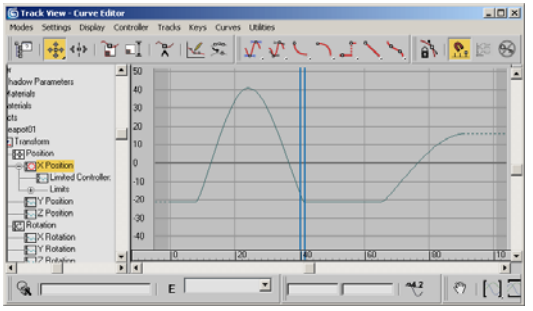
5. Scrub the animation until the X-axis position is where you want to set a limit. In this example, the position is about -20, where we'll set a lower limit.



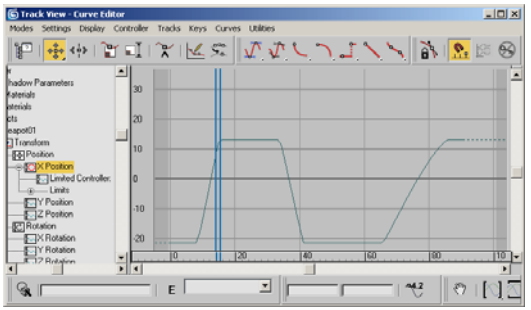
6. Highlight the X Position track, and then right-click and from the right-click menu choose Limit Controller > Set Lower Limit.



The graph curve is now clipped below the lower limit. When you scrub the animation, the object doesn't move beyond that position.



7. Again scrub the animation, stop where you want to set the upper limit, right-click the X Position track, and choose Limit Controller > Set Upper Limit.



In this example, we've set it at about 13. The graph curve is now clipped above the upper limit and below the lower limit. The resulting motion is likewise constrained.

8. Scrub the animation.

   The motion on the X axis is limited at both extremes; wherever the object originally moved beyond the limits, it now behaves as though it's hitting a wall.

9. Again compare the two tracks.

   The Limit Controller (X Position) track clearly shows the upper and lower clipping, while the Limited Controller track shows the original motion.

   The original motion still exists as the Limited Controller track; you can restore it temporarily by toggling the Limit controller.

10. Highlight the X Position track, and then right-click and choose Limit Controller > Toggle Limit.

    Now, when you scrub the animation, the object moves as before.

11. Choose Limit Controller > Toggle Limit again to restore the limits

    You can copy the Limit controller in two ways: the limits only, or the limits and motion combined. First, you'll copy the limits only.

12. Highlight the Y Position track and note how its curve differs from that of the X Position track.

13. Highlight the X Position track, and then right-click and choose Limit Controller > Copy Limit Only.

14. Highlight the Y Position track, and then right-click and choose Limit Controller > Paste Limit Only. When the Paste dialog opens, click OK to confirm the paste as a copy.

    Now the Y Position track is also limited, using the same extents as the X Position track, but retaining its original underlying motion. To verify this, compare the Y Position track with its child Limited Controller track.

    Next, you'll copy the entire limited X Position track to the Z Position track, including motion and limits. The Z Position track currently is not animated.

15. Highlight the X Position track, and then right-click and choose Copy.

16. Highlight the Z Position track, and then right-click and choose Paste.

    Compare the X and Z Position tracks and both their child Limited Controller tracks. Each corresponding pair of curves is identical. The motion appears as though the object is hitting the inside edge of a box.

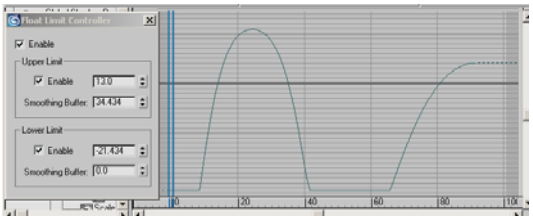
    Finally, we'll cover the Smoothing Buffer settings. By default, sharp corners are created wherever a curve is limited, causing abrupt changes in motion. You can smooth off these corners with the Smoothing Buffer parameters, resulting in more natural-looking motion.

17. Highlight the X Position track, and then right-click it and choose Properties.

    This reopens the Float Limit Controller dialog.

18. Use the Upper Limit group > Smoothing Buffer spinner to increase the Smoothing Buffer value as far as it can go.

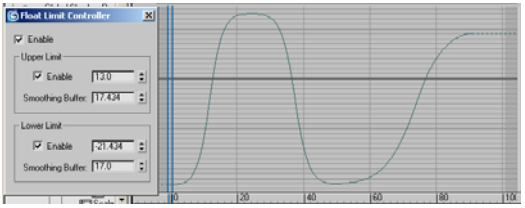
    As you increase the value, the corners of the upper limit on the graph become smoother.



Note that there's an upper limit to this value; in this case, it's 34.434. This limit is determined by the values of the other three settings on the dialog.

19. Try increasing the Lower Limit group > Smoothing Buffer. It's not possible with the upper smoothing value at its maximum.

20. Decrease the upper smoothing value, and then increase the lower smoothing value.

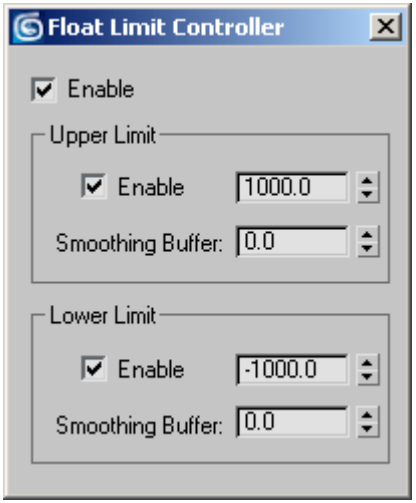


Now you get smoothing at both the upper and lower limits.

Interface

The Limit Controller interface comprises the Float Limit Controller dialog and several-right-click menu items.

Float Limit Controller dialog



The Float Limit Controller dialog opens when you first assign the Limit controller, or when you right-click a highlighted Limit-controller track and choose Properties.

Enable—Toggles the Limit controller. When off, the original values of the limited track are in effect. When on, the original values are limited by the Upper Limit and Lower Limit values.

Note: You can enable and disable all Limit controllers in the scene simultaneously with the Toggle Limits command, available from the 3ds Max Animation menu. If some Limit controllers are on and the rest are off, Toggle Limits turns them all on.

Upper Limit Group

Enable—Toggles the upper limit set by the controller. When off, no upper limit is imposed. Default=on.

[Upper Limit value]—The highest value permitted by the Limit controller. Any values above this value in the original controller are clipped; that is, they're set to this value, unless smoothing is in effect. Default=1000.0.

You can animate this value via keyframing and other standard methods, and manipulate this animation in Track View via the Upper Limit track in the controller's Limits branch.

Smoothing Buffer—Specifies a smoothing value, so that clipped values at the beginning and end of a clipped range gradually increase and decrease instead of leveling off abruptly.

The maximum total smoothing is determined by the Upper and Lower Limit values. (Upper Limit > Smoothing Buffer value) + (Lower Limit > Smoothing Buffer value) cannot exceed this total. For the smoothest possible results at the upper and lower extents of the clipping, set either Smoothing Buffer value to the maximum amount, and then back it off to half that value and set the other Smoothing Buffer to the same amount.

You can animate the Smoothing Buffer value via keyframing and other standard methods, and manipulate this animation in Track View via the Upper Smoothing track in the controller's Limits branch.

Lower Limit Group

Enable—Toggles the lower limit set by the controller. When off, no lower limit is imposed. Default=on.

[Lower Limit value]—The lowest value permitted by the Limit controller. Any values below this value in the original controller are clipped; that is, they're set to this value, unless smoothing is in effect. Default=-1000.0.

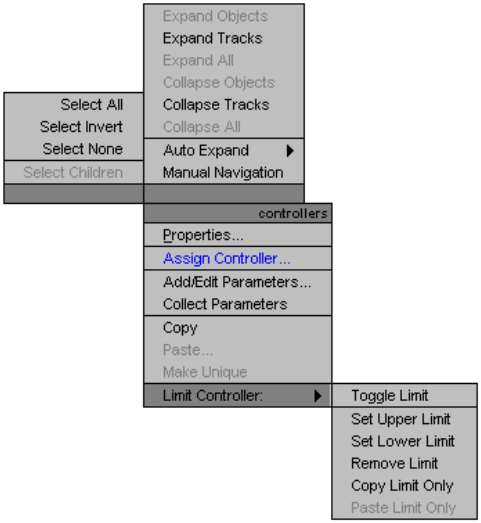
You can animate this value via keyframing and other standard methods, and manipulate this animation in Track View via the Lower Limit track in the controller's Limits branch.

Smoothing Buffer—Specifies a smoothing value, so that clipped values at the beginning and end of a clipped range gradually decrease and increase instead of leveling off abruptly.

The maximum total smoothing is determined by the Upper and Lower Limit values. (Upper Limit > Smoothing Buffer value) + (Lower Limit > Smoothing Buffer value) cannot exceed this total. For the smoothest possible results at the upper and lower extents of the clipping, set either Smoothing Buffer value to the maximum amount, and then back it off to half that value and set the other Smoothing Buffer to the same amount.

You can animate the Smoothing Buffer value via keyframing and other standard methods, and manipulate this animation in Track View via the Lower Smoothing track in the controller's Limits branch.

Limit Controller right-click menu



To access the Limit Controller right-click menu, highlight a Limit Controller track in the Track View hierarchy, and then right-click the track and move the cursor to the Controllers quadrant > Limit Controller menu item.

Note: After using Copy Limit Only on a Limit controller track, you can apply a new Limit controller with the same limits to any track by invoking Paste Limit Only.

Toggle Limit—Turns the Limit controller on and off. This is the same as the *Enable check box* on the Float Limit Controller dialog.

Set Upper Limit—Sets a value/key in the Upper Limit track at the current frame equal to the value at the current frame of the limited track. If there is only one key, the value of the limit is constant over time.

Set Lower Limit—Sets a value/key in the Lower Limit track at the current frame equal to the value at the current frame of the limited track. If there is only one key, the value of the limit is constant over time.

Remove Limit—Deletes the Limit controller, restoring the original controller without limits.

Copy Limit Only—Copies only the Limit values and ignores the limited controller of the highlighted track.

Paste Limit Only—Applies only the copied Limit values, or adds a Limit controller with the copied values if one doesn't exist, while retaining the values of the original limited track.

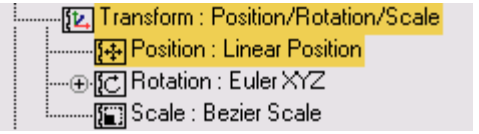
As with pasting copied controller tracks, you can paste the Limit values as a copy or an instance of the copied Limit controller, with the option to replace all instances of the paste target.

Linear Controller

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Linear

Graph Editors > Track View – Curve Editor > Select a track in the Track View hierarchy. > Track View toolbar > Controller menu > Assign > Linear

Animation menu > Position, Rotation or Scale controllers > Linear



The Linear controller interpolates between animation keys by evenly dividing the change from one key value to the next by the amount of time between the keys.

Linear controllers do not display a properties dialog. The only information stored in a linear key are the time and animation values.

Use Linear controllers whenever you want a very regular, even transition from one key to the next. For example, use a Linear controller for:

- A color parameter to change from one color to another at a constant rate of change.
- Transforms to produce mechanical, robot-like motion.

Procedure

To assign a Linear controller:

1. Select an animated object.
2. In the Motion panel > Parameters > Assign Controller rollout, select the Position track in the list window.
3.  Click Assign Controller, and then select Linear Position in the Assign Controller dialog.

   The animated object has a mechanical motion.

You can also assign this controller in Track View.

Interface

No Properties dialog is available for Linear controllers. You can, however, move keys in Track View to change the animation.

Time—Change linear key time. Move keys horizontally in Track View Edit Keys or Function Curves mode.

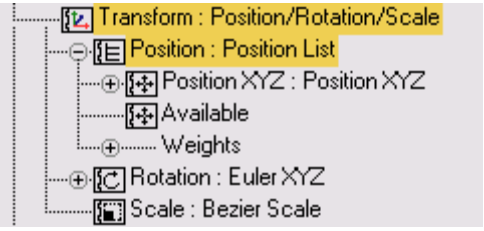
Value—Change linear animation values by moving keys vertically in Track View Function Curves mode, change the parameter value in the viewport with Auto Key on.

List Controller

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > List

Graph Editors > Track View – Curve Editor > Select a track in the Track View hierarchy. > Track View toolbar > Controller menu > Assign > List

Select an object in the viewport > Motion panel > Assign Controller rollout > Highlight either Position, Rotation or Scale controller > Assign Controller button > choose Position List, Rotation List or Scale List.



The List controller combines multiple controllers into a single effect. It is a compound controller with tools for managing the order in which its internal controllers are calculated. Controllers are evaluated in top-to-bottom order. In addition, you can specify a Weight setting for each controller in the list to determine its relative influence.

When you assign a List controller to a parameter, the current controller is moved one level below the List controller; it becomes the first controller in the list. A second parameter, named Available, is added below the List controller. This is an empty placeholder for the next controller you add to the list.

By default, each listed controller's Weight setting is set to 100.0. This setting can be changed to increase or decrease the effect the controller has on an object.

When you assign a controller using the Animation menu, a list controller is automatically assigned as a default, with the chosen controller placed first in the list. This behavior differs from when you assign a controller in the Motion Panel or Track View, where only the specified controller is assigned.

You can animate List controller weights to achieve the equivalent of a non-linear animation system. Each list controller track can hold different values from frame to frame that you can turn on or off, or blend between by animating the weights.

Procedure

Example: To use the List controller to combine two controllers:

1. Create a box, and then turn on Auto Key.
2. Create a simple animation containing three position keys.
3. On the Motion panel > Parameters > Assign Controller rollout, click the position track.
4.  Highlight the Position transform and click Assign Controller.
5. Choose the Position List in the Assign Position Controller dialog. Click OK.
6. Click the plus sign in the Position track to expand the track.

   A Position XYZ track and a track labeled Available are now visible.

7. Click the track named Available and then click the Assign Controller button.
8. Choose Noise Position in the Assign Controller dialog, and then click OK.

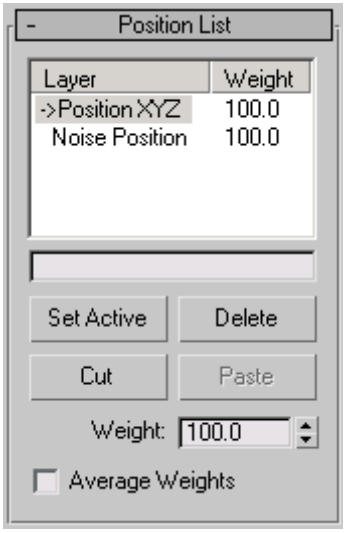
   The Noise Controller dialog opens.
9. Play the animation.

   The box follows the original animated path but vibrates randomly.

   Experiment with different Weight settings for each controller and see the results.

Interface

Motion panel > Parameters > List rollout



List Window—Displays all the controllers in a list, with evaluation order from top to bottom. The controller at the top of the list is evaluated first; lower controllers are layered onto the result of higher controllers.

Set Active—Determines which controller is affected by interactive changes in the viewport. The current active controller is marked with an arrow in the list.

For example, you have an object using a Position List, with Noise Position and Bezier Position in the list. The Bezier Position controller handles the general motion path of the object, while the Noise Position controller adds a random shake to the motion.

- If Bezier Position is the active controller, you can freely move the object in the scene.
- If Noise Position is the active controller, you cannot move the object because Noise is not an interactive controller.

Delete—Deletes the selected controller.

Cut—Removes the selected controller and stores it in a temporary clipboard. The clipboard contents are only held until you paste the controller, close the List Controller dialog, or exit the Motion panel.

Paste—Puts the contents of the controller clipboard in the position above the selected controller. The clipboard is empty after pasting.

Weight—Exaggerate or minimize the effects of a controller by increasing or decreasing its weight value. Default=100.0.

Average Weights—When on, the weight values of all the controllers in the list are averaged. Available only for the Position List and Scale List controllers. Default=off.

Pose to Pose—Enables blending among controllers in the list. Available only for the Rotation List controller. Default=off.

When off, the List controller weights each rotation individually and adds up the weighted rotations for the final effect. You animate by stacking layers on top of layers, and then weighting in and out their effect. This is the default behavior.

When on, each pose is blended with the results from the previous controllers in the list. This method is not additive, but is instead true pose-to-pose blending. Adjusting the weight of a

controller in the list will blend to and from that pose (controller in list). If the last weight is 100.0, then adjusting the weights of the previous poses in the list won't have any effect because the last pose has full effect. This is a stack-based approach that is best used for nonlinear animation (NLA) effects. You can paste motion clips or even single-frame poses onto new layers and use the weights to blend the poses together.

Editable Name Field—Select one of the controllers in the list window of the Controller list, and then enter a descriptive name in this field. This field lets you rename the controller so it's easier to identify.

Local Euler XYZ Rotation Controller

This controller is no longer available. You can, however, still edit objects that were assigned this controller previously.

The Local Euler Rotation controller is similar to the *Euler XYZ Rotation controller (page 2–318)*, but the angles are with respect to the rotated object's local coordinate system. With the standard Euler Rotation controller, angles are with respect to parent coordinates. Parameters for this controller are displayed in the Motion panel.

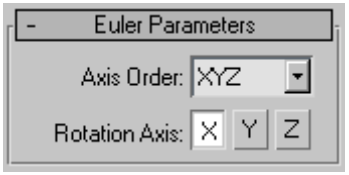
The Local Euler XYZ Rotation controller has three rotation tracks to control rotation on the X, Y, and Z axes. Setting an X axis rotation key will not create keys in the Y and Z tracks on the same frame; each track can be keyed independently. The rotation order of the axes is set in the Motion panel.

Interface

The Local Euler parameters are displayed in the Motion panel.

Euler Parameters rollout

This rollout is displayed when you select Rotation in the *PRS Parameters rollout (page 2–357)*.



Axis Order—Sets the order that rotations are calculated in the list. The default is X,Y,Z order, where the X axis is rotated first.

X—Displays controller properties for X axis rotation angle.

Y—Displays controller properties for Y axis rotation angle.

Z—Displays controller properties for Z axis rotation angle.

Each axis uses its own independent controller using the float data type. For example, the X and Y rotation axes could use Bezier Float controllers, while the Z Rotation axis uses a Noise Float controller.

Look At Controller

Create or select an object that contains a Target component, such as a target spotlight or camera. > Motion panel > Look At Parameters rollout

The Look At controller is automatically assigned as a transform animation controller upon creation of objects that contain targets, including Target Camera, targeted lights (including IES Sun/Sky) and the Tape helper. It cannot be assigned by the user. To assign the equivalent of a Look At controller to an object, use a *LookAt constraint (page 2–406)*.

Note: A targeted object's pitch and heading are adjusted by moving the target, so the only

orientation setting that can be controlled directly by the user is Roll, or bank.

Example: To prevent flipping of targeted objects during rotation:

1. Add a *Target Camera object (page 2–1371)*.
2. Activate the Move tool and use the *coordinate display (page 3–708)* to position the camera and target at **0,0,0** and **80,0,0**, respectively.
3. Add a *Dummy object (page 2–16)* and position it at **0,0,0**.
4. Use Select And Link to link the camera to the dummy, and then link the camera target to the dummy. At this point, the dummy is parent to both the camera and its target.
5. Rotate the dummy about its Y axis, and watch the camera.

   As the target passes through the zenith and nadir of its orbit, the camera flips.
6. Select the camera and go to the Motion panel. On the Look At Parameters rollout, turn on Use Target As Up Node.
7. Again rotate the dummy about its Y axis.

   The camera no longer flips.

Interface



After you create or select an object that contains a Target component, you can access the object's Look At properties on the Motion panel. In this rollout you can change the target, create and delete animation keys, set the axis, and adjust other, related parameters.

Create Key—Sets a position, roll (orientation), or scale key at the current frame, depending on which button you click.

Delete Key—Deletes a position, roll (orientation), or scale key at the current frame, depending on which button you click.

Pick Target—Lets you set a target other than the default Target object. Click this button, and then select the new object to use as a target.

Thereafter, the new target controls the object's orientation. The original target remains in the scene, and can be deleted or used as a Dummy helper.

Axis—Specifies the local axis that looks at the target. The Flip check box is used to flip the directions of the axis.

Use Target as Up Node—When turned on, the controller forces the object on which it acts (source node) to keep one of its local axes aligned with the look-at direction (the vector between the source node and the target node). it also prevents the source node from rotating around the look-at direction, to avoid flipping about the object's local Z axis.

The flipping behavior is most commonly seen when the line between a targeted object and its target is close to vertical; that is, their positions on the World Z axis are nearly the same.

The option works by aligning one of the source node's local axes with one of the target node's local axes. These axes are picked automatically by the software.

Note: This feature lets you properly manipulate Luminaire assemblies . It also provides for trouble-free operation when you *i-drop (page 3–523)* manufacturer luminaire assemblies.

Tip: In some cases, the object will flip 90 or 180 degrees even when the Use Target As Up Node option is turned on. This behavior occurs due to the fact that the axis alignment is automatic. To resolve this, apply a roll angle to the object with the coordinate display.

Position/Roll/Scale—These three buttons let you specify the other rollouts that appear for this controller. In all three cases, the *Key Info (Basic) (page 2–304)* and *Key Info (Advanced) (page 2–306)* rollouts appear. When Position is active, an additional Position XYZ Parameters rollout lets you specify the position axis.

Master Point Controller

Main toolbar > Curve Editor (Open) > Expand tracks on an object with animated vertices or vectors in the Track View hierarchy. > Master Point Controller

Graph Editors > Track View – Curve Editor > Open Track View > Expand tracks on an object with animated vertices or vectors in the Track View hierarchy. > Master Point Controller

The Master Point controller is assigned automatically whenever control points, vertices, or vectors (tangent handles) are animated in the sub-object mode of an Editable Spline, Editable Patch, Editable Mesh, Editable Poly, NURBS surface, or FFD. By allowing you to select and move all the sub keys, visually correlate keys to points in the viewports, and change key properties quickly, this controller helps to manage the numerous tracks created when animating vertices, control points, and vectors.

The Master Point controller is displayed as a track with green keys in Track View. Sub tracks below the Master track contain all the animated vertices, control points, and vectors.

Procedure

Example: To animate vertices, viewing the Master Point controller track:

1. In the Front viewport, create a sphere.
2. Right-click the sphere and select Convert to Editable Mesh in the dialog.
3. On the Editable Mesh > Selection rollout, click Vertex.
4. On the Editable Mesh > Soft Selection rollout, click Use Soft Selection.
5. Turn on Auto Key, and move the time slider to frame 10.
6. In the viewports, select and move some vertices on the sphere.

   A key appears on the track bar at frame 10.

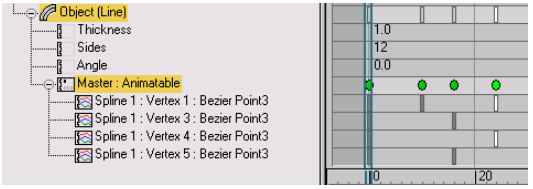
7. Right-click the sphere, and select Curve Editor in the dialog.

   Track View — Curve Editor is displayed with the sphere at the top of the Controller window.

8. Change modes to Dope Sheet and expand all the sphere tracks in the Track View hierarchy.

   The Master Point Controller track is displayed, with tracks for the animated vertices below it.

Interface



Master Track Keys—While in the Track View – Dope Sheet, selecting a master key (green) selects all the sub keys at that frame.

If hundreds of vertices are animated, collapse the master track so that only the master track is visible. Moving the green keys will move all the sub keys.

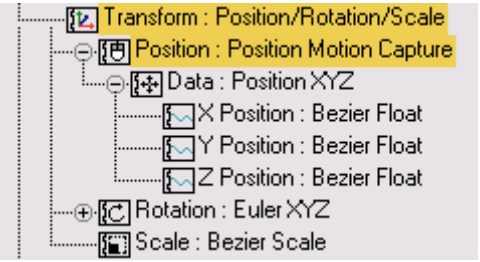
Master Track Key Info dialog

Right-click a green master key in Track View to display the *Master Track Key Info dialog (page 2–391)*.

Motion Capture Controller

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View toolbar > Controller menu > Assign > Motion Capture

Graph Editors > Track View – Curve Editor > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Motion Capture



You can control an object's position, rotation, or other parameters using an external device with the Motion Capture controller.

Note: Rotation Motion Capture uses Euler rotations with an X, Y, Z axis order and is subject to the limitations of Euler rotations. See *Euler XYZ Rotation Controller (page 2–318)*.

The Motion Capture System

To use motion capture in 3ds Max, you follow these steps:

1. In Track View — Dope Sheet or Motion panel, you assign motion capture animation controllers to the specific tracks you want controlled by external devices.
2. After assigning the motion capture controller, you open the Properties dialog for the track and bind the type of peripheral device(s) that you want. Currently supported devices are mouse, keyboard, MIDI device, and joystick. Each device has specific properties.
3. After binding the devices, adjust their settings and parameters in the lower portion of the Track Properties dialog. These controls vary depending on the type of device.

4. In the Utilities panel, access the Motion Capture utility. You can test and record your motion for any combination of tracks over any range of frames.
5. When you assign a Motion Capture controller, the previously assigned controller is maintained as a child of the Rotation controller. You can continue to adjust the rotation of the object using standard transform controls, while still making motion-capture control available.

Procedure

To use your mouse as an input device to animate the position of a box:

1. Create a box.
2. In Motion panel > Assign Controller, select the position track for the box in the list window, and then click the Assign Controller.
3. Select Position Motion Capture in the Assign Position Controller dialog, and then click OK.
4. Select X Position in the Motion Capture dialog, select mouse input device, and then click OK.

   Note: You can open the Motion Capture dialog by selecting the Position track in the Assign Controller list window, and then right-click and select Properties.
5. In the Mouse Input Device rollout, make sure that horizontal is turned on and Scale is set to 1.
6. In Utilities panel > Motion Capture, click Box01 > Position in the Track list window. Click Start in the Record Controls group.

   Real-time recording is in effect; move the mouse horizontally back and forth. The box moves back and forth along the X axis.

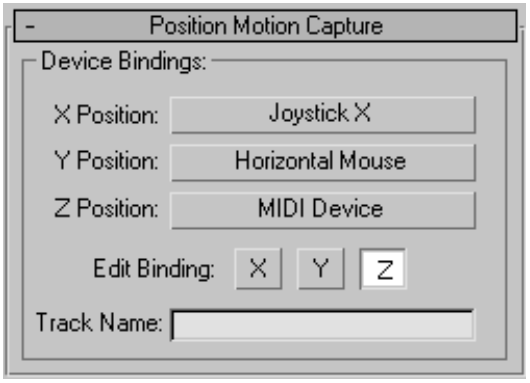
   Play the animation to view the recorded movement.

   By default, a key is generated at every frame.

Interface

Right-click the transform track in the Track View - Dope Sheet hierarchy to display the Motion Capture Properties dialog.

Position Motion Capture rollout



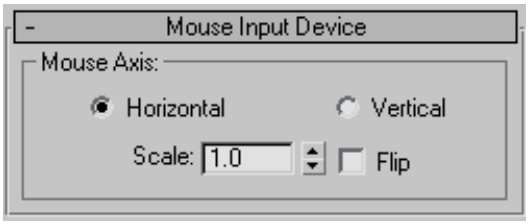
X, Y, Z Position—Click an axis's button to display a Choose Device dialog, which lets you pick a motion-capture device from a list of alternatives: keyboard, joystick, mouse, or MIDI.

Choose Device dialog—Select a device from a list of devices.

Edit Bindings XYZ—Displays device controls in the device rollout in the lower part of the Motion Capture dialog. If no device has been assigned to an axis, this option is not available.

Track Name—Names the Motion Capture Data track. This overrides the default naming convention.

Mouse Input Device rollout



This rollout controls animation using the horizontal or vertical motion of the mouse. The available settings include:

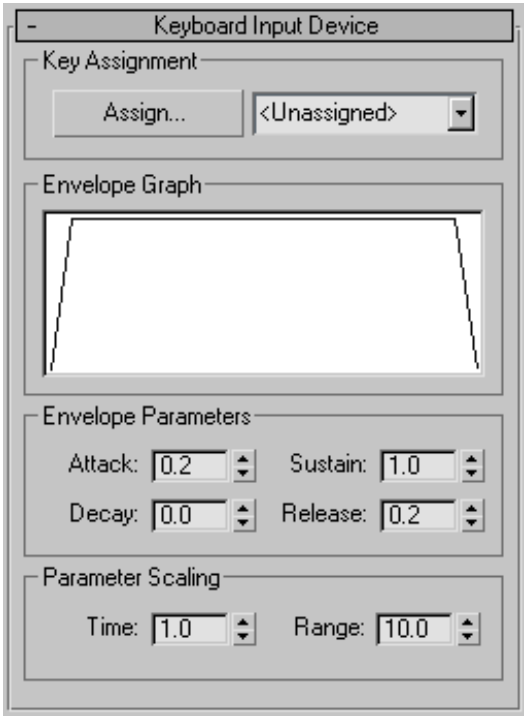
Horizontal/Vertical—Specifies which mouse motion drives the animation.

Scale—Scales the relative effect of the mouse movement to the animation response. Spinner Value=float, 0 to 999,999

Flip—Flips the direction of the response relative to the mouse movement. For example, if moving the mouse horizontally to the right produces a clockwise effect on a Rotation controller, activating Flip will reverse the rotation to counterclockwise.

Keyboard Input Device rollout

This rollout lets you assign most keyboard keys to drive the animation.



Assign—Click and then press any key. The assigned key is displayed in the list window at right.

List—Open the list to select a key.

Envelope Graph group

The Envelope Graph group displays a representation of the amplitude curve over time.

Envelope Parameters group

These options specify the time the envelope of the action takes effect. This relates to the duration of pressing and releasing the key.

Attack—Specifies the time, after pressing the key, for the value to reach its maximum level.

Decay—Specifies the time, after having reached maximum, for the value to fall to that specified by the Sustain spinner.

Sustain—After the Attack and Decay, specifies the value sustained until you release the key.

Release—After releasing the key, specifies the time for the value to fade out to zero.

Parameter Scaling group

Provides controls for setting the scale of the envelope and the range of the output value.

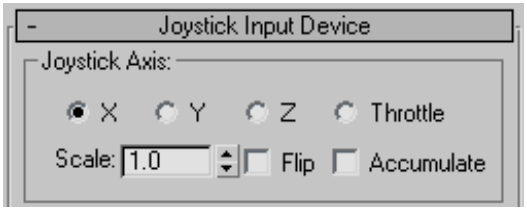
Time—Specifies the scale of the Attack, Decay, and Release parameters. The value represents the number of seconds for 1 unit. For example, if this value is 1.0, then an Attack value of 1.0 equals 1 second.

Range—Sets the maximum output value of the controller.

Note: This controller ignores the state of the Ctrl, Alt, and Shift keys.

Joystick Input Device rollout

The Joystick Input Device driver was designed for the Microsoft Sidewinder joystick, which contains more controls than the standard joystick. You can use this device driver for standard joysticks as well.



X, Y, Z—Specifies which joystick direction drives the animation. (Standard joysticks provide X and Y axes only. The Sidewinder provides the Z axis when you twist the joystick.)

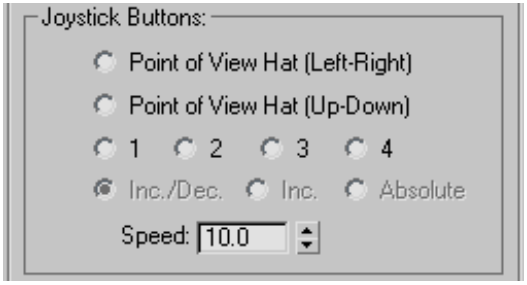
Throttle—On the Sidewinder, this is a slider next to the stick.

Scale—Scales the relative effect of the joystick action to the animation response. Spinner Value=float, 0 to 999,999

Flip—Flips the direction of the response.

Accumulate—When off, the joystick position represents an absolute position, and you're limited to the "rectangle" defined by the limits of the joystick. When the joystick returns to its rest position, the value generated returns to zero. When on, the joystick represents a change in the current position. Moving the joystick forward, for example, can cause an object to start moving, and it will continue to move until you return the joystick to its rest position.

Joystick Buttons group



Point-of-View Hat (Left-Right, Up-Down)—A mini joystick on the tip of the main joystick that specifies the direction of the animation

1, 2, 3, 4—Specifies one of four buttons in the Sidewinder joystick. They work similarly to the Point-of-View Hat, except that each button increases a direction value only while pressed. When you release the button, the value returns to zero (centered).

Inc./Dec—One of three options that are only available when one of the numbered joystick button options is turned on. This option (Increment/Decrement) means that the value is incremented when the button is down, and decremented when the button is up.

Inc—When on, the value increments when the button is down, and stays at that value when the button is released.

Absolute—When on, the assigned parameters jumps to the value set in the Speed spinner when the button is down, and then jumps back to zero when the button is released. Use this for on/off animations, such as blinking lights.

Speed—Controls the rate at which the value changes when using either the Point-of-View Hat or the four buttons. When using a button option and the Absolute option, this specifies the value that's output when the button is pressed. Spinner Value=float, -999,999 to 999,999

Increment Based On Direction group



This group provides controls that let you derive the direction of movement from a Rotation controller. These options are used primarily when you're animating a first-person flythrough such as when you're controlling a camera.

Note: The items in this group are only available when Accumulate is selected in the Joystick Axis group.

Controller—Assigns a Rotation Motion Capture controller where the direction will be derived. Typically, this would be the Rotation controller of the camera you're moving.

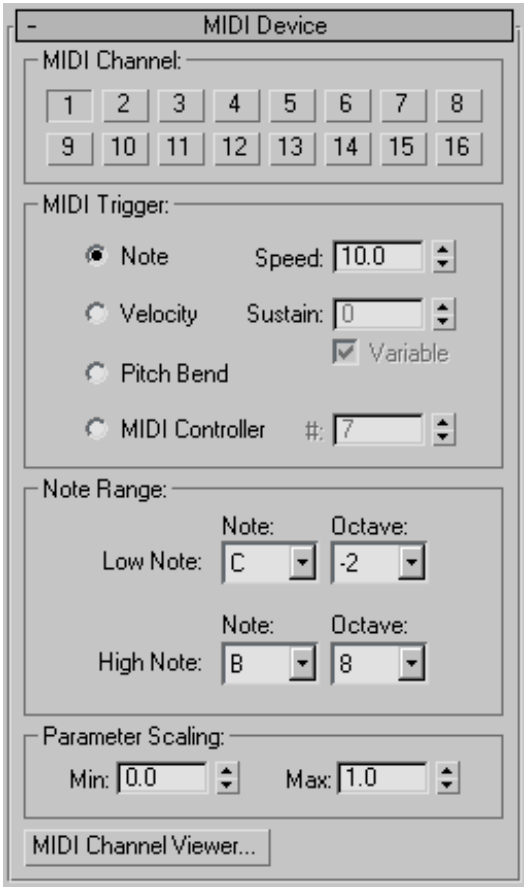
Note: You can only use a Rotation Motion Capture controller here.

Clear—Removes the assigned controller.

Direction X/Y/Z—Specifies the local axis that will be used as the direction. For a Free Camera, for example, this would be Z, because the camera points in the Z direction. However, if you had a car that pointed along its Y axis, you'd use Y.

Component X/Y/Z—Specifies which edit binding to use. Match this to the Edit Binding button in the Device Bindings group. For example, if the Y Edit Binding button is selected, choose the Y Component option.

MIDI Device rollout



MIDI Channel group

This group contains 16 buttons. You can assign a channel to your MIDI device.

MIDI Trigger group

Here, you define the type of MIDI event (message) that will drive the motion. There are four options; Note, Velocity, Pitch Bend and MIDI controller.

Note: Turn on to let the note number or pitch define the output value. The value is derived from where the note falls within the Note Range, specified in the group below. When the note is at the bottom of the range, the value takes on

the Min value specified in the Parameter Scaling group. When the note is at the top of the range, the value takes on the Max value from the same group. Anything in between is interpolated between the Min and Max values. (Note that Min doesn't have to be less than Max.) The generated value will slide around as different keys are pressed. The harder a key is pressed, the faster the value changes.

Speed—Defines how fast the value changes as keys are pressed.

Velocity—Determines the output value based on the velocity that the note pressed. The notes set in the Note Range group specify which notes are valid to press. The value takes on the Min value until a key within the Note Range is pressed. When the key is pressed, the value approaches the Max value based on how hard the key was pressed. (The value actually travels along a parabola toward the Max value.)

Sustain—Defines how long it takes the value to move through the parabola.

Variable—Sustain duration is scaled by how hard the key is pressed.

Pitch Bend—Value is defined by the MIDI instrument's pitch bend knob. The Note Range doesn't apply in this case and is not available.

MIDI Controller—Specifies a note event when you're hooked up to a different type of MIDI controller than the typical keyboard. For example, if you're using a MIDI slider box, you would select the MIDI Controller option, and then use the # spinner to specify the note event for the specify slider.

Note Range group

Turn on Note or Velocity, and then set the note range here. A value is derived from where the note falls within the Note Range.

Low Note—Set a note and octave for the lower range.

High Note—Set a note and octave for the higher range.

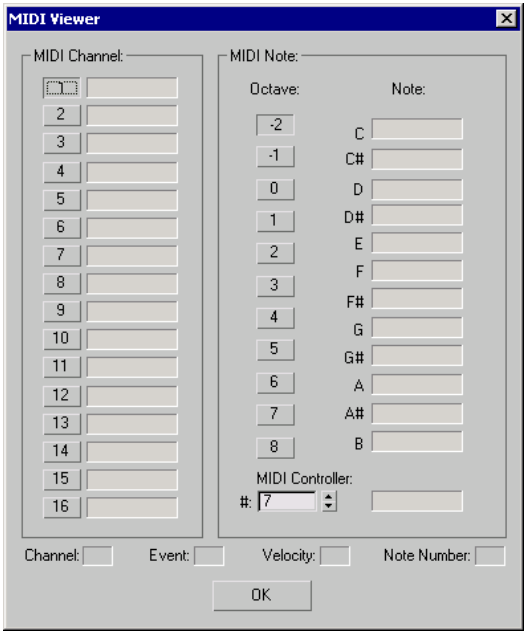
See **Note** parameter in the MIDI Trigger group, above.

Parameter Scaling group

Contains the Min and Max spinners, which specify the range of generated values. See Note and Velocity above.

MIDI Channel Viewer—Displays a dialog that lets you test your MIDI device to see which MIDI channel is receiving events and which notes are being triggered.

Midi Viewer Dialog



MIDI Channel group

Provides a column of 16 buttons and progress bars representing the 16 MIDI channels. Select the channel where you want to view note activity. The channel progress bars light up when any channel has an event.

MIDI Note group

The 11 Octave buttons let you select which octave you want to view. When a note is played in that octave, a corresponding progress bar lights up in the Note column.

MIDI Controller—When using a different type of MIDI controller, such as a slider box, you can specify a note event, and then watch the corresponding progress bar light up when you activate that event. (You can find the correct note number by activating the event while watching the Note Number field in the group below.)

Channel—Displays the currently selected channel. This is one of four text fields that display all of the values being generated by the MIDI device as you activate an event.

Event—Displays the type of MIDI event being sent. For example:

Note On: 9

Note Off: 8

Pitch Bend: 14

MIDI Controller: 11

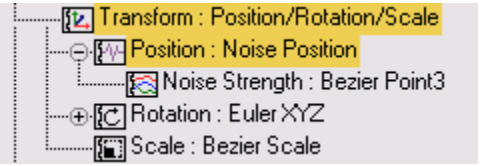
Velocity—Displays the velocity, which has a different meaning depending on the event. For the most common event, a note being pressed, this value represents the velocity at which the key was struck. Other events, however, might generate a continuous value. For example, a pitch bend event transmits the position of the pitch bend.

Note Number—Displays the corresponding note number for the event. When you're using a non keyboard MIDI device, such as a slider box, you can use this to identify the note number of a specific slider.

Noise Controller

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Noise

Graph Editors > Track View – Curve Editor > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Noise



The Noise controller produces random, fractal-based animation over a range of frames. Noise controllers are parametric; they work on a range of frames, but do not use keys.

When you assign a Noise controller in Track View or the Motion panel, it is initially applied to all frames in the current time segment. You can change the range of frames by dragging the Noise range bar in Track View.

- Use Noise whenever you need completely random animation around a given value. For example, use a Noise Rotation controller when you want an object to wobble in place.
- Use Noise as part of a compound List controller to apply noise variations to the result of another controller. For example, use a List controller to combine Noise Position with Bezier Position. The Bezier controller moves the object through your scene while the Noise controller makes the object shake and stray a little from the path.

Procedure

Example: To apply a Noise controller to a rotation track:

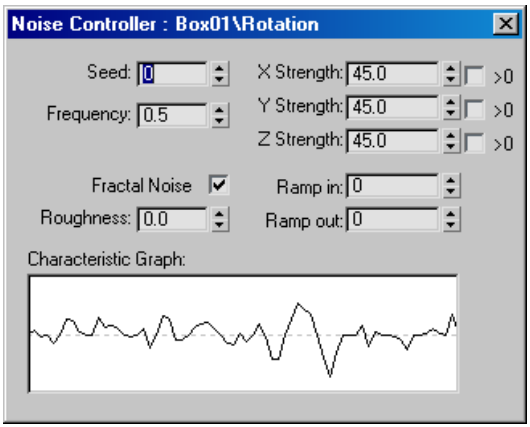
1. Create a box in any viewport.

2. In the Motion panel > Parameters > Assign Controller rollout, select the rotation track in the list window.
3.  Click Assign Controller.
4. Select Noise rotation from the list of controllers.

   The Noise Controller dialog automatically opens.
5. Play the animation.

   The box rotates randomly around all three axes. Because the non-modal Noise Controller dialog is still open, you can adjust parameters and see the results in real time.

Interface



Characteristic Graph—Shows a stylized graph of how changing Noise properties affects the Noise curve.

Seed—Starts the noise calculations. Changing the seed creates a new curve.

Frequency—Controls the peaks and valleys of the noise curve. The useful range is from 0.01 to 1.0. High values create jagged, heavily oscillating noise curves. Low values create soft, gentle noise curves.

Strength fields—Sets the value range for noise output. These values can be animated.

The number and meaning of the fields vary with different parameter data types. Some typical field arrangements are:

- XYZ fields for Position and Point3 keys
- XYZ percentage fields for Scale keys
- A single value field for Float keys

>0 Value constraint—Forces Noise values to stay positive. Each strength field has its own >0 constraint.

When on, the application of the Strength field is changed. Noise values will range from 0 to the value of Strength; most values will hover around Strength/2.

Ramp In—Sets the amount of time Noise takes to build to its full strength. A value of 0 causes Noise to start immediately at full strength at the start of its range. Any other value causes Noise to start at 0 strength and then build to full strength by the elapsed time set in the Ramp In field.

Ramp Out—Sets the amount of time Noise takes to fall to 0 strength. A value of 0 causes Noise to stop immediately at the end of its range. Any other value causes Noise to fall off to 0 strength by the end of its range. The value in the Ramp Out field sets the amount of time before the end of the range that Noise begins to fall off.

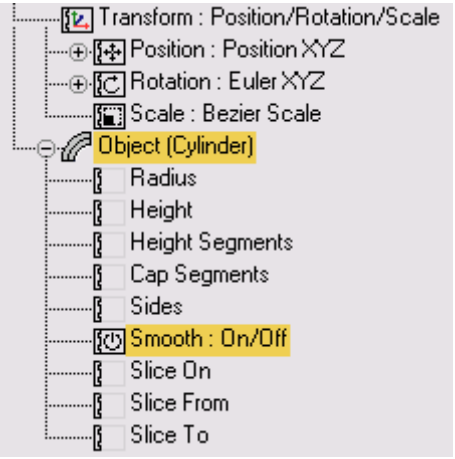
Fractal Noise—Generates noise using a fractal Brownian motion. The main value of using Fractal Noise is that it activates the Roughness field.

Roughness—Changes the roughness of the Noise curve (when Fractal Noise is turned on). Where Frequency sets the smoothness of the overall Noise effect, Roughness changes the smoothness of the Noise curve itself.

On/Off Controller

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > On/Off

Graph Editors > Track View – Curve Editor > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > On/Off



The On/Off controller provides binary on and off control. It is similar to the *Boolean controller (page 2–316)*. For example, you can apply an On/Off controller to the Visibility track of an object.

You can view the curve for this controller in the Function Curve display, but it's not editable. In addition, there's no Properties dialog for this controller.

Note: When you load a scene from version 1.x of the software that contains a Visibility controller, the controller is automatically converted to an On/Off controller.

See also

Boolean Controller (page 2–316)

Procedure

Example: To use the On/Off controller to control an object's visibility:

1. Create a cylinder, and open Track View.
2. In Track View, select the Cylinder track and click Tracks > Visibility Track > Add.

   A Visibility and Transform track are now added.
3. Select the track named Visibility .
4. From the Track View toolbar > Controller menu > Assign and select On/Off in the Assign Controller dialog. Click OK.

   If you change modes to Dope Sheet, note that the entire track region is filled by a blue bar.
5. Click the Add Keys button in Track View — Dope Sheet and add a key at frame 20.

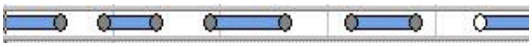
   The blue bar stops at frame 20.
6. Add another key at frame 60.

   The track is now blue again following frame 60.
7. Add some additional keys in the Visibility Track. Each added key reverses the on/off blue pattern following the added key.

Tip: You can select two or more keys and copy them to different positions. This allows you to create even intervals for the On/Off effect.

Interface

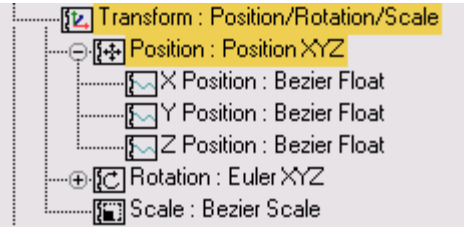
You assign the On/Off Controller in Track View.



The On/Off track displays a solid blue color in frames that are On, and no blue in frames that are Off. When you add a key to an On section of the track, the section following that key is turned off. When you add a key to an Off section, the following section is turned on and colored blue.

Position XYZ Controller

Main toolbar > Curve Editor (Open) > Select a position track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Position XYZ



The Position XYZ controller splits the X, Y, and Z components into three separate tracks (similarly to the *Euler XYZ Rotation controller (page 2–318)*). This provides separate control of the three tracks when referenced from Expression controllers.

The Position XYZ controller assigns three keys (one for each axis), by default.

In earlier versions of 3ds Max, you had to edit the keys manually in order to create explicit axis keys. However, there are now actions available from the *Customize User Interface dialog (page 3–792)* that let you create explicit axis keys *using Set Key Mode (page 2–280)*.

Procedure

Example: To use the Position XYZ and Noise controller together:

1. Create a sphere.
2. In the Motion panel > Parameters > Assign Controller rollout, select the Position track of the sphere.
3. Click Assign Controller and choose Position XYZ.
4. Turn on Auto Key and create three keys that move the sphere on the XY plane.
5. In the Assign Controller rollout, expand the position track and select Z Position.
6. Click Assign Controller, and then select Noise Float.
7. Play the animation. The sphere moves around the keys that were created earlier. The random up and down movement in the Z axis is generated by the noise controller on the Z track.

   This effect can also be created using the List Controller to combine controllers.

Interface

Motion panel > Parameters > Position XYZ Parameters rollout



X—Displays controller properties for X axis transformation.

Y—Displays controller properties for Y axis transformation.

Z—Displays controller properties for Z axis transformation.

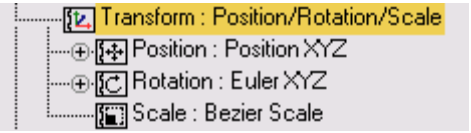
When you select an axis, a Key Info (Basic) and Key Info (Advanced) rollout is available to change values.

A Bezier Float Controller is assigned to each track by default. Use Key Info (Basic) and Key Info (Advanced) to control how the keys behave. See *Bezier Controller (page 2–310)*.

PRS Controller

Main toolbar > Curve Editor (Open) > Select a transform track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Position/Rotation/Scale

Graph Editors > Track View – Curve Editor > Select a transform track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Position/Rotation/Scale



The Position/Rotation/Scale (PRS) controller is the default Transform controller for most objects. Use it for all general-purpose transforms.

Procedure

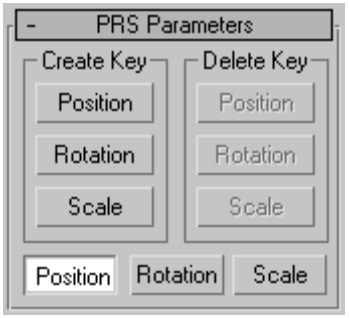
To create PRS Transform keys:

1. Select an object.
2. Click Motion panel > Parameters.
3. Drag the time slider to the frame where you want to place a key.
4. On the PRS Parameters rollout, in the Create Key group, click one of the following:
   - Click Position to create a Position key.
   - Click Rotation to create a Rotation key.
   - Click Scale to create a Scale key.

   If a particular Position, Rotation, or Scale controller does not use keys, then that button is not available in the Create Key group. For example, you cannot create Position Keys if a Noise Position controller is being used.

Interface

Motion panel > PRS Parameters rollout



Create Key/Delete Key—The six buttons in the PRS Parameters rollout let you create or delete a transform key at the current frame. These buttons become active or inactive depending on the existence of a key type at the current frame.

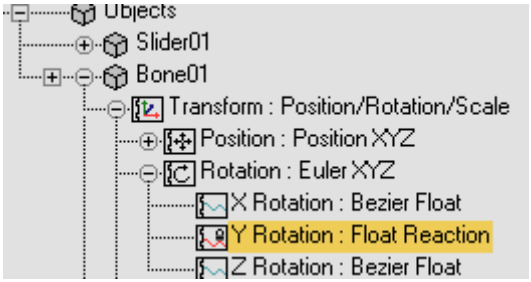
For example, if you're on a frame containing a Scale key, the Scale button is inactive in the Create column, because a key already exists. At the same time, the Position and Rotation buttons are inactive in the Delete column because there are no keys of that type to delete.

Position/Rotation/Scale—Determines the contents of the Key Info rollouts that appear below the PRS Parameters rollout in the Motion panel.

Reaction Controllers

Track View > Highlight a track in the Track View hierarchy. > Assign a Reaction-type controller (e.g., Position Reaction).

Select an object. > Motion panel > Assign Controller rollout > Highlight a track. > Click Assign Controller > Choose a Reaction-type controller (e.g., Position Reaction).



The Reaction controller is a procedural controller that lets a parameter react to changes in any other parameter in 3ds Max. Typically, most of the setup involving Reaction controllers is done with the *Reaction Manager dialog (page 2–361)*. You use the dialog to define a *master*, which is an object that controls other objects, and, for each master, any number of *slaves*, which are objects the master controls. Alternatively, you can assign a Reaction controller directly to a slave object using Track View or the Motion panel (as you would any other controller), and then use Reaction Manager to specify its master and other parameters.

Reaction controllers come in five different types: Position Reaction, Rotation Reaction, Point3 Reaction, Scale Reaction, and Float Reaction. You can assign a Reaction controller to any animatable track in the scene. Reaction is not based on time, but rather on other variables in your scene such as position or rotation.

You can use a Reaction controller to turn on a light as an object nears a given point. Muscles can bulge as an arm bone rotates. A ball can squash automatically as the ball's Z position nears the ground plane. Feet can rotate as their heels are lifted from the floor.

A particle system can be triggered by any given event. Morph target percentages can be controlled by events.

Note: The Reaction controllers replace the Reactor controllers found in previous versions of 3ds Max. If you load a file containing a Reactor controller into this version of 3ds Max, the Reactor controller is converted to a Reaction controller, with all settings and reactions fully intact.

See also

Using Manipulators with Reaction Controllers (page 2–369)

Procedures

Example: To make the position of a sphere react to the position of a box:

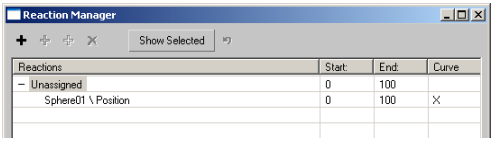
This procedure shows an alternate way to use a Reaction controller. The recommended workflow is to do most of the setup in the Reaction Manager dialog: See *Example: To use the Reaction Manager dialog: (page 2–361)*.

Note: Objects don't need to be animated to use Reaction controllers. This procedure starts with animation to make the reaction easier to see.

1. On the left side of the Top viewport, create a box about 30 units on a side.
2. Animate the box from position **(-100,0,0)** at frame 0 to position **(100,0,0)** frame 100.
3. In the Front viewport, create a sphere slightly above the box. Position the sphere at **(-100,0,50)**.
4.  Go to the Motion panel, and on the Assign Controller rollout, click the Position track to highlight it.

5.  On the Assign Controller rollout, click the Assign Controller button.

6. On the Assign Position Controller dialog, click Position Reaction to highlight it, and then click OK.



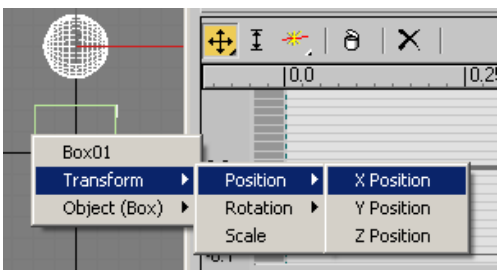
This opens the *Reaction Manager dialog (page 2–361)*. The Reactions list shows that the sphere is assigned as a slave, but no master is assigned. A Reaction controller uses a master's motion to control any number of slaves.

7. Right-click the "Unassigned" line at the top of the Reactions list. From the menu, choose Replace Master.

This places you in a Pick mode where you can choose a motion track from any object in the scene to act as master.

8. In any viewport, click the box (Box01).

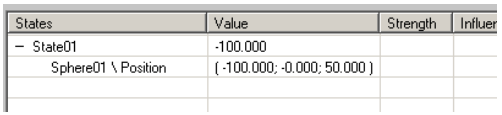
A pop-up menu appears.



9. From the pop-up menu, choose Transform > Position > X Position.

The "Unassigned" text is replaced by a master track labeled "Box01 / X Position".

Also, a new state, State01, appears in the States list.



This is an example of the basis of the control mechanism used by the Reaction controller: For each master/slave combination, you can specify any number of states defined by values for the master and slave tracks. The first state, created automatically when you added the master, specifies that when the box is at -100 on the X axis, the sphere should be at (-100,0,50). In this case, you're controlling three parameters (the sphere position on all three axes) with one (the box position on the X axis). More typically you'd use one-to-one master/slave-parameter ratios.

Note: Reaction Manager also defines a state automatically when you assign a new slave to a master.

Next you'll create a second state that tells the sphere how to move on multiple axes as the box moves on one.

10. Move the time slider to frame 50.

The box moves to the midpoint of its animated trajectory. The sphere remains where it is.

11.  On the Reaction Manager dialog, click the Create State button.

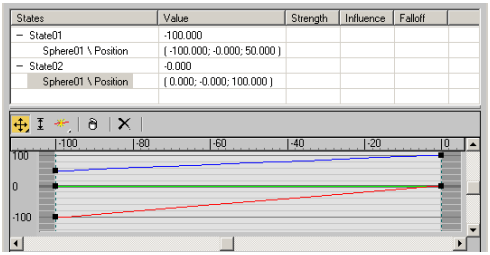
This creates a new state (State02) using the current positions of the box and sphere.

12. Try moving the sphere in the Front viewport.

You can't move the sphere because it's under the full control of its master, the box. Similarly, you can't change its position using the Coordinate Display fields on the status bar. However, you can change the sphere's position using the Reaction Manager's editing tools.

13. On the Reaction Manager dialog, click the Sphere01 / Position entry to highlight it, if necessary, and then click the Edit Mode button.

14. Position the sphere at **(0,0,100)**.



The position updates in the Sphere01 / Position line under State02. Also, the Reaction Manager graph shows the change in values of the slave's (sphere) X and Z position values as the master box's X position changes. On the graph, as in Track View, red=X, green=Y, and blue=Z.

15. Click the Edit Mode button again to turn it off, and then scrub the time slider.

As the box moves between frames 0 and 50, the sphere rises to the position specified in the second state.

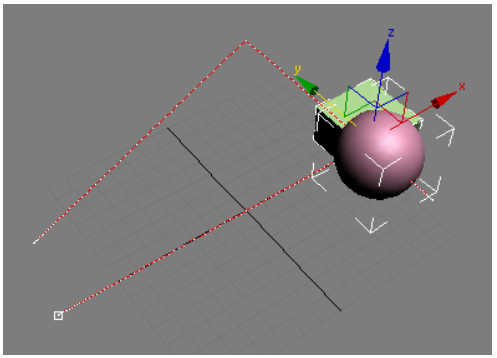
For the third and final state, you'll have the slave move in a different direction over the second half of the master's motion.

16. Go to frame 100, and then create a new state.

17. In the States list, click the new state or its slave track to highlight it.

18. Turn on Edit Mode. Position the sphere at **(0,–100,100)**, and then turn off Edit Mode.

19. Scrub the time slider between frames 0 and 100.



Over the first 50 frames, the sphere moves between the first two states, and over the second half of the animation, the sphere moves between the second and third states.

To conclude the procedure, you'll demonstrate that the sphere is responding only to the box's X position, regardless of animation.

20. Go to frame 0. Select the box, and then press and hold the Alt key and right-click the box.

This opens the Animation quad menu.

21. From the Set (lower-right) quadrant, choose Delete Selected Animation.

This command deletes all animation keys for the current selection.

22. Now move the box in all three dimensions.

Any change in the Y or Z position has no effect on the sphere. However, moving the box between -100 and 100 on the X axis, no matter what the Y or Z position, results in changing the sphere's position as specified in the Reaction controller.

This procedure gives you a hint of the Reaction controller's power. You can find another procedure showing additional aspects of the Reaction Manager in: *Example: To use the Reaction Manager dialog: (page 2–361)*.

Interface

After assigning a Reaction controller, right-click the track and then choose Properties, or choose Animation > Reaction Manager, to open the *Reaction Manager dialog (page 2–361)*.

Reaction Manager Dialog

Animation menu > Reaction Manager
Track View > Highlight a track in the Track View hierarchy. > Assign a Reaction-type controller (e.g., Position Reaction).
Select an object. > Motion panel > Assign Controller rollout > Highlight a track. > Click Assign Controller > Choose a Reaction-type controller (e.g., Position Reaction).

Use the Reaction Manager dialog for setting up and modifying *Reaction controllers (page 2–358)*. Reaction Manager lets you add and delete masters and slaves, define states for reactions, use a graph to view and modify reactions with curves, and more.

Note: If you assign a Reaction controller while Reaction Manager is open, the dialog closes. To reopen it, choose Animation menu > Reaction Manager. If you use Reaction Manager often, it's recommended that you assign a keyboard shortcut to the command.

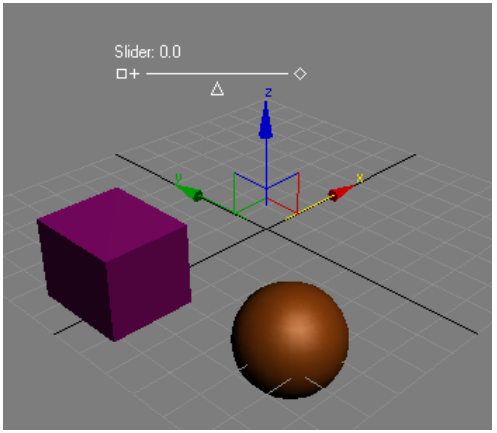
Procedure

Example: To use the Reaction Manager dialog:

The purpose of this procedure is to show you different ways to use Reaction Manager. Before trying it, we suggest that you first follow the introductory procedure in the Reaction Controllers topic: *Example: To make the position of a sphere react to the position of a box: (page 2–358)*.

1. Start or reset 3ds Max.
2. In the Perspective viewport, add a box and a sphere object. Also add a *Slider manipulator (page 2–31)*, and give it the label **Slider**.

   The viewport should resemble the following illustration:

   You'll use the slider to control two slaves separately: Negative values for the box X position, and positive values for the sphere's Hemisphere parameter.
3. On the Create panel > Parameters rollout, set Minimum to **-100.0**.

   The slider's current value is 0.
4. Open the Reaction Manager dialog from the Animation menu and position it so you can see the Perspective viewport.

   At this point, the dialog is empty.
5.  On the upper toolbar, click Add Master, and then click the Slider manipulator in the Perspective viewport.
6. On the pop-up menu that appears, choose Object (Slider) > value.

   In the Reactions list, the master entry Slider01 / value appears.

7.  On the upper toolbar, click Add Slave, and then click the box in the Perspective viewport.

8. On the pop-up menu that appears, choose Transform > Position > X Position.

   In the Reactions list, the slave entry Box01 / X Position appears under the master entry. Also, a new state appears in the States list for the newly added slave track. This simply says that, when the slider value is 0, the box's X position should be what it is currently. However, a reaction needs at least two states to be meaningful. You'll learn how to accomplish this shortly.

   Next, you'll discover a different way to add a slave track.

9. In the Perspective viewport, select the sphere.

10.  On the upper toolbar, click Add Selected.

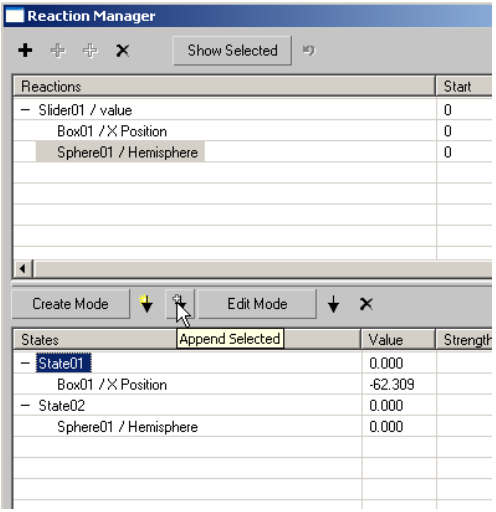
    This causes the pop-up menu to appear at the mouse cursor.

11. From the pop-up menu, choose Object (Sphere) > Hemisphere.

    The sphere's Hemisphere track is added as a slave of the slider, and a new state for the two is added to the States list.

    You now have two states in the States list, but in both cases the state value (that is, the slider value) is the same, so you really need only one. You can resolve this by adding the Hemisphere slave track to the first state and then eliminating the second state.

12. In the Reactions list, click the Sphere01 / Hemisphere entry, and in the States list, click the State01 entry.



13. On the center toolbar, above the States list, click the Append Selected button.

    This adds the Hemisphere track to State01.

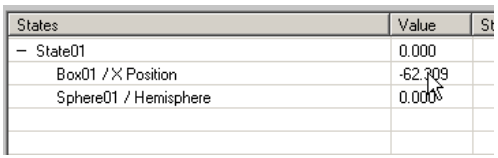
    Now you can delete the second, superfluous state.

14. In the States list, click the State02 entry, and then on the toolbar above the States list, click the Delete State button.

    State02 goes away.

    Currently, the box is where you created it, but you might want it to start out at a different location. However, because the box is a slave of the slider, you can't move it freely on the X axis. You'll use Reaction Manager to set the X position of the box to a specific value.

15. In the States list, position the click and hold the mouse cursor over the Box01 / X Position > Value entry.



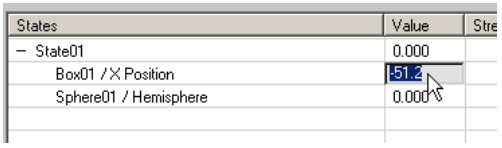
Note: Your value probably differs from that shown in the illustration.

16. Drag upward to increase the value, or downward to decrease the value. Change the value to about **-50.0**.

    You might find it difficult to set an exact value by dragging, which changes only the tenths (and higher) value. You can use the keyboard to specify a more exact value.

17. Double-click the value.

    The value highlights and becomes an keyboard-editable field.



18. Enter **-50** and press **Enter** or **Tab**.

    The value now displays as -50.000.

    You can use the same mouse and keyboard methods to change state (master) values as well.

    Next you'll learn how to use Create Mode to create new states freely.

    First you'll set a new state to move the box leftward for negative slider values.

19. Try moving the box on the X axis.

    You can't do it, because it's slaved to the slider value.

20. Above the States list, click the Create Mode button.

    When Create Mode is on, you can change master and slave properties and use the current properties to create new states.

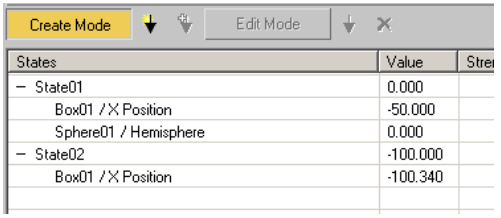
21.  On the main toolbar, click Select And Manipulate, and then drag the slider pointer all the way to the left, so the slider value reads -100.0.

22. In the Reactions (upper) list, click the Box01 / X Position entry to highlight it, and then move the box about -50 units on the X axis, to about X=-100.

    Highlighting the slave entry lets you manipulate the slaved value interactively in the viewport while Create Mode is on. Then, when you click Create State, a new state is created for that slave only. If you highlight the master instead, you can manipulate all slaves freely, but then clicking Create State creates a new state for all slaves. Reaction Manager is a complex tool, so in general it's best to keep the data as simple as possible.

23.  Click Create State.

    Reaction Manager adds a new state, State02, with the state (slider) value at -100.0 and the box at its new X position.



24. Turn off Create Mode and move the slider left and right to see how it affects the box position.

    Next you'll set a new state so the right half of the slider range affects the sphere's Hemisphere value.

25.  Select the sphere and go to the Modify panel.

    The Hemisphere parameter is unavailable, because it's slaved to the slider value.

26. Turn on Create Mode, and then and in the Reactions list, click the Sphere01 / Hemisphere entry.

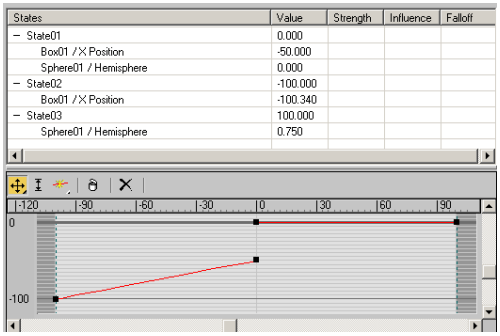
    The Hemisphere parameter is now available.

27. On the Modify panel, set Hemisphere to **0.75**.

28. In the Perspective viewport, drag the slider all the way to the right.

29. Click Create State.

    Reaction Manager adds a third state, State03, with the state (slider) value at 100.0 and the Sphere01 / Hemisphere value at 0.750.



    Also, if you click Slide01 / value in the Reactions list, you can see in the graph showing the slider value on the horizontal axis and the two reactions on the vertical axis. The scale between the two differs greatly, so the Hemisphere reaction on the right side of the graph looks like a straight line, but if you click the Sphere01 / Hemisphere slave entry in the Reactions list, the graph shows a steady rise from 0.0 to 0.75.

30. Turn off Create Mode and then drag the slider in both directions.

    Over the left half of the slider range, the box moves along the X axis, and over the right half of the range, the sphere's Hemisphere value changes.

    To finish the procedure, you'll experiment with using the graph to change a reaction.

31.  If you haven't already, click the Sphere01 / Hemisphere slave entry in the Reactions list, and then click Zoom Extents button below the graph.

32.  Click the Add Point button, and then click the middle of the graph curve (it's the red, diagonal line).

    This adds a new point to the curve and a new state in the States list.

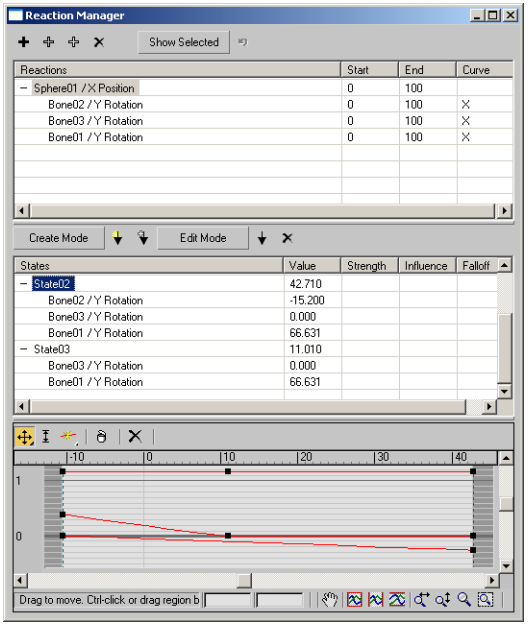
33.  Click the Move button, and then drag the new point upward so it's above the right-hand endpoint.

34. Drag the slider again.

    As the slider moves from 0 to 100, the Hemisphere value increases so much that the sphere disappears, and then decreases so you can see the upper fourth of the sphere.

    This has been a brief tour of some of the capabilities of the Reaction Manager dialog. For further information, consult the reference, below.

Interface



The Reaction Manager dialog interface consists of three parts: the Reactions list, the States list, and

the graph. Each has its own toolbar or toolbars. Each list contains rows and columns; the column headings employ standard Windows functionality. To resize a column, drag the divider to the right of its heading, or auto-size a column to fit its widest entry by double-clicking the right-side heading divider. To change the vertical size of a list section, drag the horizontal divider at the bottom of the list, and to scale all three sections proportionally, drag the bottom of the dialog. You can also resize the dialog by dragging a corner or a side.

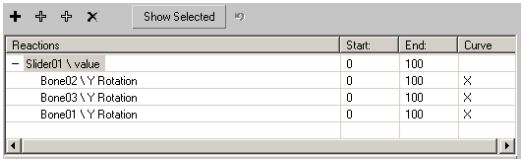
You can use standard methods to highlight list items for further actions. Highlight a single item by clicking it. Highlight multiple consecutive items by clicking the first item and then Shift+clicking the last. Use Ctrl+click to highlight non-contiguous items and to toggle highlighting of a single item.

In most cases, if a cell in the Reactions list or the States list contains a single numeric value, you can change the value by dragging vertically on the cell. Click and hold on the number in the cell, and then move the mouse upward to increase the value or downward to decrease it. The exception to this are the Start and End values for master objects; you can change these values only by modifying a slave's Start or End value. You can also modify a one-value field by double-clicking it and then editing it with the keyboard. If a field contains multiple numeric values, you must manipulate the object directly in the viewports or the appropriate dialog.

The dialog also has a *right-click menu (page 2–368)* for managing reactions and states, available by right-clicking anywhere on the dialog except in the graph section.

Note: A *master* is an object that you can use to control other objects via Reaction controllers, and a *slave* is an object controlled by a master.

Reactions list



The Reactions list displays a hierarchical view of the master-slave relationships of objects with Reaction controllers in the scene. You can use its controls to add and delete masters and slaves, change start and end frames, and toggle usage of the graph curve. These controls are available both on the toolbar above the list and in the list itself.

Add Master—Places you in Add Master mode. Click an object in a viewport, and then use the pop-up menu to choose an animation track to use as master. Alternatively, choose a track in Track View or from the Motion panel > Assign Controller rollout list.

Using the pop-up menu, you can choose only a track at the end of the menu hierarchy; for example, Transform > Position > X Position. If you choose the object name from the top of the menu, Reaction Manager lists it as World Space Position, so that you can use its three-dimensional position in the World coordinate system to control its slaves.

To use a Position or Rotation track for the master, click Add Master and then click the track in Track View or the Motion panel > Assign Controller rollout list. In the latter case, the object must already be selected in the viewports before clicking Add Master in order for it to show up on the Assign Controller rollout list. If using Track View, you can open the Track View dialog before or after clicking Add Master.

Tip: If you use Track View or Motion panel to assign a Reaction controller to an animation track, the object shows up in Reaction Manager without a master ("Unassigned"). To assign a master,

right-click the highlighted "Unassigned" label, and then choose Replace Master from the context menu. Finally, choose the master track.

 **Add Slave—**Places you in Add Slave mode, letting you assign a slave track to the currently highlighted master in the Reactions list. Highlight a master and then click Add Slave. Click an object in a viewport, and then use the pop-up menu to choose an animation track to use as slave. Alternatively, choose a track in Track View or from the Motion panel > Assign Controller rollout list. Available only when a master is highlighted in the Reactions list.

To add multiple slaves, click Add Slave and then press H to open the Pick Object dialog. Highlight the objects to add, and then click Pick. The pop-up menu then appears, prompting you to choose an animation track to use as slave. The same track is used from each object. Alternatively, select the objects and then use Add Selected (see following).

If multiple masters are highlighted when you add a slave or slaves, the slave is added to the first master in the list.

If any states are highlighted in the States list, the new slave or slaves are added to the highlighted states. If no states are highlighted in the States list, Reaction Manager creates a new state with the added slave or slaves.

 **Add Selected—**Lets you add slave tracks to the current master for any selected objects in the viewports. Opens a pop-up menu from which you can choose a track. If multiple objects are selected, this track is added as a slave for each selected object. Available only when a master is highlighted in the Reactions list and at least one object is selected in the viewports.

If multiple masters are highlighted, the slaves are added to the first master in the list.

If any states are highlighted in the States list, the new slave or slaves are added to the highlighted states. If no states are highlighted in the States list, Reaction Manager creates a new state with the added slave or slaves.

 **Delete Selected—**Deletes any items highlighted in the Reactions list.

If you delete a slave, any related state definitions are deleted from the States list, although the states themselves remain. If you delete a master, all of its slaves and any related states are also deleted.

Show Selected—When on, displays data in the Reactions list only for objects selected in the viewports. This lets you restrict display reaction setups to objects of immediate interest, without having to scroll through long lists of parameters.

When off, shows all objects and states, regardless of viewport selection status.

If the master is selected in the viewport, Show Selected shows all slaves in the Reaction list whether or not they're selected.

 **Update Selection—**Refreshes the list contents based on the current selection. Available only when Show Selected is on and the viewport selection has changed.

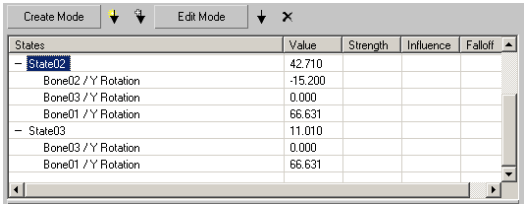
Start/End—These columns show the first and last animation frames for which the master-slave relationship is in effect. These default to the *active time segment (page 3–904)*. Use *standard methods* to change these values.

If multiple slaves assigned to the same master have different Start or End values, the master's Start and End fields are blank.

Curve—This column shows an X when the *graph curve (page 2–368)* is used to modify Strength, Influence, and Falloff. To specify these numerically

in the States list, click the field to toggle the X. Available for slaves only.

States list



A state describes interaction between a master and its slaves. In each state, you specify a value or values for the master track, and corresponding values for each slave. Then, as a master-track value changes in the scene, the slaved values follow along according to the states' specifications.

The States list shows all states for the highlighted reaction in the Reactions list. Any member of a reaction (master and/or slaves) can be highlighted in the Reactions list for its states to appear in the States list. If members of multiple reactions are highlighted, the States list reflects values for the first highlighted reaction only.

You can use tools in the States list to add and modify states and their names. To change a state name, click the text twice, *slowly* (don't double-click). Or, if the state name already highlighted, just click it once. The text highlights and a box appears around it to show that you can now edit it. You can rename only state names, not the indented slave tracks.

You can change values for single-numeric-value slaves in this list by dragging vertically on the value.

The following reference describes changing values interactively in the viewports using the States list controls. Other controls are described in the introduction to this topic.

Create Mode—Use primarily to create new states interactively. To use it, turn on Create Mode, adjust values for the current reaction (master and slaves) in the viewports, and then click Create State. You can repeat this to create any number of new states quickly and efficiently.

The difference between creating a state in or out of Create mode is that, in Create mode you can manipulate the reaction tracks for all objects, whereas with it off you can manipulate the reaction track only for the master.


Create State—Adds a new state to the States list for the current reaction, using the current values for the master and slave parameters.

When you create a state, the software gives it a default name (State##) and lists its slave tracks, indented, under the state name.

The contents of the new state depend on what's highlighted in the Reactions list. If the master is highlighted, a new state specifies values for the state (i.e., master track) and all of its slaves. If one or more slaves are highlighted in the Reactions list, the new state specifies values for the master and only the highlighted slaves.


Append Selected—Adds a slave or slaves to a state in which they don't currently participate. To use, highlight one or more slave tracks in the Reactions list, highlight a state in the States list, and then click Append Selected.

Edit Mode—When on, you can change slave values interactively in the current state by manipulating slave objects in the viewports or by editing their values on a rollout. Available only when at least one item in the States list is highlighted.

If you highlight a slave name, you can edit its value for the current state. If you highlight a state name and then click Edit Mode, you can edit any of its slave tracks. To exit this mode and test the new values, click Edit Mode again.

As you manipulate the object, the value is updated immediately and the new value is displayed in the States list. If slaves in multiple states are highlighted, values are updated for those in the first state only.



Set State—Updates the highlighted state (master track) to its current value. To use this option, highlight the master track in the Reactions list *and* the state in the States list. Manipulate the master object in the viewports, and then click Set State to update the state value. Available only when a state is highlighted in the States list.



Delete State—Deletes any items highlighted in the States list.

Value—This column shows the current value or values for the state: the master parameter(s) on the State## line and the slave parameter(s) on each slave line within a state. You can edit a single value (not groups of three in parentheses) by dragging vertically on the value in the table.

The following three parameters, which appear as columns in the States list, apply to slaves and are normally controlled via the graph curve. They're available only if the slave has Curve turned off in the Reactions list.

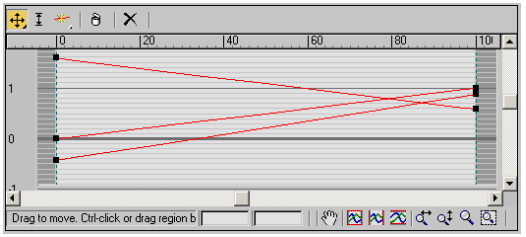
These parameters are based on the fact that states can influence one another, so that you get a blending between them.

Strength—This is a biasing factor that affects a state's relative influence when one or more states overlap. A greater Strength value influences a value towards a state and a smaller one influences the value away from it.

Influence—The distance from the master state's value to the maximum extent of its range of influence.

Falloff—The rate of change in the slave state (its speed) as the master goes from a state to the maximum extent of its range of influence.

Graph



Reaction Manager's graph displays curves that graph the master values on the horizontal axis and slave values on the vertical axis. It shows data of all states for items highlighted in the Reactions list. If a master is highlighted, the graph shows states curves for all its slaves. If one or more slaves are highlighted, but the master isn't, the graph shows data for those slaves only. You can edit data using the graph by dragging points.

The graph controls are similar to those found in other graphs in 3ds Max, such as the *Material Editor Output rollout (page 2–1621)*. The toolbar above the graph offers functions for moving and scaling points on the graph, as well as inserting new ones (that is, new states). The same functions are available by right-clicking the graph, and if you right-click a graph point, you can set it to Corner and two different Bezier types. If you select a Bezier point, you can reshape the curve by moving its handles. You can drag a region to select multiple points, and then edit them together. Use the toolbar under the graph for panning and zooming the graph.

Right-click menu

Reaction Manager's right-click menu provides quick access to a number of commonly used functions. It also contains the Replace Master function, which is unavailable elsewhere in the

dialog interface. To open the right-click menu, right-click anywhere on the Reaction Manager dialog except in the graph section.

Add Master—Places you in Add Master mode. Click an object in a viewport, and then use the pop-up menu to choose an animation track to use as master. Alternatively, choose a track in Track View or from the Motion panel > Assign Controller rollout list.

For more information, see *Add Master*.

Replace Master—Lets you replace the current master. Use the standard method to choose another master as described in *Add Master*.

Use this function to specify a master when you open Reaction Manager by assigning a Reaction controller to a slave.

Note: If you replace a master with multiple states, Reaction Manager deletes all existing states except the default one. In essence, replacing a master is like starting over.

Add Slave—Lets you add a slave to the current master. For details, see *Add Slave*.

Add Selected—Lets you add selected objects as slaves to the current master. For details, see *Add Selected*.

Delete Selected—Deletes any items highlighted in the Reactions list. For details, see *Delete Selected*.

Create States Mode—Lets you create states on the fly by manipulating objects in the viewports. For details, see *Create Mode*.

Create State—Adds a new state to the States list for the current reaction. For details, see *Create State*.

Append Selected—Adds a slave to a state in which the slave doesn't currently participate. To use, highlight a slave in the Reactions list, highlight a state in the States list, and then choose Append Selected.

Set State—Updates the highlighted state (master track) to its current value. For details, see *Set State*.

Delete State—Deletes any highlighted items in the States list.

Edit Slave State Mode—Lets you change slaves' values in the current state. For details, see *Edit Mode*.

Show in Track View—Opens Track View, displaying reaction tracks for highlighted slaves and masters only.

Show in Schematic View—Opens Schematic View, with reaction tracks expanded for highlighted slaves and masters.

Using Manipulators with Reaction Controllers

The Position Reaction and Rotation *Reaction controllers (page 2–358)* let you perform certain functions via manipulators in the viewports.

Procedure

To use manipulators for a position or rotation Reaction controller:

(Point3, scale, and float Reaction controllers don't have manipulators.)

1. Select the object that has the Reaction controller assigned to it.
2.  On the default Main toolbar, click Select And Manipulate.

   The manipulators appear in viewports.
3. In the active viewport, drag or click manipulators to adjust the reaction parameters.

   When you move the mouse over a manipulator, the manipulator turns red to show that dragging or clicking it will have an effect.

Tip: Use the reactor manipulators in a wireframe viewport. (See *Viewport Right-Click Menu (page 3–731)*.) Also, for the Position Reaction controller, you might have to change viewports to see a particular manipulator clearly.

The specific manipulators and their use are described below.

Manipulators for Reactor Controllers

Reaction controllers for position and rotation tracks have graphic manipulators to help you adjust their settings.

Other Reaction controllers, such as point3, scale, and float, do not have manipulators.


Manipulators are visible and usable while the *Select and Manipulate button (page 2–15)* is turned on. This button is on the default *main toolbar (page 3–686)*. When you move the mouse over a manipulator, the manipulator turns red to show that dragging or clicking it will have an effect.

Tip: Use the Reaction controller manipulators in a wireframe viewport. (See *Viewport Right-Click Menu (page 3–731)*.) Also, for the Position Reaction controller, you might have to change viewports to see a particular manipulator clearly.

Position Reaction controller

The Position Reaction controller has a set of manipulators to control (or simply display) different reaction values.

Reaction cycle manipulator: Click this to cycle among the various reactions in the Reactions list. The display of the manipulators and indicators changes to show their settings for the active reaction.

This manipulator is a small yellow circle that appears in the lower-left corner of the active viewport, near the icon for coordinate axes.

Influence manipulator: Drag in a viewport to change the manipulator's size, which changes the reaction's range of influence.

This manipulator is a yellow sphere that turns red when you can drag it. At the center of the Influence manipulator is the Reaction Value manipulator.

Reaction Value manipulator: Drag in a viewport to change the reaction value, which for a position track is an XYZ location.

This manipulator is a light blue sphere that appears at the center of the Influence manipulator.

Object Pivot indicator: Shows the pivot point location of the master object.

This indicator is a dark blue sphere. You can't manipulate it.

Reaction State manipulator: Drag to adjust the reaction state. For position reactions, the reaction state is an XYZ location that the slave object approaches, depending on the animation of the master object and the Influence value.

This manipulator is a small green square.

Reaction State indicator: Shows the reaction state value. This indicator is independent of the square green Reaction State manipulator, which you can drag in viewports. As you drag the green square, the indicator follows it.

This indicator is a small dark-blue dot.

Rotation Reactors

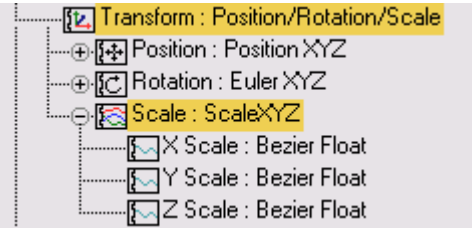
The Rotation Reaction controller has a single manipulator, which controls the reaction's Influence parameter.

Influence manipulator: Drag in a viewport to change the manipulator angle, which changes the reaction's range of influence.

Scale XYZ Controller

Main toolbar > Curve Editor (Open) > Select a scale track in the Track View hierarchy > Track View menu bar > Controller menu > Assign > Scale XYZ

Graph Editors > Track View – Curve Editor > Select a scale track in the Track View hierarchy > Track View menu bar > Controller menu > Assign > Scale XYZ



The Scale XYZ controller has independent float controllers for each scale axis of an object's transform. With three separate tracks for scale, you can create scale keys for each axis independently, change interpolation settings for a single axis, or assign a controller on an axis. For example, after applying a Scale XYZ controller, you could apply a Noise or Waveform controller to an axis to animate that axis independently.

The Scale XYZ controller assigns three keys (one for each axis), by default. In previous versions of the software, you had to edit the keys manually in order to create explicit axis keys. However, there are now actions available from the *Customize User Interface dialog (page 3–792)* that let you create *explicit axis keys (page 2–297)*.

Tip: Open Track View to view and edit the three scale tracks.

Procedure

Example: To assign the Scale XYZ controller to an object:

1. In the Top viewport, create a sphere.
2. Right-click over the sphere and choose Curve Editor from the quad menu.

   Track View - Curve Editor opens with the sphere as the first item in the hierarchy.
3. In the Track View hierarchy, click the Scale track.
4. From the Track View menus, choose Controller > Assign and select Scale XYZ from the Assign Controller dialog.

   Three tracks are displayed and are available to control the object's scale.
5. In the Track View hierarchy, click the Z Scale track, click Controller > Assign from the Track View menus, and then choose Noise Float from the dialog.

   The Noise Controller dialog automatically displays.
6. Click Play. The sphere changes scale along the Z axis, controlled by the Noise controller.

Interface

After applying the Scale XYZ controller to an object, Scale XYZ parameters appear in a rollout on the Motion panel.



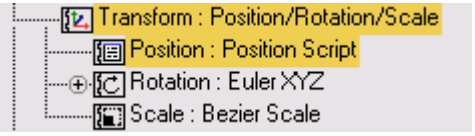
Scale Axis—Adjusts an axis using controls in the Motion panel > Key Info rollout.

An axis should already have a key to adjust. To create keys, either turn on Auto Key and scale the object, or click Scale in the Create Key group on the PRS Parameters rollout.

Script Controller

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Script

Graph Editors > Track View - Curve Editor > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Script



The following Script controllers are available in the software:

- Float Script - float controller
- Position Script - position Point3 controller
- Point3 Script - Point3 controller
- Rotation Script - rotation Quaternion controller
- Scale Script - scale Point3 controller
- *Transform Script (page 2–379)* — matrix3 PRS controller

The Script controllers work similarly to *Expression controllers (page 2–320)*. They provide a Script Controller dialog where a script can be entered for computing the controller value. The primary advantages of Script controllers are:

- They can use all the features of the MAXScript language, including loops, scripted functions, and path names.
- Almost any property of any object in a scene can be used to help compute controller values, including mesh vertices, values of properties at arbitrary frame times, and other non-animatable properties that are not accessible in Expression controllers.
- They can use MAXScript global variables to communicate and coordinate with other controllers and scripts in the software.

Refer to the MAXScript Reference for a complete explanation of this scripting language.

Writing Controller Scripts

3ds Max interprets the text you type into the Script text box as the body of a MAXScript block expression. You can type as many expressions as you want on as many lines as you want, and they are evaluated in turn. The value of the last expression is taken as the controller value. This value must yield the right type for the controller: Float for float, Point3 for position, Quat for rotation, and so on.

Because the text is inside a block expression, you can declare local variables that are visible only within the script and are temporary for one evaluation. You can also declare or access global variables that are shared with all other scripts in MAXScript and hold their values from one evaluation to the next.

A controller is always evaluated by 3ds Max with respect to a specific animation time. This might be the current time slider or incrementing frame time if an animation is playing, or a rendering is under way. In the case of Script controllers, the time being evaluated is used to establish an automatic "at time" context around the controller script, so any properties you access (outside of other explicit “at time” expressions) yield the correct values for the current controller evaluation time. This means you don’t have to do anything special in your scripts to work at the correct time. You can access the evaluation time with the standard MAXScript variable, currentTime. You can also reference scene property values at other times by using "at time" expressions in your scripts, as in regular MAXScript programming.

Assign Nodes and Tracks to Variables

When you need to refer to nodes in your scene or to animation tracks, it is recommended that

you use the Script controller variable toolset to create variables to assign to any particular node or controller track. This way, if you decide to later on rename your scene objects, the script controllers using these objects are preserved because the variables maintain the link to the nodes.

Otherwise, if you assign, for example, a node to a variable manually in the dialog's Expression window, that link becomes broken as soon as you rename that particular node.

Note: It is especially important to assign nodes and tracks to variables with the corresponding buttons when using XRef scenes and objects.

Procedure

Example: To keep an object centered relative to other objects in the scene during an animation:

1. Name the object that should remain centered **foo** and assign a Script controller to its Position track.
2. Enter *foo* in the Name field and click Create.

   The new variable is automatically added to the Variables list.
3. With the variable highlighted, click Assign Node.

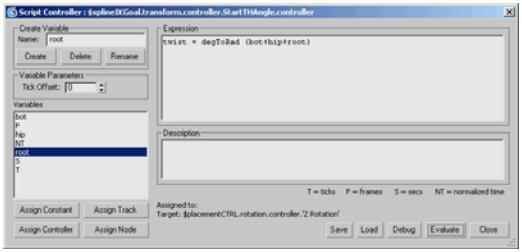
   The Track View Pick dialog opens, listing the contents of your scene.
4. Expand the Objects hierarchy until you locate foo. Highlight it and click OK.

   The foo node is assigned to your variable.
5. Enter the following script in the Script Controller dialog's Expression window:

   ```
   local pos=[0,0,0]
   for o in objects where o != foo do
   pos += o.pos
   pos / (objects.count - 1)
   ```

   This script computes the average position of all objects, except the current one (written as foo here) by setting up a local, iterating over all objects except foo, accumulating a total position vector, and computing the average in the last line, which is the final result of the script.

Interface



Assigning a Script controller automatically opens a Script Controller dialog where you can enter a script. You can open the dialog subsequently by right-clicking the track on the Motion panel or in the Track View hierarchy and choosing Properties, or clicking the Properties button on the Track View toolbar.

Tip: You can resize the dialog by dragging an edge or a corner.

Create Variable group

Name—Lets you enter and edit the name of user variables.

Create—Creates a variable and adds it to the Variables list.

Delete—Removes the highlighted variable from the Variables list. You can also delete a variable by tying its name into the Name field and click Delete.

Rename—Renames the highlighted variable.

Variable Parameters group

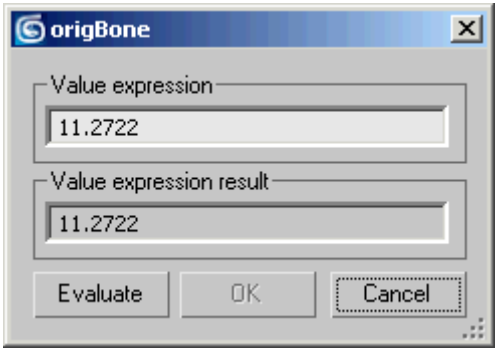
Tick Offset—Specifies a time offset in ticks for the current variable. When the script is evaluated, the variable's value is set from the current time plus the Tick Offset value.

Note: The Time Offset has no effect on variables that are assigned a constant.

Variables List—Lists all available variables in the controller. The following pre-defined constant variables are available in every script Controller and cannot be deleted or renamed:

- **F**—the current time in frames
- **NT**—the normalized time
- **S**—the current time in seconds
- **T**—the current time in ticks

Assign Constant—Opens a dialog which lets you assign a constant to the highlighted variable.



Value expression—Enter any MAXScript value (such as an integer, a float, an array, etc.) or expression.

Value expression result—Displays the results of the expression above.

Evaluate—Analyzes your expression at the current time.

OK—Assigns the value expression result to the current variable.

Cancel—Disregards the current value expression result and closes the dialog.

Assign Track—Lets you assign a controller to the highlighted variable. The controller's value is taken at the current time plus the variable's Tick Offset.

Assign Controller—Lets you assign a track to the highlighted variable.

Assign Node—Lets you assign a node to the highlighted variable.

Expression text field—Enter the script expression here to calculate the controller value.

Description text field—Enter comments regarding the functionality of the controller.

Load/Save—Load and save scripts to text files.

Debug—Opens the Script Controller Debug Window, which displays the value of all variables used in your script.

Evaluate—Evaluates the script expression. The evaluation is computed for the current position of the time slider.

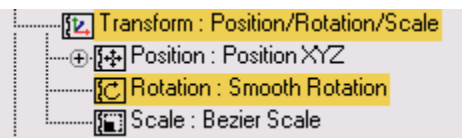
Close—Compiles and checks the controller script for errors. If no errors are found, the dialog is closed. Any problems result in a query box asking whether you want to revert the expression to the original value of the current track and close the dialog (OK), or to return to editing the expression (Cancel).

Smooth Rotation Controller

Main toolbar > Curve Editor (Open) > Select a rotation track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Smooth Rotation

Graph Editors > Track View – Curve Editor > Select a rotation track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Smooth Rotation

Animation menu > Rotation controllers > Smooth



Use Smooth Rotation when you want rotation to have a smooth and natural look. Smooth rotation

functions the same as Linear Rotation, but uses non-adjustable curved interpolation, and has the following characteristics:

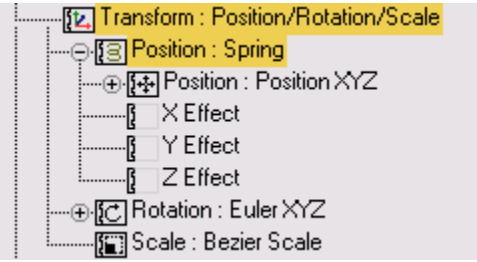
- Displays no controllers, key properties, or function curves
- Move keys in Track View to change timing
- Directly rotate objects in the viewports to change rotation values

Spring Controller

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Spring

Graph Editors > Track View – Curve Editor > Select a track in the Track View hierarchy. > Track View toolbar > Controller menu > Assign > Spring

Animation menu > Position controllers > Spring



The Spring controller adds secondary dynamics effects to any point or object position. The end result is secondary mass/spring dynamics similar to Flex. This constraint adds realism to generally static animations.

When you apply Spring to an animated object, its original motion is preserved and secondary, velocity-based dynamics are applied. You can control the object's mass and drag. When you first apply the controller, it constructs a virtual spring between the object's original position and where it would end up after forces are applied to it.

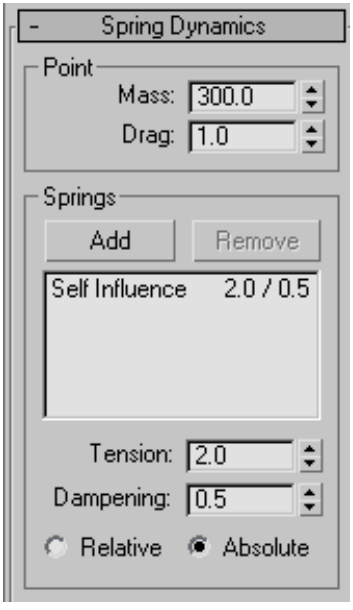
You can adjust spring tension and dampening. Increasing the tension creates a tighter spring, while increasing the dampening smoothes out jitters in the motion. You can add external forces like *Gravity (page 2–73)* and *Wind (page 2–75)* to affect the motion, and also add additional objects whose motions act as springs.

Note: You can also assign the Spring controller to animated vertices in Track View: Open the object's Master Point Controller hierarchy, select the vertices, and assign the Spring controller.

Interface

When you assign a Spring controller to an object, its rollouts automatically appear in the Motion panel. In addition, the rollouts open in a floating dialog. If you close the dialog, you can reopen it by right-clicking the track in the Motion panel Assign Controller rollout or in the Track View hierarchy and choosing Properties, or clicking the Properties button on the Track View toolbar.

Spring Dynamics rollout



Point group

Mass—The mass of the object to which the Spring controller is applied. Increasing the mass causes the "bouncing" spring motion to become more exaggerated.

Drag—Acts as air friction on the spring motion. A low Drag setting results in a greater "bouncing" effect, while a high Drag results in subdued bouncing. Default=1. Range=0 to 10.

Springs group

Add springs to an object using the Spring controller when you want the object to be affected by the motion of other objects in your scene. This effectively creates virtual springs between objects.

Add—Click this button, and then select one or more objects whose motion relative to the spring-controlled object will act as springs on that object. To finish adding springs, click Add again, right-click in a viewport, or press Esc.

Remove—Removes highlighted spring objects from the list.

Note: You cannot remove the Self Influence list entry, which is the spring-controlled object's influence on itself.

(List)—Lists all spring objects by name, and displays each object's Tension and Dampening settings separated by a slash (/). The spring-controlled object's influence on itself appears as Self Influence, and cannot be deleted. However, you can effectively remove the object's influence on itself by setting Tension and Dampening to 0, resulting in no motion.

To set Tension and Dampening for one or more objects in the list, highlight them in the list and then adjust the settings.

Tension—The "stiffness" of the virtual spring between the controlled object and the highlighted spring object(s).

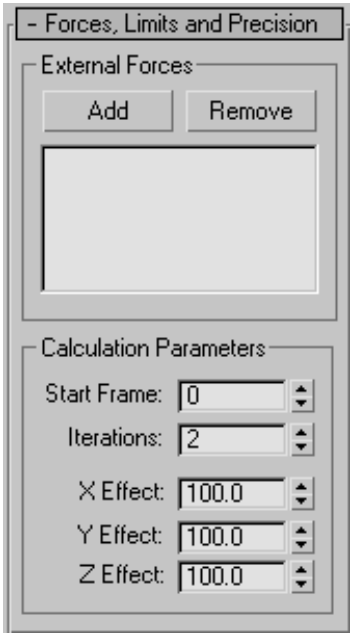
Dampening—Acts as a multiplier of an internal factor that determines how quickly the object comes to rest.

With the Self Influence spring, changing Dampening has the same effect as changing Drag. With other springs, Dampening affects only the movement caused by that spring.

Internally, the dampening value is proportional to the tension, so as you increase the tension and make the solution more stiff, the dampening is increased to maintain system stability.

Relative/Absolute—With Relative chosen, changing the Tension and Dampening settings causes the new settings to be added to the existing values. With Absolute chosen, the new settings replace the existing values.

Forces Limits and Precision rollout



External Forces group

Add—Click this button, and then select one or more space warps in the Forces category that are to

affect the object's motion. To finish adding space warps, click Add again, right-click in a viewport, or press Esc.

Remove—Removes highlighted space warps from the list.

(List)—Lists all spring space warps by name.

Calculation Parameters group

Start Frame—The frame at which the Spring controller first takes effect. Default=0.

Iterations—The accuracy of the controller application. If you get unexpected results, try increasing this setting. Default=2. Range=0 to 4.

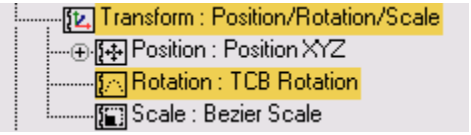
X/Y/Z Effect— These settings let you control the percentage of the effect on the individual world axes. Default=100. Range=0 to 200.

TCB Controllers

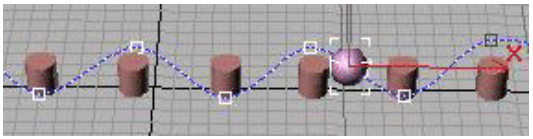
Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > TCB

Graph Editors > Track View - Curve Editor > Select a track in the Track View hierarchy. > Track View toolbar > Controller menu > Assign > TCB

Animation menu > Position, Rotation, or Scale > Quaternion (TCB)



TCB Controllers produce curve-based animation much like Bezier controllers. However, TCB controllers do not use tangent types or adjustable tangent handles. They use fields to adjust the Tension, Continuity, and Bias of the animation.



This controller works well as a Position controller in conjunction with the Motion panel > Trajectories display of an object. With sub-object keys turned on and the Key Info dialog for a key displayed (right-click and select Key Info on the selected key to display Key Info properties), you can change the parameters in the Key Info dialog and see the trajectory path change. This allows for very precise control of an object's trajectory with visual feedback.

Euler XYZ is now the default rotation controller in 3ds Max. The TCB Rotation controller was the default rotation controller for files created in versions 4 and earlier.

Tip: Euler XYZ behaves very differently than TCB rotations. Users upgrading from versions 4 or earlier might have difficulty adapting to the differences between these controllers. If you are used to working with TCB rotations, you can reassign it to be the default rotation controller.

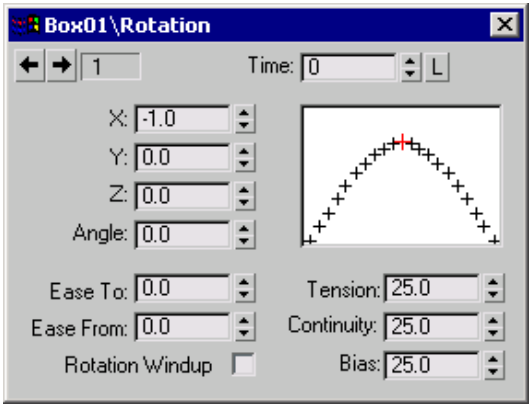
Procedure

To change an Euler XYZ Rotation controller into a TCB Rotation controller:

1. Select an object animated with an Euler XYZ rotation controller.
2. In Motion panel > Parameters > Assign Controller rollout, select the Rotation track in the list window.
3. Click Assign Controller, and then select Quaternion (TCB).

   You can also assign this controller in Track View or from the Animation menu.

Interface



TCB controllers display their properties in a Key Info dialog in Track View and a single Key Info rollout in the Motion panel.

Time field—Specifies when in time the key occurs.

Time Lock—Prevents horizontal key motion in Track View edit modes.

Key Value fields—Stores the animation values for the key. The number and meaning of the fields vary with different parameter data types. Some typical field arrangements are:

- XYZ fields for Position and Point3 keys
- XYZ percentage fields for Scale keys
- A single value field for Float keys

TCB Graph—Charts the effect that changing the controller properties will have on the animation. The red mark at the top of the curve represents the key. The marks to the left and right of the curve represent an even division of time to either side of the key.

The TCB graph is a stylized representation of the animation around a single key. If you want to view the effect of changing TCB properties on the true animation curve, use the Function Curves mode of Track View (works only with position and scale).

The following descriptions refer to both the TCB graph and the function curve as the animation curve.

Note: When you are changing the properties of a selection of multiple keys, the TCB graph is blank unless all properties are equal.

Ease To—Slows the velocity of the animation curve as it approaches the key. Default=0.

High Ease To causes the animation to decelerate as it approaches the key.

The default setting causes no extra deceleration.

Ease From—Slows the velocity of the animation curve as it leaves the key. Default=0.

High Ease From causes the animation to start slow and accelerate as it leaves the key.

The default setting causes no change of the animation curve.

Tension—Controls the amount of curvature in the animation curve.

High Tension produces a linear curve. It also has a slight Ease To and Ease From effect.

Low Tension produces a very wide, rounded, curve. It also has a slight negative Ease To and Ease From effect.

The default value of 25 produces an even amount of curvature through the key.

Continuity—Controls the tangential property of the curve at the key. The default setting is the only value that produces a smooth animation curve through the key. All other values produce a discontinuity in the animation curve causing an abrupt change in the animation. Default=25.

High Continuity values create curved overshoot on both sides of the key.

Low Continuity values create a linear animation curve. Low continuity creates a linear curve

similar to high tension except without the Ease To and Ease From side effect.

The default setting creates a smooth continuous curve at the key.

Bias—Controls where the animation curve occurs with respect to the key. Default=25.

High Bias pushes the curve beyond the key. This produces a linear curve coming into the key and an exaggerated curve leaving the key.

Low Bias pulls the curve before the key. This produces an exaggerated curve coming into the key and a linear curve leaving the key.

The default setting distributes the curve evenly to both sides of the key.

Rotation Windup (TCB rotation only)—When on, rotation keys can be greater than 180 degrees. When off, rotation keys are always less than 180 degrees. This is useful if you want to set a single rotation key to create multiple revolutions of an object. Default=off.

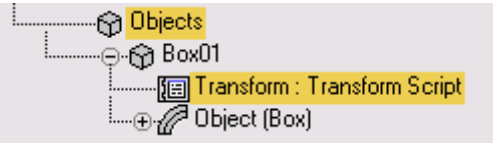
Warning: If you have Rotation Windup on and you insert a rotation key between two previously created rotation keys, you might get unexpected directional shifts. Rotation Windup works best if you set keys using a straight-ahead workflow.

Transform Script Controller

Main toolbar > Curve Editor (Open) > Select a transform track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Transform Script

Graph Editors > Track View - Curve Editor > Select a transform track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Transform Script

Animation menu > Transform Controllers > Script



The Transform Script controller contains all of the information contained in a *Position/Rotation/Scale (PRS) controller (page 2–357)* in one scripted matrix value. Instead of having three separate tracks for position, rotation, and scale, all three values can be simultaneously accessed from one script controller dialog. Because the transform values are defined by a script, they are easier to animate.

The value of the controller script must be a matrix3 value. A matrix3 value is a 4x3 3D transformation matrix. For more information, see the Matrix3 Values topic in the MAXScript reference.

Writing Controller Scripts

The software interprets the text you type into the Script text box as the body of a MAXScript block expression. You can type as many expressions as you want on as many lines as you want, and they are evaluated in turn. The value of the last expression is taken as the controller value. This must yield a matrix3 value.

Since the text is inside a block expression, you can declare local variables that are visible only within the script and are temporary for one evaluation. You can also declare or access global variables that are shared with all other scripts in MAXScript and hold their values from one evaluation to the next.

A controller is always evaluated with respect to a specific animation time by the software. This might be the current time slider or incrementing frame time if an animation is playing, or a rendering is under way. In the case of Script controllers, the time being evaluated is used to establish an automatic "at time" context around the controller script, so any properties you access (outside of other explicit “at time” expressions) yield the correct values for the current controller evaluation time. This means you don’t have to do anything special in your scripts to work at the correct time. You can access the evaluation

time with the standard MAXScript variable, current Time. You can also reference scene property values at other times by using "at time" expressions in your scripts, as in regular MAXScript programming.

Assign Nodes and Tracks to Variables

When you need to refer to nodes in your scene or to animation tracks, it is recommended that you use the Script controller variable toolset to create variables to assign to any particular node or controller track. This way, if you decide to later on rename your scene objects, the script controllers using these objects are preserved because the variables maintain the link to the nodes.

Otherwise, if you assign, for example, a node to a variable manually in the dialog's Expression window, that link becomes broken as soon as you rename that particular node.

Note: It is especially important to assign nodes and tracks to variables with the corresponding buttons when using XRef scenes and objects.

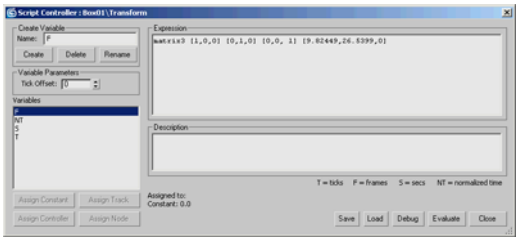
Procedure

To use a Transform Script controller:

1. Select an object.
2. Right-click to open the quad menu.
3. In Transform quadrant, click Curve Editor.
4. Select the Transform track of the selected object in Track View.
5. Select Transform Script from the Assign Transform Controller dialog and click OK.

   The Script Controller dialog opens.
6. Edit the matrix3 value in the Script Controller dialog, and click Evaluate when you are finished.

Interface



After assigning a Transform Script controller, a Properties dialog is available by right-clicking the track in the Track View hierarchy and choosing Properties.

Create Variable group

Name—Lets you enter and edit the name of user variables.

Create—Creates a variable and adds it to the Variables list.

Delete—Removes the highlighted variable from the Variables list. You can also delete a variable by tying its name into the Name field and click Delete.

Rename—Renames the highlighted variable.

Variable Parameters group

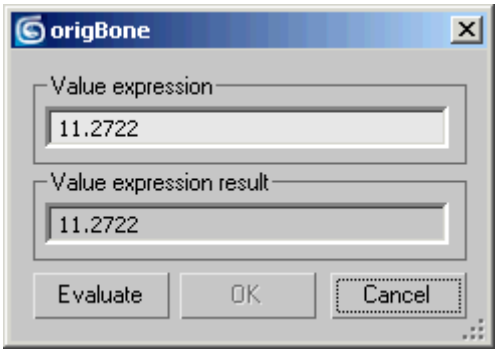
Tick Offset—Specifies a time offset in ticks for the current variable. When the script is evaluated, the variable's value is set from the current time plus the Tick Offset value.

Note: The Time Offset has no effect on variables that are assigned a constant.

Variables List—Lists all available variables in the controller. The following pre-defined constant variables are available in every script Controller and cannot be deleted or renamed:

- **F—**the current time in frames
- **NT—**the normalized time
- **S—**the current time in seconds
- **T—**the current time in ticks

Assign Constant—Opens a dialog which lets you assign a constant to the highlighted variable.



Value expression—Enter any MAXScript value (such as an integer, a float, an array, etc.) or expression.

Value expression result—Displays the results of the expression above.

Evaluate—Analyzes your expression at the current time.

OK—Assigns the value expression result to the current variable.

Cancel—Disregards the current value expression result and closes the dialog.

Assign Track—Lets you assign a controller to the highlighted variable. The controller's value is taken at the current time plus the variable's Tick Offset.

Assign Controller—Lets you assign a track to the highlighted variable.

Assign Node—Lets you assign a node to the highlighted variable.

Expression text field—Enter the script expression here to calculate the controller value.

Description text field—Enter comments regarding the functionality of the controller.

Load/Save—Load and save scripts to text files.

Debug—Opens the Script Controller Debug Window, which displays the value of all variables used in your script.

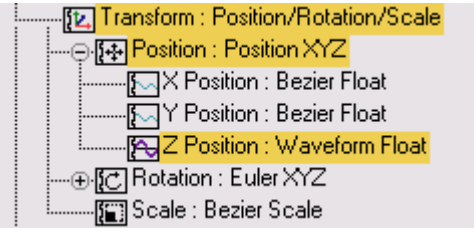
Evaluate—Evaluates the script expression. The evaluation is computed for the current position of the time slider.

Close—Compiles and checks the controller script for errors. If no errors are found, the dialog is closed. Any problems result in a query box asking whether you want to revert the expression to the original value of the current track and close the dialog (OK), or to return to editing the expression (Cancel).

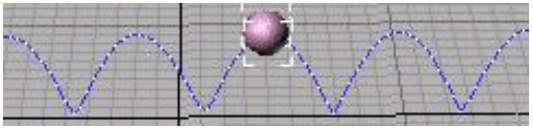
Waveform Controller

Main toolbar > Curve Editor (Open) > Select a track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > Waveform Float

Graph Editors > Track View – Curve Editor > Select a track in the Track View hierarchy. > Track View toolbar > Controller menu > Assign > Waveform Float



The Waveform controller is a float controller that provides regular, periodic waveforms. Created originally to control blinking lights, you can use it on any float value.



To view how the trajectory of an object is being affected by the Waveform controller, turn on Motion panel > Trajectories.

Procedure

To use the Waveform controller to animate the visibility of a sphere:

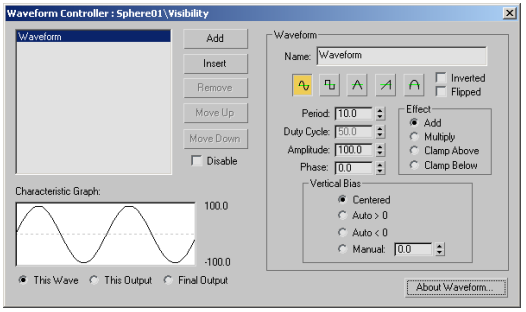
1. Create a sphere.
2. Open Track View, and select the Sphere track.
3. On the Tracks menu choose Visibility Track > Add.
4. Select the Visibility track, click Controller > Assign, and select Waveform Float.
5. Play the animation to see the effect.

To use the Waveform controller to animate the radius of a sphere:

1. Create a sphere.
2. Open Track View — Dope Sheet, and expand the Object (Sphere) track.
3. Select the Radius track, click Controller > Assign, and choose Waveform Float from the Assign Controller dialog.
4. Play the animation to see the effect.

Interface

After assigning the Waveform controller, in Track View, right-click its track to display its properties in the Waveform Controller dialog.



List Window—Displays waveforms in a list.

Add—Adds a new waveform to the end of the list. The default waveform is a sine wave with a period of 10 frames and with an amplitude of 100.

Insert—Inserts a new waveform before the selected waveform slot.

Remove—Deletes the selected waveform. This is not available when only one waveform remains.

Move Up and Move Down—Shifts the selected waveform up and down in the list, allowing the waveform order to be altered.

Disable—Disables the selected waveform.

Waveform group

This group provides control parameters for the currently selected waveform.

Tip: Using the Sphere example, you can observe the waveform in the Function Curve displays of the Track View, and watch changes in the waveform while testing the parameters described below.

Name—Contains the name of the waveform.

Waveform Icons—Specifies five waveform types: Sine, Square, Triangle, Sawtooth, and Half Sine.

Inverted—Flips the waveform vertically.

Flipped—Flips the waveform horizontally.

Period—Sets the number of frames to complete one waveform pattern. Spinner Value=float, 0.01 to 9,999,999

Duty Cycle—For square waves only, specifies the percentage of time the square wave is "on." Default=50, Spinner Value=float, 0 to 100

Amplitude—Sets the height of the wave. Spinner Value=float, 0 to 9,999,999

Phase—Sets the offset of the wave. Spinner Value=float, 0 to 1

Vertical Bias group

A waveform's output value can be altered by adjusting the Vertical Bias parameter. Spinner Value=float, -9,999,999 to 9,999,999

Centered—Centers the waveforms vertical bias about zero. For example, a 100 unit amplitude waveform would range from -100 to 100.

Auto > 0—Shifts the waveform above the zero line.

Auto < 0—Shifts the waveform below the zero line.

Manual—Lets you set the waveform manually by adjusting the spinner.

Effect group

Each waveform has an Effect parameter, which can be one of the following:

Add—Causes the current waveform's value to be added to the previous waveform's output. This is the default.

Multiply—Multiplies the previous waveform's output by this waveform's value.

Clamp Above—Limits the previous waveform's output to be above this waveform's value.

Clamp Below—Limits the previous waveform's output to be below this waveform's value.

Characteristic Graph

Characteristic Graph shows one of three displays, controlled by the radio buttons below the display.

This Wave—Shows the current waveform only, independent of all other waveforms.

This Output—Shows the output of the current waveform, including all previous waveforms.

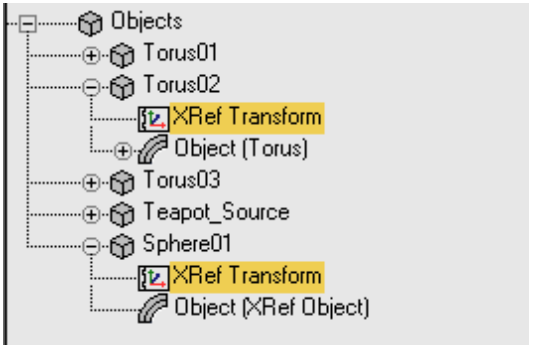
Final Output—Shows the output of all the waveforms in the list. You can also see the final output by turning on the Track View — Curve Editor display for the track which has the waveform controller assigned.

All displays show the output graph of the controller waveforms in a solid dark line, with the zero line drawn in a dotted gray line. Two numeric displays to the right show the high and low range of the waveform curve. The graph automatically scales vertically to fit the waveform output curve, and scales horizontally to two times the period of the longest-period waveform.

XRef Controller

Main toolbar > Curve Editor (Open) > Select a transform track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > XRef Controller

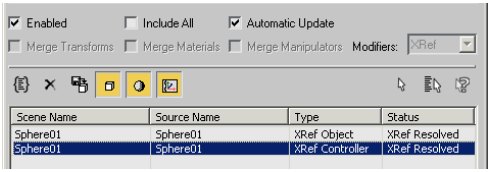
Graph Editors > Track View - Curve Editor > Select a transform track in the Track View hierarchy. > Track View menu bar > Controller menu > Assign > XRef Controller



This controller lets you externally reference any type of *Transform controller (page 2–298)* from another scene file. The XRef controller follows the same concepts and principles as the *XRef objects (page 3–394)* and *XRef material (page 2–1616)*; it can improve upon the collaborative synergy within a production environment by letting you access animation data currently in another scene while someone else updates it. For example, while an animator works locally on a character's walk cycle, a scene integrator can externally reference it into his master scene comprised of several characters. The integrator can now regularly update his master scene with the latest changes from the source scene; if he notices a problem, the animator can make the appropriate modifications, which reflect in the master scene.

You can create an XRef controller with or without an XRef object:

- You can externally reference a source object's animation without the object itself by *assigning an XRef controller to your current selection (page 2–384)*.
- You can create an XRef object, which automatically nests the source object's Transform controller in an XRef controller. The XRef Objects dialog's *XRef Entities list (page 3–402)* lists both XRefs as separate entries.



Note: The Merge Transforms option must be off to externally reference a controller; otherwise, the controller merges with the current scene.

To remove an XRef controller, you can either right-click the entry in the XRef Entities list and choose Merge in Scene from the contextual menu, or assign a new Transform controller to your object.

Note: Replacing the XRef controller with another type removes any animation it contains.

Although XRef controllers are read-only in your master scene, you can *add a local offset to the object's animation (page 2–385)*.

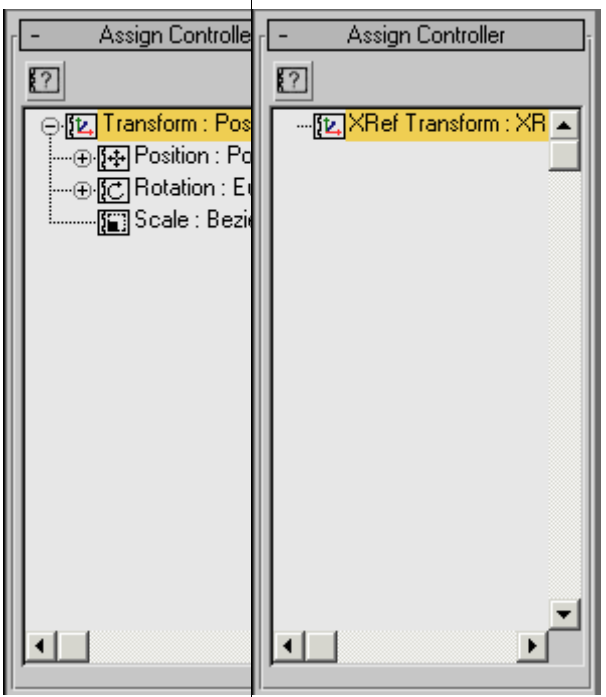
Procedures

To assign an XRef controller using the Motion Panel:

1. Select an object in the current screen.
2. On the Motion Panel, click Parameters, and open the Assign Controller rollout.
3. Highlight the Transform track in the Assign Controller list.
4. Click the Assign Controller button, and then select XRef Controller from the Assign Transform Controller dialog.
5. In the Open File dialog, navigate to the file from which to externally reference the transform controller and then click Open.
6. Use the XRef Merge dialog to designate the object whole transform controller you want to XReference.

   The software references the scene object's transform controllers from the external object's controllers.

   Tip: You can also XReference a different controller from either another file or another object via the Parameters rollout of the XRef controller.



Object's Transform Position highlighted. *Object's Transform Position XReferenced.*

To add and reset a Transform offset to an XRef Object's Animation

1. Create a sphere at [0,0,0] and animate it from [0,0,0] to [5,5,5]
2. Save your scene, then choose File > Reset.
3. Choose File > XRef Objects.
4.  In the XRef Objects dialog, make sure Merge Transforms is off, then click the Create XRef Record From File and select the scene you just saved.

   The sphere appears in your master scene at [0,0,0].
5.  Move the sphere to [10,10,10].

   The sphere now translates from [10,10,10] to [15,15,15].
6. Right-click the XRef Controller entry in the XRef Entities list of the XRef Objects dialog.
7. From the contextual menu, choose Reset PRS Offset.

   This removes the sphere's local offset, which now returns to its original [0,0,0] position.

Interface



These controls, similar to those on the *XRef Object rollout (page 3–414)*, let you change the file path, file name, and object name of the source of the XRef controller.

 **Hightlight corresponding XRef Record in the XRef Object dialog—**Opens the *XRef Objects dialog (page 3–397)* and highlights the corresponding XRef record within.

File Name controls

File Name field—Displays the path and file name of the scene containing the source controller. You can edit this to point to a different path and file.

File Name display—Displays the file name only, without the path.

 **Path button—**Opens a file dialog from which you select a new file (either .max or .chr).

Object controls

Object Name field—Displays the name of the source object pointed to in the source file which holds the source controller.

Object Name display—Displays the file name of the source object.

Path button—Opens an *XRef Merge dialog (page 3–406)* pointing to the scene in the XRef File Name field. Here, you can specify a different object to be used as the XRef object.

status line—Displays status information regarding the XRef controller:

- "Unresolved XRef, File Not Found": Cannot find the scene file containing the specified object name.
- "Unresolved XRef, Object Not Found": Cannot find the object in the specified scene.
- "XRef Resolved": Both scene and object match.

Controller Subdialogs

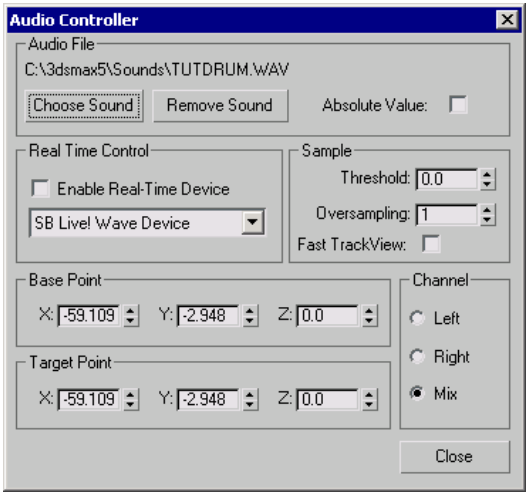
Audio Controller Dialog

Track View or Motion panel > Highlight a track using the Audio controller in the Track View hierarchy. > Right-click the track. > Properties

Track View or Motion panel > Double-click a track using the Audio controller in the Track View hierarchy. > Properties

Adding a sound track to your animation allows you to synchronize animation to sound. Add a sound track to lip-sync a character speaking, for example.

Interface



Audio File group

Use this group to add and remove sound files from the controller, and to adjust amplitude.

Choose Sound—Displays a standard file selector dialog. You can choose WAV and AVI files.

Remove Sound—Removes any sound file associated with a controller.

Absolute Value—Controls the interpretation of sound amplitude. The value returned by the Audio controller is sample amplitude divided by maximum amplitude.

Value=(sample amp.)/(maximum amp.)

When Absolute Value is turned on, the maximum amplitude equals the maximum sampled amplitude from the waveform. This ensures that potential for the output value reaches the target value.

When off, the maximum amplitude equals the maximum potential amplitude of the waveform. Output reaches the target value only if the waveform reaches its maximum potential amplitude.

The maximum potential amplitude for an 8-bit file is 128; for a 16-bit file, it's 32768.

Real Time Control group

Use this group to create interactive animation that's driven by sound captured from an external audio source, such as a microphone. Use these options only for interactive presentations. You can't save the real-time sound or save the animation produced by the controller.

Enable Real Time Device—Sets whether sound is captured from an external audio source. This option is inactive if a sound capture device is not installed in your system.

When on, any selected audio file is ignored, and the controller uses sound captured by the selected device.

When off, the controller uses the selected audio file.

Real Time device list—Displays all available real-time sound devices installed in your system. Select the device you want to use for real time sound capture.

Sample group

This group contains controls to filter out background noise, smooth out the waveform, and control display in Track View.

Threshold—Sets the bottom cut-out level as a percentage of total amplitude. Any amplitude below the threshold drops to 0.0.

Threshold range is from 0.0 to 1.0.

A Threshold of 0.0 has no effect on amplitude output values.

A Threshold of 1.0 drops all amplitude output values to 0.0.

You can use low threshold values to filter out background noise from the controller.

Oversampling—Smoothes the waveform. Multiple samples are averaged to remove peaks and valleys. Enter a number in the Oversampling field to calculate the average

Fast Track View—Controls the display of oversampling.

When turned on, oversampling is ignored for the Track View display.

When turned off, oversampling is applied to the Track View display. High oversampling values can slow the display of the waveform.

Base & Target Scale groups

Here you enter minimum and maximum parameter values returned by the controller. The fields you see vary by the type of parameter using the Audio controller.

Float parameters are displayed in the Controller Range group with Min. and Max. fields.

Vector parameters (3 components), such as Transforms, display Base and Target groups with fields for X, Y, and Z values.

Base Scale—Defines the float value or X, Y, and Z values returned for an amplitude of 0.0.

Target Scale—Defines the float value or X, Y, and Z values returned for maximum amplitude.

Channel group

With this group you select which channel drives the controller output value. These options are only available if you have chosen a stereo sound file.

Left—Uses the left channel amplitude.

Right—Uses the right channel amplitude.

Mix—Combines the two channels so that the returned amplitude is the greater value of either channel.

Attach Controls Dialog (Block Controller)

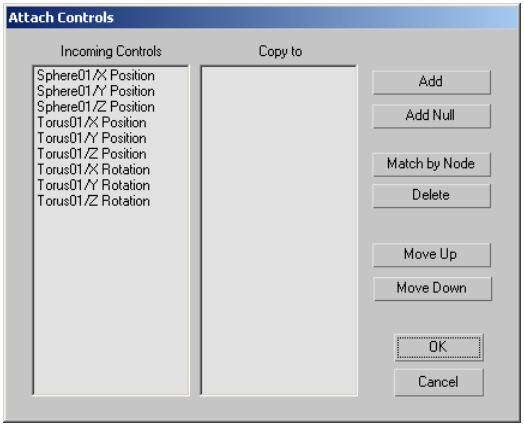
Track View > Global Tracks > Block Control > Click and then right-click MasterBlock track. > Properties > Load > Choose a file. > OK > Attach Controls dialog

Track View > Global Tracks > Block Control > Double-click MasterBlock track. > Properties > Load > Choose a file. > OK > Attach Controls dialog

This dialog is displayed when you click Load in the *Master Block Parameters dialog (page 2–390)*. Tracks in blocks previously saved can be mapped to tracks in the current scene.

The Incoming Controls list on the left contains all the incoming tracks. The Copy To list on the right contains all the tracks in the current scene that you will attach to.

Interface



Add—Selects tracks in the current scene in the *Track View Pick dialog (page 2–392)*.

Add Null—Allows a space to be taken if you don't want to use a particular track.

Match by Node—Select tracks on the left that you want to match. Click Match By Node. Then select a group of tracks that you want to match to. The system attempts to match the controls based on their names.

Delete—Deletes a track from the list.

Move Up, Move Down—Move tracks up and down to align them before attaching them.

OK—Imports the tracks and closes the dialog.

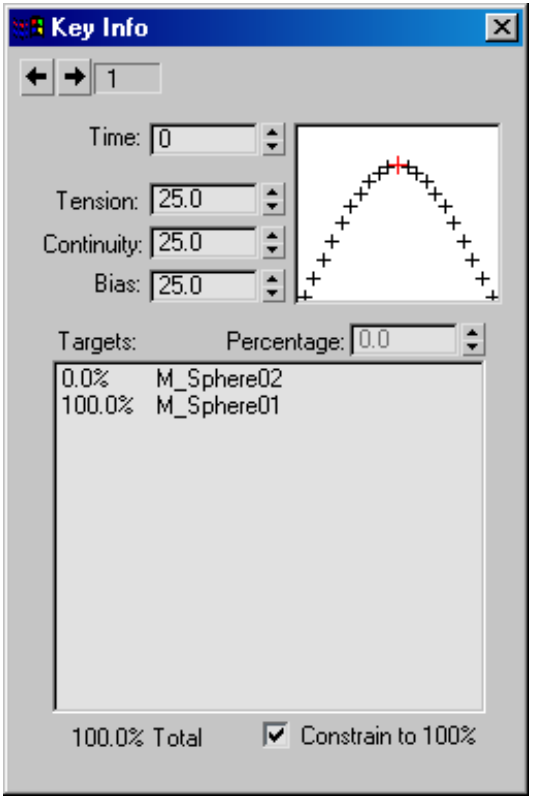
Cancel—Cancels the Load operation.

Barycentric Morph Controller Key Info Dialog

Select a Morph object > Track View - Dope Sheet > Right-click a morph key. > Barycentric Morph Controller Key Info dialog

You can change morph target weighting using controls in the Barycentric Morph Controller Key Info dialog.

Interface



Current Key—Specifies the current key that you're adjusting.

Time—Specifies where, in time, the current key is placed.

Tension/Continuity/Bias—Adjusts the TCB parameters of the key spline interpolation.

TCB Spline—Displays the key interpolation spline.

Targets—Lists all morph targets for the object, along with the percentage of their influence at the current key. The total percentage of all targets is displayed below the window. You can select any target in this window, and then adjust its percentage of influence using the spinner at right.

Percentage—Sets the percentage of influence for the target selected in the list window. Percentage values can be negative as well as greater than 100 percent.

Constrain to 100%—Limits total weight of all targets to 100 percent. As you increase or decrease the percentage of a selected target, the other targets adjust accordingly. The adjustment takes the form of balancing all of the target percentages so that their relative weights remain the same.

When this check box is turned off, you can adjust the weight of each target without affecting the others, and create a total that's greater or less than 100 percent. Totals greater than 100 percent cause the morph object to scale up in size, while totals less than 100 percent cause the morph object to scale down. Note that this check box is global, and not confined to a single key or track.

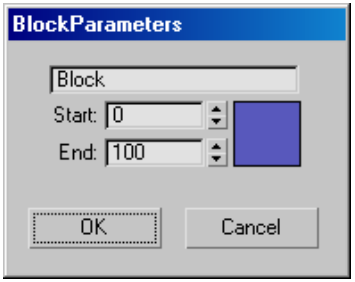
Block Parameters Dialog (Block Controller)

Track View > Global Tracks > Block Control > Click and then right-click MasterBlock track. > Properties > Master Block Parameters dialog > Add > Choose tracks in Track View Pick dialog > Click OK. > Block Parameters dialog

Track View > Global Tracks > Block Control > Double-click MasterBlock track. > Properties > Master Block Parameters dialog > Add > Choose tracks in Track View Pick dialog > Click OK. > Block Parameters dialog

After clicking Add in the Master Block Parameter dialog and selecting tracks in the Track View Pick dialog to create a Block, clicking OK in the Track View Pick dialog displays the Block Parameters dialog.

Interface



Name Field—Names the block.

Start, End—Set the range of the block in frames.

Color—Selects a color for the Block. Displays a color selection dialog.

OK—Saves the parameters and closes the dialog.

Cancel—Closes the dialog without saving the parameters.

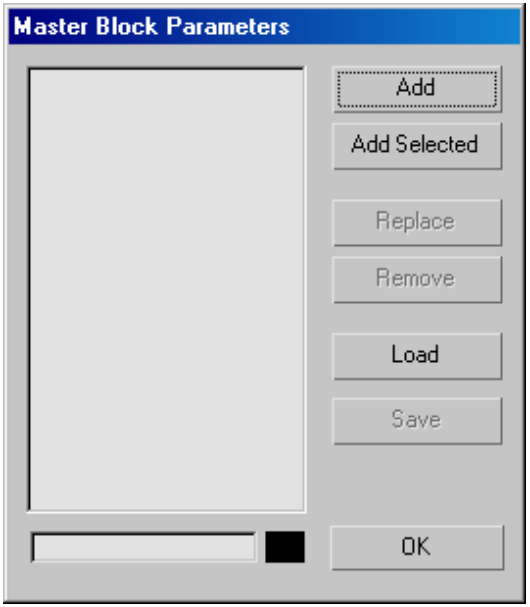
Master Block Parameters Dialog (Block Controller)

Track View > Global Tracks > Block Control > Master Block > Click and then right-click the Master Block track. > Master Block Parameters dialog

Track View > Global Tracks > Block Control > Master Block > Double-click the Master Block track. > Master Block Parameters dialog

After assigning a Master Block, right-click the Master Block track to display the *Master Block Parameters dialog (page 2–390)*. This is the first step in creating a block.

Interface



Add—Displays the Track View Pick dialog. In the Track View Pick dialog, select tracks to include in a Block. Valid tracks are displayed in darker text.

Add Selected—Creates a block using the tracks already selected in the Track View hierarchy. Any invalid tracks are ignored.

Select the tracks (include the Master Block track), and then right-click Master Block in the Track View hierarchy.

Replace—Replaces the currently selected block.

Remove—Remove the selected block from the list.

Load—Loads a block from disk. Displays the Attach Controls dialog.

Map incoming tracks to tracks in your scene using controls in the Attach Controls dialog.

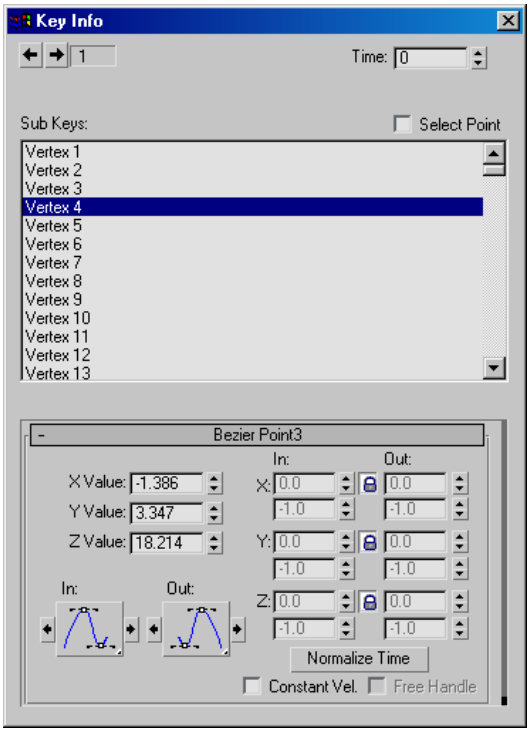
Save—Saves the current block to disk.

Master Track Key Info Dialog (Master Point Controller)

Track View - Dope Sheet > Right-click a master track key in the Dope Sheet Key window. > Master Track Key Info dialog

With Track View in Dope Sheet mode, right-click a green master key of a *Master Point controller (page 2–346)* to display the Master Track Key Info dialog.

Interface



[key arrows]—Moves to the next or previous key.

Time—Moves the master key to a different frame or time.

Sub Keys list window—Displays all the animated sub-object keys at a particular frame or time.

Select Point—Select a node in the List Window and the corresponding node in the viewports is selected, if the corresponding sub-object level is active.

Key Frame Properties—Displays and lets you adjust interpolation and position. See *Bezier Controller (page 2–310)*.

Slave Parameters Dialog (Block Controller)

Track View > Click and then right-click a slave track in the Track View hierarchy. > Properties > Slave Parameters dialog

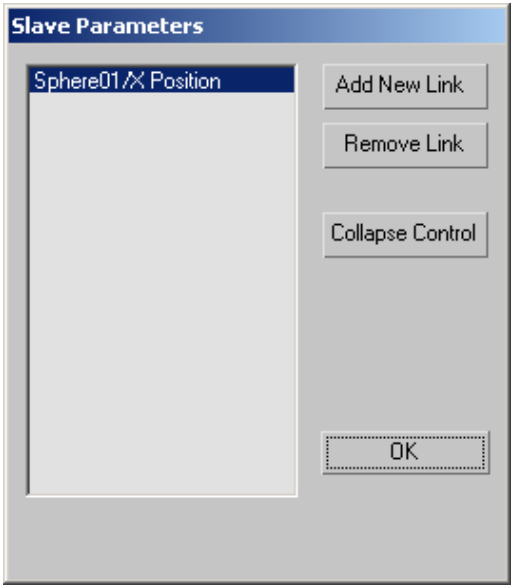
Track View > Double-click a slave track in the Track View hierarchy. > Properties > Slave Parameters dialog

Every time you create a block, all tracks within the block are assigned a Slave controller, which allows the Master Block to transfer key data.

The Slave controller can be assigned without being connected to a Master Block controller. Controls in the Slave Parameters dialog allow you to attach the slave to a master. If the slave is not attached to a master, a dialog is displayed listing the available master controls you can attach to. The *Track View Pick dialog (page 2–392)* is displayed, listing the tracks you can add.

Right-click a Slave controller in the Track View hierarchy to display the Slave Parameters dialog. The list window displays any assigned tracks.

Interface



Add New Link—Adds a link. Displays the Add New Link dialog with available tracks. Click a track and then click OK.

Remove Link—Removes the highlighted link.

Collapse Control—Collapses the Slave controller to a standard controller.

Existing keys are copied to the new controller.

OK—Closes the dialog.

Track View Pick Dialog (Block Controller)

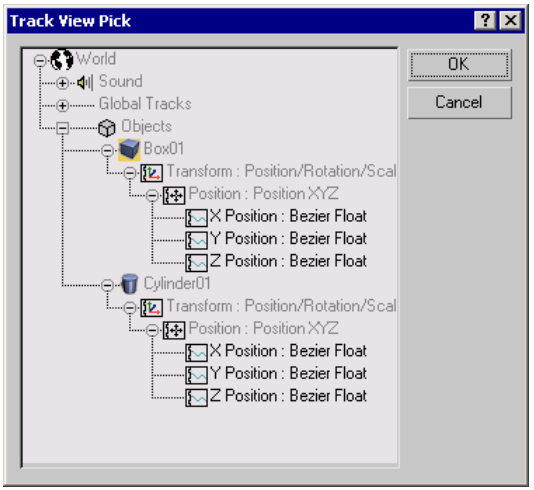
Track View > Global Tracks > Block Control > Click and then right-click MasterBlock track. > Properties > Master Block Parameters dialog > Add > Track View Pick dialog

Track View > Global Tracks > Block Control > Double-click MasterBlock track. > Properties > Master Block Parameters dialog > Add > Track View Pick dialog

This dialog is displayed when track selection is necessary, by clicking Add in the *Master Block Parameters dialog (page 2–390)*, for example.

Select tracks in the dialog that you want to include in a Block. Valid tracks are darker.

Interface



Cancel—Exits the Track View Pick dialog with no changes.

If tracks are highlighted in the Hierarchy list, then the *Block Parameters dialog (page 2–389)* is displayed when you click OK, to allow you to set the duration for the Block.

Animation Constraints

An animation constraint is a special type of controller that can help you automate the animation process. You can use constraints to control an object's position, rotation, or scale through a binding relationship with another object.

A constraint requires an animated object and at least one target object. The target imposes specific animation limits on the constrained object.

For example, to quickly animate an airplane flying along a predefined path, you can use a *Path constraint (page 2–398)* to restrict the airplane's motion to a spline.

You can use keyframe animation to toggle the constraint's binding relationship with its targets over a period of time.

Common uses for constraints include:

- Linking one object to another over a period of time, such as a character's hand picking up a baseball bat
- Linking an object's position or rotation to one or several objects
- Keeping an object's position between two or more objects
- Constraining an object along a path or between multiple paths
- Constraining an object to a surface
- Making an object point toward another object
- Keeping an object's orientation in relation to another

There are seven types of constraints:

- *Attachment constraint (page 2–393)* attaches an object's position to a face on another object
- *Surface constraint (page 2–396)* restricts an object's position along the surface of another object
- *Path constraint (page 2–398)* restricts an object's movement along a path
- *Position constraint (page 2–401)* causes the constrained object to follow the position of another object
- *Link constraint (page 2–403)* links the constrained object from one object to another
- *LookAt constraint (page 2–406)* constrains an object's orientation so that it's always looking at another object
- *Orientation constraint (page 2–409)* causes the rotation of the constrained object to follow the rotation of another object

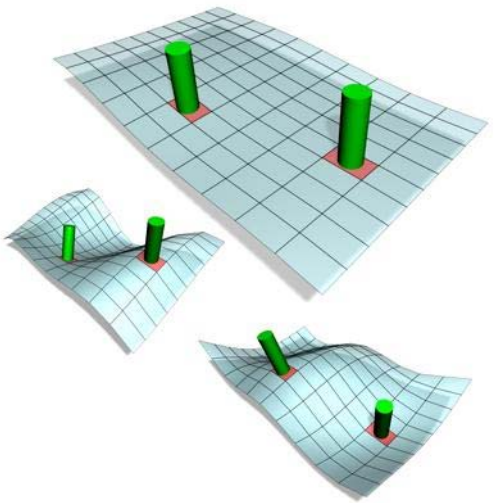
Tip: You can use Schematic View to see all the Constraint relationships in a scene.

Using Constraints with Bones

Constraints can be applied to bones as long as an IK controller is not controlling the bones. If the bones have an assigned IK controller, you can only constrain the root of the hierarchy or chain.

Attachment Constraint

Animation menu > Constraints > Attachment Constraint



Attachment constraints keep the cylinders on the surface.

The Attachment constraint is a position constraint that attaches an object's position to a face on another object (the target object doesn't have to be a mesh, but must be convertible to a mesh).

By keying different attachments over time, you can animate the position of an object over the irregular surface of another object, even if that surface is changing over time.

Procedures

Example: To attach a cone to a bending cylinder:

1. In the Perspective viewport, create a cylinder with a radius of **20**, a height of **30**, and **10** height segments.
2. In the Perspective viewport, create a cone with a radius 1 of **15**, a radius 2 of **5**, and a height of **30**.
3. Select the cylinder, apply a Bend modifier, and set the bend angle to **-70** degrees.
4. Turn on Auto Key, move to frame 100, and set the bend angle to **70** degrees.

   The cylinder bends from one direction to the other over 100 frames.
5. Turn off Auto Key.

Example continued: To assign the Attachment constraint and adjust the cone:

1. Select the cone.
2.  On the Motion panel, open the Assign Controller rollout, click the Position track, click Assign Controller, and choose Attachment.

   The cone moves to the origin of the scene, and the Attachment Parameters rollout is displayed.
3. Click the Pick object, and then click the cylinder.

   The name of the cylinder is displayed above the Pick Object button.

   Go to frame 0. Arc Rotate the Perspective viewport so you can see the top surface of the cylinder.
4. Click Set Position, and click and drag over the faces on the top surface of the cylinder.

   The cone jumps to the top of the cylinder. As you drag the mouse, it jumps to whichever face you drag over.
5. Release the mouse when the cone is on the top surface of the cylinder.

Example continued: To adjust the position of the cone relative to the face:

1. Drag in the face display window to position the red x relative to the triangle representing the face. (Because of the radial arrangement of cap faces in a cylinder, the upper-left corner of the displayed triangle is the center of the cylinder cap. You can turn off Edges Only for the cylinder to see this.)
2. Adjust the A and B spinners to move the cone across the surface of the face.
3. Drag the time slider to various frames.

   As the cylinder bends back and forth, the cone remains attached to its upper surface. Continue adjusting the A and B spinners and dragging in the face display window to adjust the cone's position.
4. Remember the number in the Face spinner, and then lower the spinner value until the cone leaves the cylinder cap and begins jumping around various areas of the cylinder.

   The Face spinner specifies which face the cone is attached to. As you change its values, the cone moves to different faces on the cylinder.
5. Re-enter the original value in the Face spinner to return the cone to the top of the cylinder.
6. Play the animation.

   The cylinder bends back and forth with a cone attached to its upper cap.

Interface

Parameters for this constraint are available on the Motion panel after the controller has been assigned.

Attach To group

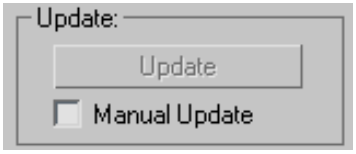


Object Name text—Specifies the target object to which the source object is attached.

Pick Object—Selects and picks the target object in the viewports for the attachment.

Align to Surface—Fixes the orientation of the attached object to the face where it's assigned. When this is turned off, the orientation of the attached object is not affected by the orientation of the face on the target object.

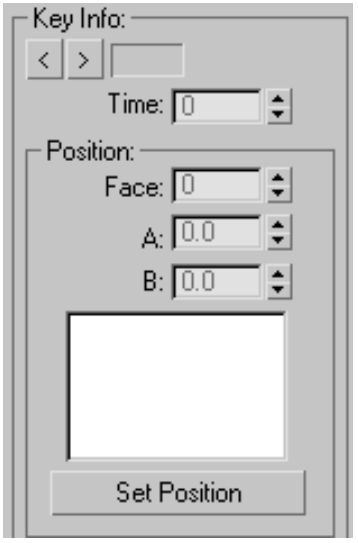
Update group



Update—Updates the display.

Manual Update—Enables Update.

Key Info group



Current Key—Displays the current key number and lets you move to another key.

Time—Displays the current frame, and lets you move the current key to a different frame.

Position group

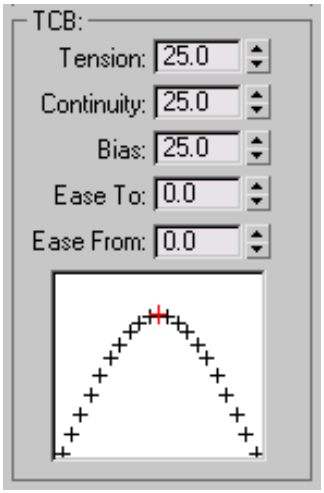
Face—Provides the index of the face to which the object is attached. Range=0 to 268435455.

A/B—Contains the barycentric coordinates defining the position of the attached object on the face. Range=-999,999 to 999,999.

(display window)—Shows the position of the source object within the attachment face. You can drag within this window to adjust the position of the object relative to the face.

Set Position—Adjusts the placement of the source object on the target object. Drag over the target object to specify a face and a position within the face. The source object moves accordingly over the target object.

TCB group



All of the items in this group are the same as in other *TCB controllers (page 2–377)*. The orientation of the source object is also interpolated and affected by these settings.

Tension—Range=0 to 50

Continuity—Range=0 to 50

Bias—Range=0 to 50

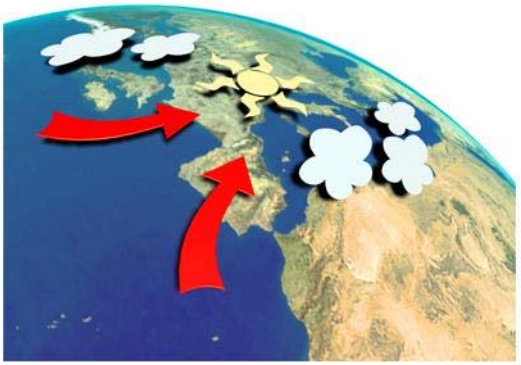
Ease To—Range=0 to 50

Ease From—Range=0 to 50

Surface Constraint

Animation menu > Constraints > Surface Constraint

Animation menu > Track View > New/Open Track View > Select a position track in Track View Hierarchy. > Track View toolbar > Assign Controller > Surface



Surface constraints position the weather symbols on the globe.

The Surface constraint positions an object along the surface of another object.

The type of object that can be used as the surface object is limited to those whose surfaces can be represented parametrically. Use the Surface constraint with the following types of objects:

- Sphere
- Cone
- Cylinder
- Torus
- Quad Patches (single quad patches)
- Loft object
- NURBS object

The surface used is a "virtual" parametric surface, and not the actual mesh surface. Objects with a low number of segments might have a mesh surface quite different than the parametric surface.

The parametric surface ignores Slice and Hemisphere options. So if the object is sliced, for

example, the controlled object will position itself as if the missing portion were still there.

Since the Surface constraint only works on parametric surfaces, if you apply a modifier that converts the object to a mesh, the constraint will no longer work. For example, you can't use it with a cylinder with a bend modifier applied.

Procedure

Example: To animate a sphere over the surface of a cylinder:

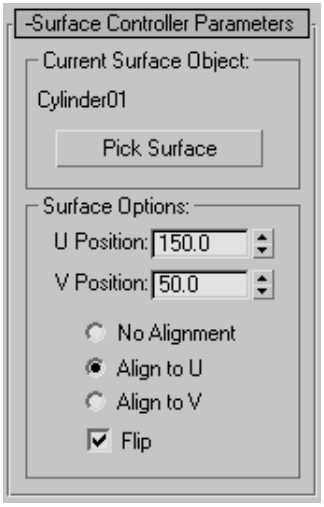
1. In the Top viewport, create a cylinder and a sphere.
2. Select the sphere, open the Motion panel, expand the Assign Controller rollout, and then expand the Transform heading in the list.
3.  In the list window, click the Position item and then click Assign Controller.
4. In the Assign Position Controller dialog, choose Surface and then click OK.

   The Surface Controller Parameters rollout replaces the Key Info rollout.
5. Click Pick Surface, and then select the cylinder.
6. Turn on Auto Key, and place the time slider at frame 0.
7. Use the V Position spinner to move the sphere to a starting position at the bottom of the cylinder.
8. Place the time slider at frame 100.
9. Use the V Position spinner to place the sphere at the top of the cylinder.
10. Set U Position to **300**.

    Turn off Auto Key and play the animation. The sphere moves over the surface of the cylinder in a helical path.

Interface

The Surface Constraint Parameters rollout is on the Motion panel.



Current Surface Object group

This group provides a method for selecting and then displaying the selected surface object.

Text—Displays the name of the selected object.

Pick Surface—Selects the object you want used as a surface.

Surface Options group

This group provides controls for adjusting the position and orientation of the object along the surface.

U Position—Adjusts the position of the controlled object along the U coordinates of the surface object.

V Position—Adjusts the position of the controlled object along the V coordinates of the surface object.

No Alignment—When on, the controlled object is not reoriented, regardless of its position on the surface object.

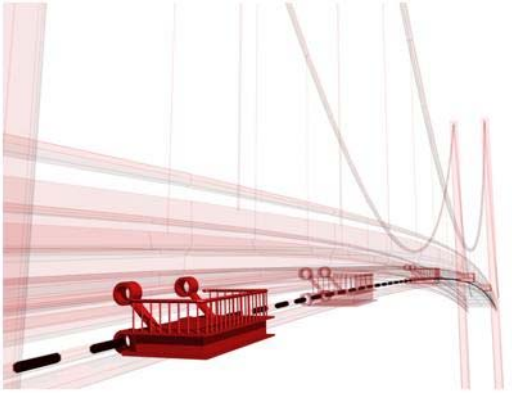
Align to U—Aligns the local Z axis of the controlled object with the surface normal of the surface object, and the X axis with the U axis of the surface object.

Align to V—Aligns the local Z axis of the controlled object with the surface normal of the surface object, and the X axis is aligned with the V axis of the surface object.

Flip—Flips the alignment of the local Z axis of the controlled object. This check box is not available if No Alignment is turned on.

Path Constraint

Animation menu > Constraints > Path Constraint



Path constraint positions the service platform along the side of the bridge.

A path constraint restricts an object's movement along a spline or at an averaged distance between multiple splines.

A *path target* can be any type of spline. The spline curve (target) defines a path of motion for the constrained object. Targets can be animated using any of the standard translation, rotation, scale tools. Setting keys on the path's sub object level, such as vertex, or segment will animate the path while affecting the constrained object.

Multiple Targets and Weighting

A constrained object can be influenced by several target objects. When using multiple targets, each target has a weight value that defines the degree by which it influences the constrained object, relative to other targets.

Using Weight is meaningful (and available) only with multiple targets. A value of 0 means the target has no influence. Any value greater than 0 causes the target to influence the constrained object relative to other targets' Weight settings. For example, a target with a Weight value of 80 will have twice the influence of a target with a Weight value of 40.

Procedures

To assign a Path constraint:

1. Create a Sphere with a radius of **10** and a Circle with a radius of **60**.
2. With the Sphere selected, choose Animation menu > Constraints > Path Constraint.

   You are now in select target mode.
3. In the viewport, select the Circle shape.

To access the Path constraint's parameters through the Motion panel:

1. Select the Sphere.
2. Open the Motion panel.
3. In the PRS Parameters rollout, click the Position button.

   The Path constraint's settings are located in the Path Parameters rollout.

To assign a Path constraint through the Motion panel:

1. Create a Sphere with a radius of **10** and a Circle with a radius of **60**.
2. Select the Sphere.
3. On the Motion panel and click Parameters.

4. Open the Assign Controller rollout and select the Position controller.
5.  Click the Assign Controller button.
6. Choose Path Constraint from the Assign Position Controller dialog.
7. On the Motion panel, click Parameters.
8. In the Path Parameters rollout, click Add Path.
9. In the viewport, select the Circle.

To edit weight values:

1. Open the Create > Shapes panel and create a Line that is about **120** units long.

   Tip: Use the diameter of the Circle to gauge the length of the Line.
2. Select the Sphere.
3. On the Motion panel, open the Path Parameters rollout.
4. Click the Add Path button and select the line.

   The Sphere is now affected equally by both paths since the path weighting defaults to 50
5. Adjust the Weight spinner or enter a numerical value for the weight value.

To animate weight values:

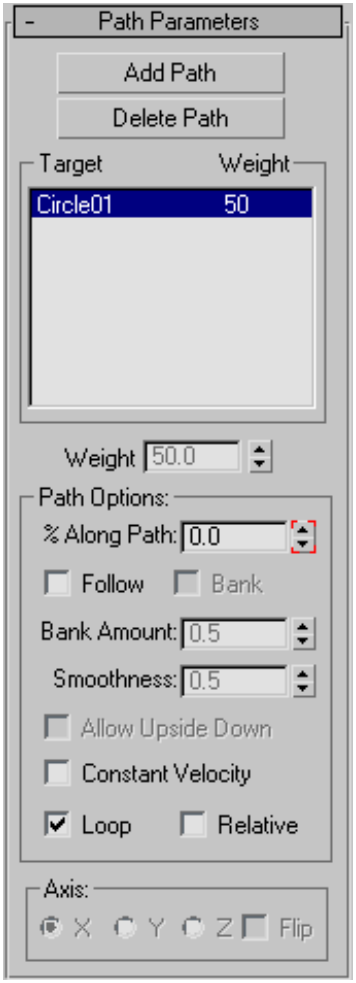
1. Select the Sphere.
2. On the Motion panel, open the Path Parameters rollout.
3. Turn on the Auto Key button.
4. Select the Line01 from the Path list.
5. Place the time slider at frame 50 and change the Weight of Line01 to **75**.
6. Move the time slider at frame 100 and change the Weight of Line01 to **10**.
7. Select the Circle01 path and change its Weighting to **25**.
8. Turn off Auto Key and play the animation.

To correct path constrained object flipping:

When an object is assigned a path constraint and the follow box is turned on, the object will rotate as it moves along the path. Sometimes the object is subject to unwanted flipping.

1. Select the object that is flipping.
2. On the Animation menu choose Constraints > Orientation constraint, then constrain the object to another object's orientation.
3. Use the control object to adjust the flipping. Animate the orientation of the control object, while watching the flipped object at the problematic frames.

Interface



Once you assign a Path constraint, you can access its properties on the Path Parameters rollout in the Motion panel. In this rollout you can add or delete targets, assign weighting, and animate each target's weight value.

Note: When you assign a Path constraint via the Animation menu, the software assigns a Position List controller to your object. In the Position List rollout list you will find Path Constraint. This is the actual path constraint controller. To view the Path Parameters rollout with the constraint settings, double-click Path Constraint in the list.

Add Path—Adds a new spline path that influences the constrained object.

Delete path—Removes a path from the target list. Once removing the path target, it will no longer influence the constrained object

Weight—Assigns and animates weight values for each target.

% Along Path—Sets the percent that the object is positioned along the path. This duplicates the Value spinner in the track Properties dialog for the Percent track in Track View. If you want to set keys to place an object at a certain percent along the path, turn on Auto Key, move to the frame where you want the key set, and adjust the % Along Path spinner to move the object.

Note: The % Along Path value is based on the parameterization of the spline path's U value. A NURBS curve might not have an evenly spaced U value, so a value of 50 % Along Path might not translate visually to 50 per cent of the NURBS curve's length.

Follow—Aligns the object to the trajectory as it follows the contour.

Bank—Allows the object to bank (roll) as it negotiates the curves of the spline.

Bank Amount—Adjusts the amount of the banking to one side or the other, depending on whether the value is positive or negative.

Smoothness—Controls how rapidly the roll angle changes as the object moves through bends in the trajectory. Smaller values will make the object more responsive to subtle changes in the curve, while larger values smooth out jerking. The default value is a good value for general damping along the curve. Values below 2 tend to make the action jerky, but values around 3 can be very useful for simulating a certain degree of realistic instability.

Allow Upside Down—Turn on to avoid the situation in which an object flips when going around a vertically oriented path.

Constant Velocity—Provides a constant velocity along the path. When off, the velocity of the object along the path varies depending on the distance between the vertices on the path.

Loop—By default, when the constrained object reaches the end of a path it can no longer move past the end point. The loop option changes this behavior so that when the constrained object reaches the end of the path it loops back to the starting point.

Relative—Turn on to maintain the original position of the constrained object. The object will follow the path with an offset distance based on its original world space position.

Axis—Defines which axis of the object is aligned to the trajectory of the path.

Flip—Turn on to flip the direction of the axis.

The following controls are located on the Hierarchy panel while the IK button is active:



Active—Activates an axis (X/Y/Z). Allows the selected object to animate along the activated path.

Limited—Limits the range of motion allowed on an active path. Use in conjunction with the From and To spinners.

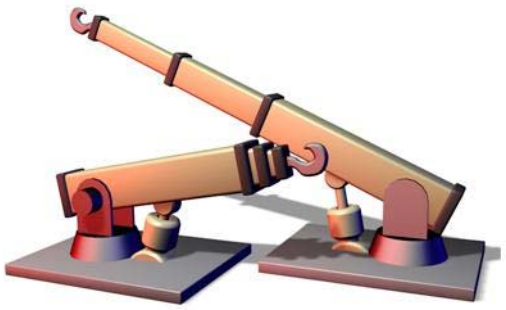
Ease—Causes a joint to resist motion as it approaches its From and To limits. Simulates an organic joint, or worn mechanical joint, moving or rotating freely in the middle of its range of motion but moving less freely at the extremes of its range.

From and To Spinners—Determine for path limits. Use in conjunction with the Limited function.

Damping—Applies resistance over a joint's motion along the path. Simulates the natural effect of joint friction or inertia.

Position Constraint

Animation menu > Constraints > Position Constraint



Position constraints align the elements of the robot assembly.

A position constraint causes an object to follow the position of an object or the weighted average position of several objects.

In order to activate, a position constraint requires an object and a target object. Once assigned the object becomes constrained to the target object's position. Animating the target's position causes the constrained object to follow.

Each target has a weight value defining its influence. A value of 0 is equal to off. Any value greater than 0 will cause the target to influence the constrained object. Weight values can be animated to create effects such as a ball being picked up from a table.

Multiple Targets and Weighting

A constrained object can be influenced by several target objects. When using multiple targets, each target has a weight value that defines the degree by which it influences the constrained object, relative to other targets.

Using Weight is meaningful (and available) only with multiple targets. A value of 0 means the target has no influence. Any value greater than 0 causes the target to influence the constrained object relative to other targets' Weight settings. For example, a target with a Weight value of 80 will have twice the influence of a target with a Weight value of 40.

For example, if a sphere is Position-constrained between two targets and each target's weight value is 100, the sphere will maintain an equal distance between both targets even when they are in motion. If one of the weight values is 0 and the other is 50, then the sphere is influenced only by the target with the higher value.

Procedures

To assign a Position constraint:

1. Select the object you want to constrain.
2. Choose Animation menu > Constraints > Position Constraint.
3. Select the target object.

To access the Position constraint's parameters through the Motion panel:

1. Select the Position-constrained object.
2. On the Motion panel, on the Position list, double-click Position Constraint.

   The Position constraint parameters are located under the Position Constraint rollout.

To edit weight values:

1. Select the constrained object.
2. On the Motion panel, on the Position list, double-click Position Constraint.

   The Position constraint parameters are located under the Position Constraint rollout.
3. Click a target from the list.
4. Adjust the Weight spinner or enter a numerical value for the weight value.

To animate weight values:

1. Select the constrained object.
2. On the Motion panel, on the Position list, double-click Position Constraint.

   The Position constraint parameters are located under the Position Constraint rollout.
3. Click a target from the list.
4. Turn on the Auto Key button.
5. Adjust the Weight spinner or enter a numerical value for the weight value.

Example: To assign a Position constraint with two targets and editing weights:

1. In the Top viewport, create a sphere, a box, and a cylinder so that the box is between the sphere and the cylinder.
2. Click to select the box, assign a Position constraint, and select the sphere as the target.
3. Click to select the box, assign a Position constraint, and select the cylinder as the target.

   The box is now position-constrained between the two targets.
4. In the Top viewport, move the sphere around.

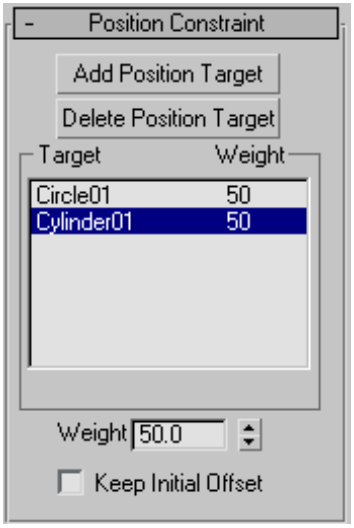
   As the sphere moves, the box maintains an equal distance between the sphere and the cylinder. This is because the weight values for both targets are equal. By default the values are 1.00. If the sphere had a higher weight value than the cylinder, the sphere would influence the box more than the cylinder.

5. To edit the weight values, select the box.
6. Open the Motion panel and view the Position Constraint rollout.
7. Click the Cylinder's name in list of targets.
8. Using the Weight spinner, change the value from 50 to **20**.

   As the value decreases, the box moves closer to the sphere.
9. In the Top viewport. select the cylinder and move it around.
10. In the Top viewport, select the sphere and move it around.

    The sphere has more influence over the box's movement than the cylinder.

Interface



Once you assign a Position constraint, you can access its properties on the Position Constraint rollout in the Motion panel. In this rollout you can add or delete targets, assign weighting, and animate each target's weight value.

Note: When you assign a Position constraint via the Animation menu, the software assigns a Position List controller to your object. In the Position List rollout list you will find Position Constraint. This is the actual Position Constraint controller. To view the Position Constraint rollout, double-click Position Constraint in the list.

Add position target—Adds new target objects that influence the position constrained object.

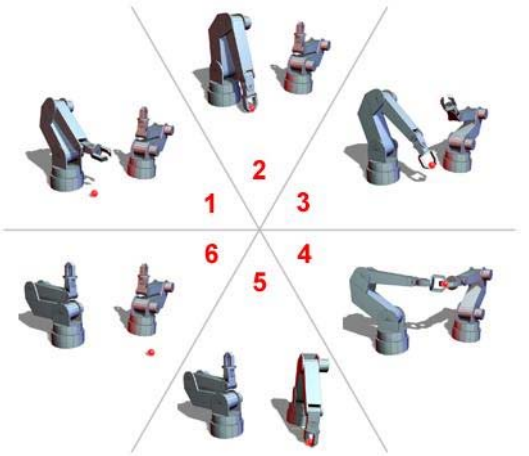
Delete position target—Removes targets. Once a target is removed, it will no longer influence the constrained object.

Weight—Assigns and animates weight values for each target.

Keep Initial Offset—Use Keep Initial Offset to preserve the original distance between the constrained object and the target object. This prevents the constrained object from snapping to the target object's pivot. The default is Off.

Link Constraint

Animation menu > Constraints > Link Constraint



A link constraint enables the robot arms to pass a ball.

A Link constraint is used to animate an object linking from one target object to another.

The Link constraint causes an object to inherit the position, rotation, and scale of its target object.

An example of using a link constraint is to pass a ball from one hand to another. Assume that at frame 0 the ball is in the right hand. The hands are animated to meet at frame 50, where the ball is passed to the left hand, and then spread apart until frame 100.

Procedures

Example: To assign a Link constraint:

1. In the Top viewport, create a sphere and a box.
2. Select the sphere.
3. Open the Motion panel.
4. Expand > Assign Controller rollout.
5. Select the Transform: Position/Rotation/Scale controller.
6.  Click Assign Controller.
7. Choose Link Constraint.

   It is recommended that you add a Link To World before linking to any target objects. This will allow you to animate the object on its own, before the first link to a target object takes place.

   You can also assign Link Constraints through the Animation > Constraints submenu.

To access the Link constraint's parameters through the Motion panel:

1. Select the Link-constrained object.
2. Choose Motion panel > Link Params rollout.

Example: To assign the Link constraint and animate links:

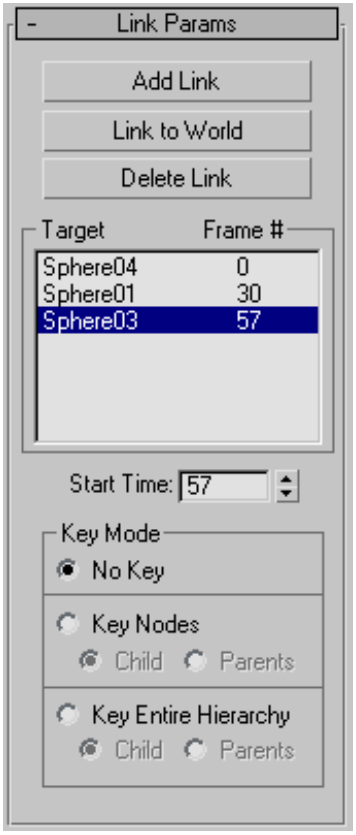
1. Move the time slider to frame 0.
2. In the Top viewport, create a sphere, a cylinder and a box.
3. Select the sphere.
4. Open the Motion panel.
5. Expand the Assign Controller rollout.
6. Select the Transform: Position/Rotation/Scale controller.
7. Click Assign Controller.
8. Choose Link Constraint.
9. In the Link Params rollout, click Link To World.
10. Move the time slider to frame 1.
11. Click Add Link and select the cylinder.

    It now becomes a target. The Link Constraint relationship is now active between the sphere and the cylinder.
12. Move the time slider to frame 50.
13. Turn on the Auto Key button and move the cylinder a good distance away from where it is.
14. Play back the animation.

    The sphere follows the cylinder as it moves. The sphere is link-constrained.
15. Select the sphere.
16. Move the time slider to frame 25.
17. On the Motion panel > Link Param rollout, click Add link and select the box.

    You have now added another target that becomes active at frame 25.
18. Turn off the Auto Key button and play back the animation. The sphere is linked to the cylinder from frame 0 to 24. It follows the cylinder until frame 25 where it links to the box.

Interface



Once you assign a Link constraint, you can access its properties on the Link Params rollout in the Motion panel. In this rollout you can add and delete targets and animate the time at which each target becomes the active parent of the constrained object.

Add Link—Adds a new link target.

Link to World—Links the object to the world (the scene as a whole). We recommend this be the first target in the list. This prevents the object from reverting to its independent creation or animation transforms if other targets are deleted from the list.

Delete Link—Removes a link target. Once a link target is removed, it will no longer influence the constrained object.

Start Time—The start time spinner is used to assign or edit the frame value of a target. Select one of the target object names in the list window and see the frame where the object becomes a parent. You can adjust the value to change when the link transfer takes place.

Key Mode group

Note: The options Key Nodes and Key Entire Hierarchy have no effect unless the object you are constraining is already part of a hierarchy. If you add objects to the hierarchy after you apply the Link constraint, you have to reapply the Link constraint using the key options you desire.

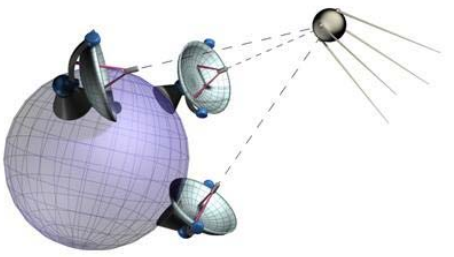
No Key—When used, no keyframes are written to the constrained object or its targets. The link control happens without inserting any keys. This is similar to how the Link Controller worked in 3ds Max 3.1 and earlier versions.

Key Nodes—When used, keyframes are written to the specified option. There are two options: **Child** and **Parents**. Child sets a keyframe only on the constrained object. Parents sets keyframes for the constrained object and all of its targets.

Key Entire Hierarchy—Sets keyframes up the hierarchy for the specified option. There are two options: **Child** and **Parents**. Child sets a keyframe only on the constrained object and its parents. Parents sets keyframes for the constrained object, its targets, and their upper hierarchy.

LookAt Constraint

Animation menu > Constraints > LookAt Constraint



LookAt constraints enable the antenna dishes to track the satellite.

The LookAt constraint controls an object's orientation so that it's always looking at another object. It locks an object's rotation so that one of its axes points toward the target object. The LookAt axis points toward the target, while the Upnode axis defines which axis points upward. If the two coincide, a flipping behavior may result. This is similar to pointing a target camera straight up.

An example of a LookAt constraint's use would be to constrain the eyeballs of a character to a point helper. The eyes will then always be aimed at the point helper. Animate the point helper, and the eyes follow. Even if you rotate the character's head, the eyes maintain their lock on the point helper.

Multiple Targets and Weighting

A constrained object can be influenced by several target objects. When using multiple targets, each target has a weight value that defines the degree by which it influences the constrained object, relative to other targets.

Using Weight is meaningful (and available) only with multiple targets. A value of 0 means the target has no influence. Any value greater than 0 causes the target to influence the constrained object relative to other targets' Weight settings. For example, a target with a Weight value of 80 will have twice the influence of a target with a Weight value of 40.

Procedures

To assign a LookAt constraint:

1. Select the object you want to constrain. This is the object that will be always looking at its target.
2. Choose Animation menu > Constraints > LookAt Constraint.
3. Select the target object.

To access the LookAt constraint's parameters through the Motion panel:

1. Select the LookAt Constrained object.
2. On the Motion panel > Rotation list, double-click LookAt Constraint.

   The LookAt constraint parameters are located under the LookAt Constraint rollout.

To edit weight values:

1. Select the constrained object.
2. On the Motion panel > Rotation list, double-click LookAt Constraint.

   The LookAt constraint parameters are located under the LookAt Constraint rollout.
3. Click a target from the list.
4. Use the Weight spinner or enter a numerical value to adjust the weight value.

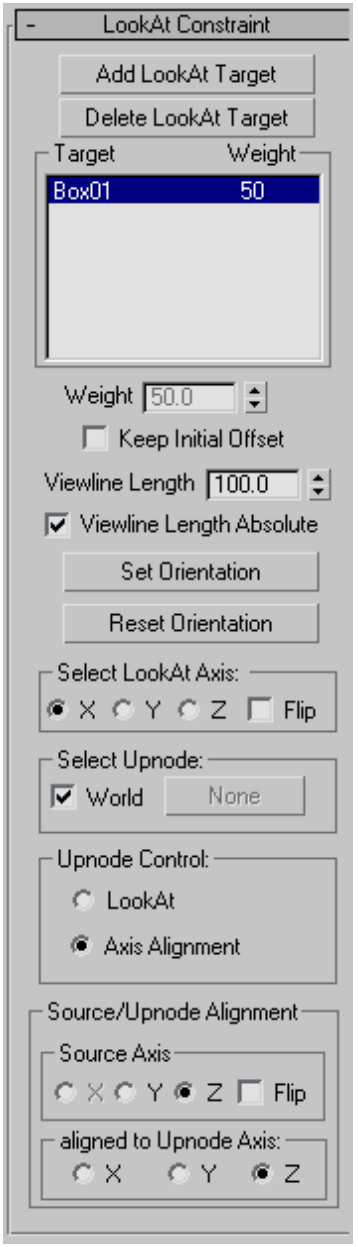
To animate weight values:

1. Select the constrained object.
2. On the Motion panel > Rotation list, double-click LookAt Constraint.

   The LookAt constraint parameters are located under the LookAt Constraint rollout.

3. Click a target from the list.
4. Turn on the Auto Key button.
5. Use the Weight spinner or enter a numerical value to adjust the weight value.

Interface



Once you assign a LookAt constraint, you can access its properties on the LookAt Constraint rollout on the Motion panel. On this rollout, you can add or delete targets, assign weighting, assign

and animate target weight values, and adjust other related parameters.

Note: When you assign a LookAt constraint via the Animation menu, 3ds Max assigns a Rotation List controller to your object. In the list on the Rotation List rollout, you will find LookAt Constraint, which is the constraint you assigned. To view the LookAt Constraint rollout, double-click the LookAt Constraint entry in the list.

Add LookAt Target—Use to add new targets that influence the constrained object.

Delete LookAt Target—Use to remove target objects that influence the constrained object.

Weight—Use to assign and animate weight values for each target. Available only when multiple targets are used.

Keep Initial Offset—Maintains the constrained object's original orientation as an offset to its constrained orientation.

Viewline Length—Defines the length of the main viewline drawn from the pivot of the constrained object to the pivot of its target (or the average, in case of multiple targets). A negative value draws the line from the constrained object in the opposite direction of the target or targets.

With a single target, the length of the viewline is determined by the distance between the constrained object and the target, as well as the Viewline Length setting. However, if Viewline Length Absolute is on, the distance between the two has no effect on the length.

The color of the viewline is defined by the Target Line element in the Gizmos category of the *Colors panel (page 3–799)* in the Customize > *Customize User Interface dialog (page 3–792)*.

Note: When multiple targets are assigned, additional viewlines drawn from the constrained object to each target object inherit the color of the respective targets. If Viewline Length Absolute is on, the length of each target-specific line is determined by its target's Weight setting and the Viewline Length value. If Viewline Length Absolute is off, the length of each line is determined by the distance between the constrained object and the respective target, as well as the Viewline Length value. An additional (main) viewline, whose length and color are determined as specified above, indicates the actual, calculated orientation.

Viewline Length Absolute—When on, the software uses only the Viewline Length setting for the length of the main viewline; the distance between the constrained object and the target(s) has no effect.

Set Orientation—Lets you define the offset orientation of the constrained object manually. When on, you can use the Rotation tool to set the constrained object's orientation. This orientation is then maintained as the constrained object looks at its target.

Reset Orientation—Sets the orientation of the constrained object back to the default. This is useful if you want to reset the constrained object's orientation after having set the orientation manually.

Select LookAt Axis—Use to define the axis that looks at the target. The X,Y,Z check boxes reflect the constrained object's local coordinate system. The Flip check box reverses the directions of the local axes.

Select Upnode—The default Upnode is the World. Turn off World to manually select an object that defines the Upnode plane. This plane is drawn from the constrained object to the Upnode object. If the LookAt Axis and the Upnode axis coincide, the constrained object will flip. Animating the position of the upnode object will move the upnode plane.

Upnode Control group:

Lets you quickly flip between LookAt Upnode Control and Axis Alignment.

LookAt—When selected the Upnode matches the LookAt target.

Axis Alignment—When this is selected the Upnode Aligns to the object axis. Choose which axis (X, Y or Z) in the Source Upnode Alignment group directly below Upnode Control.

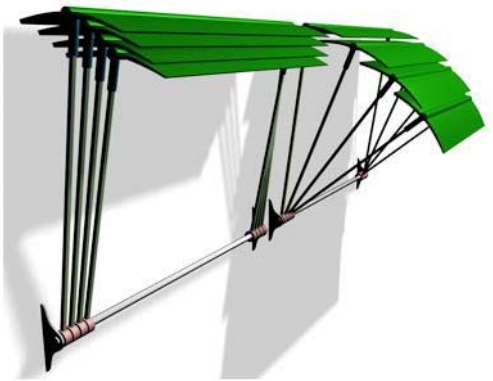
Source/Upnode Alignment group

Source Axis—Chooses the constrained object's axis that is to be aligned to the Upnode Axis. The Source Axis reflects the constrained object's Local Axis. The Source Axis and LookAt Axis work together therefore the Axis used to define the LookAt Axis will be unavailable.

Aligned to Upnode Axis—Chooses the Upnode axis that the selected Source Axis aligns to. Note that the selected Source axis may or may not be able to completely align to the Upnode Axis.

Orientation Constraint

Animation menu > Constraints > Orientation Constraint



Orientation constraints align the awning vanes to the supporting rod.

An Orientation constraint causes an object's orientation to follow the orientation of an object or averaged orientation of several objects.

An Orientation Constrained object can be any rotatable object. When constrained it will inherit its rotation from a target object. Once constrained you can not rotate the object manually. You may move or scale the object as long as its not constrained in a manner that effects the object's position or scale controller.

The target object can be any type of object. The rotation of a target object drives the constrained object. Targets can be animated using any of the standard translation, rotation, and scale tools.

Multiple Targets and Weighting

A constrained object can be influenced by several target objects. When using multiple targets, each target has a weight value that defines the degree by which it influences the constrained object, relative to other targets.

Using Weight is meaningful (and available) only with multiple targets. A value of 0 means the target has no influence. Any value greater than 0 causes the target to influence the constrained object relative to other targets' Weight settings. For example, a target with a Weight value of 80 will have twice the influence of a target with a Weight value of 40.

Procedures

To assign an Orientation constraint:

1. Select the object you want to constrain.
2. Choose Animation menu > Constraints > Orientation Constraint.
3. Select target object.

To access the Orientation constraint's parameters through the Motion panel:

1. Select the Orientation-constrained object.

2. On the Motion panel > Rotation list, double-click Orientation Controller.

   The Orientation constraint parameters are located on the Orientation Constraint rollout.

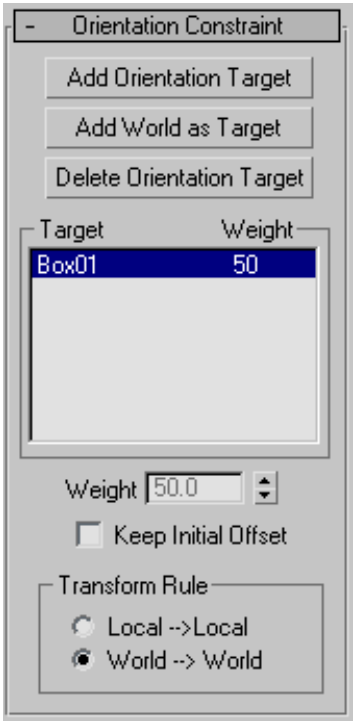
To edit weight values:

1. Select the constrained object.
2. Open the Motion panel > Orientation Constraint rollout.
3. Click a target from the list.
4. Use the Weight spinner or enter a numerical value to set the weight value.

To animate weight values:

1. Select the constrained object.
2. Open the Motion panel > Orientation Constraint rollout, and choose a target from the list.
3. Turn on the Auto Key button.
4. Use the Weight spinner or enter a numerical value to set the weight value.

Interface



Once you assign an Orientation constraint, you can access its properties on the Position Constraint rollout in the Motion panel. In this rollout you can add or delete targets, assign weighting, assign and animate target weight values, and adjust other, related parameters.

Note: When you assign an Orientation constraint via the Animation menu, the software assigns a Rotation List controller to your object. In the Rotation List rollout list you will find Orientation Constraint, which is the constraint you assigned. To view the Orientation Constraint rollout, double-click Orientation Constraint entry in the list.

Add Orientation Target—Adds new target objects that influence the constrained object.

Add World as Target—Aligns the constrained object to the world axis. You can weight the amount

of influence that the world target has on the constrained object as you would any other target object.

Delete Orientation Target—Remove targets. Once removing the target, it will no longer influence the constrained object.

Weight—Assigns and animates weight values for each target.

Keep Initial Offset—Preserves the original orientation of the constrained object. When you turn off Keep Initial Offset, the object adjusts itself to match the orientation of its target or targets. Default=off.

Transform Rule—When an orientation constraint is applied to an object that is part of a hierarchy, this determines whether the local node transform or the parent transform will be used for the orientation constraint.

Local —>—Local—When selected, the local node transform is used for the orientation constraint

World —>—World—When selected, the parent or world transform will be applied, rather than the local node transform.

Wire Parameters

Select an object. > Animation menu > Wire Parameters > Wire Parameters

Select an object. > Right-click quad menu > Transform (lower-right) quadrant > Wire Parameters

Wire Parameters lets you link parameters from one object to another in the viewport, so that adjusting one parameter changes the other automatically. This enables you to set up one- and two-way connections between specified object parameters, or to control one or more objects with a dummy object containing the desired parameters. By wiring parameters, you can set up custom constraints directly without having to go to Track View and assign controllers.

Parameter wiring is accessible from the Animation menu and the *quad menus (page 3–694)*. The Wire Parameters command is available only if a single node is selected. Choosing the Wire Parameters command displays a hierarchical pop-up menu with levels and items corresponding to the animatable tracks that would be visible for that object in Track View.

When you make a parameter selection, a dashed line, similar to that displayed in the *Select and Link mode (page 2–422)*, is drawn from the selected object to the mouse cursor. The cursor changes from an arrow to a cross whenever it is over a valid destination object. You can right-click at any time to cancel the parameter wiring.

While the dashed line is displayed, you can click a destination node or on empty space in the viewport to customize the relationship between two wired parameters. If you click empty space, the *Parameter Wiring dialog (page 2–412)* opens showing the first parameter in the left Tree View and the full scene in the right tree. Clicking a destination node (which can be the same as the originating node) displays a hierarchical pop-up menu for the destination object so you can choose the destination parameter. After you choose a second parameter, the Parameter Wiring dialog opens.

You can wire one-way and two-way connections between parameters. For one-way wires, one parameter is effectively slaved to the other and its value changes as the controlling parameter changes, according to a user-defined transfer expression. The controlling parameter can be animated and adjusted as needed using all the animation tools in 3ds Max. This includes making it a controlled parameter in another

parameter-wiring setup, so that you could potentially set up a chain of controlled parameters.

For two-way wiring, the software assigns Wire controllers of the appropriate kinds to each of the parameters and they are cross-linked so that changing either parameter cause linked changes in the other.

Note: You should establish all object hierarchies before wiring parameters in a scene. If you change the hierarchy of an object that has a wired parameter, it will take on new parameters, which could introduce undesired results in your wired parameters.

Tip: Use *Schematic View (page 3–640)* to see all the wired parameter relationships in a scene. You can also use Schematic view to wire parameters.

Procedure

To attach wire parameters:

1. Select the first object you want to use.
2. Choose Animation menu > Wire Parameters > Wire Parameters, or right-click the object and choose Wire Parameters.

   A pop-up menu displays the parameters you can link.

   Note: Wire Parameters works only with parameters that can be animated. The menu options displayed depend upon the selected object and are the same options you would see for the selected object in Track View.
3. Choose the parameter you want to link to another object from the displayed menu.

   A dashed line appears.
4. Move the cursor to the second object you want to use. When the cursor changes from an arrow to a selection cross, you are over a valid destination object. Click to select the destination object.

   A pop-up menu displays the parameters to which you can link.
5. Choose the parameter you want to link from the displayed menu.

   The *Parameter Wiring dialog (page 2–412)* opens, with the two parameters you selected highlighted in yellow.
6. If you like, change the parameter highlighting in the dialog, and then click the direction buttons between the two panels to determine the control direction: one way either way, or both directions (changes to either object's selected parameters affect the other object).
7. Click Connect to complete the wiring.

Parameter Wiring Dialog

Animation menu > Wire Parameters > Choose parameters to wire.
Select an object. > Right-click quad menu > Transform (lower-right) quadrant > Wire Parameters > Choose parameters to wire.
Select an object. > Modify panel > Right-click text box portion of animatable parameter spinner. > Choose Show In Parameter Wire dialog.
Select an object with a wired parameter. > Modify panel > Right-click text box portion of *two-way-wired* parameter > Edit Wire

The Parameter Wiring dialog allows you to define the relationships for *Wire Parameters (page 2–411)*. In this dialog, you can create new one and two-way control relationships between object parameters, edit existing relationships, and create or edit expressions which define the parameter relationships.

Only parameters that can be animated can be wired. Sub-objects, such as vertices, must be animated before they can be wired.

Parameter wires can be used to establish connections from *Manipulators (page 2–27)*

and *Custom Attributes (page 1–129)* to objects, materials, and modifiers.

Note: Parameter Wiring is a *modeless dialog (page 3–973)*, and you can have several dialogs open at the same time.

Procedures

Example: To use an expression with wire parameters:

1. Create a box and a sphere in your scene.
2. Right-click the box, and from the quad menu choose Transform quadrant > Wire Parameters.
3. From the pop-up menu, choose Object (Box) > Length.

   A rubber-band dashed line now connects the box and the mouse cursor.
4. Click the sphere and choose Transform > Position > Z Position from the pop-up menu.

   The position of the sphere and the length of the box are now wired.
5. In the Parameter Wiring dialog, create a one-way wire by clicking the arrow pointing toward the box's tree view.
6. Type **abs(** before Z_position, followed by a closing parenthesis, **)**, in the expression text box for Box01, so that the entire expressions is *abs(Z_Position)*.
7. Click Connect.
8. Move the sphere upward along the Z axis in the viewport.

   Note that as the sphere moves in the viewport, the length of the box changes. Using abs() ensures that the box never has a negative length.

To control several slave parameters from one master parameter:

1. Open the Parameter Wiring dialog.
2. Choose the master parameter in the tree view.
3. Choose the first slave parameter in the other tree view.
4. Click the arrow pointing toward the slave parameter.
5. Enter the desired relationship expression in the expression text box.
6. Click Connect.
7. Keeping the master parameter selected, choose another slave parameter from the tree view.
8. Click the arrow pointing toward the slave parameter.
9. Enter the desired relationship expression in the expression text box.
10. Click Connect.

    You can repeat this cycle to set up as many slaves as you want. The result is a "fan" of parameter wires, so that you can control all of the slave parameters at once as you change the master parameter.

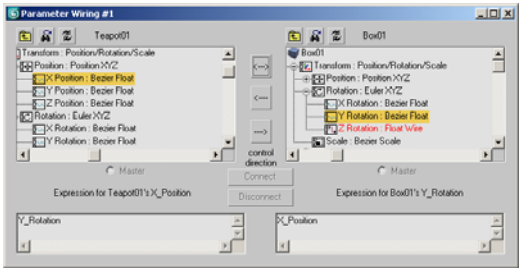
To create a chain of wired parameters:

1. Open the Parameter Wiring dialog.
2. Choose the first controlling parameter in one tree view.
3. Choose the first slave parameter in the other tree view.
4. Click the arrow pointing toward the slave parameter.
5. Enter the desired relationship expression in the expression text box.
6. Click Connect.
7. Keeping the original slave parameter selected, choose the next slave parameter from the tree view.
8. Click the arrow pointing toward the new slave parameter.
9. Enter the desired relationship expression in the expression text box.

10. Click Connect.

You can repeat this cycle, continuing to alternate slaves to masters as many times as you like. The result is a "daisy chain" of parameter wires, so that as you modify the original master parameter, each slave parameter maintains a constant relationship with the next in a line of parameters.

Interface



Tree Views

The dialog presents two tree views that display the animatable parameters of all of the visible objects in the scene. The names of the currently selected objects appear at the top. The tree views display the objects in the scene and allow you to select and wire all of the animatable parameters of the objects in the scene and the scene itself.

The tree views are color-coded to show existing wiring. A parameter with a wire controller assigned to it (either as the member of a two-way pair or as the slave in a one-way wire) displays in red. When a wire-controlled parameter is selected, all of the parameters wired directly to it are displayed in green text in the other Tree View. In either case, if the parameter is inside a track that is not expanded, the enclosing track will display in red or green, so you can expand the tracks to find the wires.

Note: If you open the Parameter Wiring dialog by wiring two parameters, the tree views initially display only the selected objects, highlighting the parameters that you chose from the pop-up menus during the wiring procedure. If you want to expand the trees, to include all other animatable parameters in the scene, click Show All Tracks.

The buttons above the tree views are:

 **Show All Tracks—**Brings you to the top of the scene object list.

 **Find Next Parameter—**Finds the next wired parameter in the scene, expanding the tree if necessary to show and select it. You can browse through all existing wires by clicking this button repeatedly.

Note: If you click a wired parameter, its connections are shown in green on the other tree view. While the wired parameter is highlighted, the Show Next Wire button on the other tree cycles through the green connections only, so you can quickly find the selected wire's connections.

 **Refresh Tree View Content to Selected Node—**Shows only the node or nodes that are selected in the scene in the tree view. The nodes are shown as currently exanded.

Control Direction—The Parameter Wiring dialog provides three direction-control buttons between the tree views. You can choose only one of these at a time. These buttons determine the direction of control, either one-way or two-way:

-  **Two-way connection [two-headed arrow]—**Click this to link both parameters to each other, so that changing either parameter affects the other.
-  **One-way connection: right parameter controls left parameter [left arrow]—**The right parameter controls the left parameter.

- 

One-way connection: left parameter controls right parameter [right arrow]—The left parameter controls the right parameter.

Note: You can change the direction of a wire at any time by clicking the desired direction button and then clicking Connect or Update.



Connect/Update—The Connect/Update button changes its label depending on whether clicking it would add a new wire to previously unwired parameters (Connect) or change an existing wire's expressions or direction (Update). Changes to existing wires are only applied when you click Update.



Disconnect—The Disconnect button is enabled when you've selected parameters with an existing wire between them. It will remove the wire controller(s) and replace them with the Master parameter's animation track (if two-way) or with default controllers (if one-way).

Master parameters

The Parameter Wire system provides a way for the pair of wired parameters to be animated as a single system. It does this by setting up a subcontroller on one of the parameters; any animation on this controller drives the Wire Controller pair. The animation subcontroller is always assigned to the master parameter of the wired pair. By default, this is the parameter that is clicked first in the wiring interaction. The master parameter can also be designated with the Master radio buttons beneath the tree views.

The animation subcontroller appears as a nested track inside the master parameter's track in Track View and its values directly drive and match the master parameter values. If the subcontroller is keyframable (which is the default when a two-way wiring is first established) it can be keyframed by adjustments of *either* of the wired parameters. This means that if you keyframe the wired pair by adjusting the non-master parameter, the values keyframed into the animation subcontroller are derived from the master parameter transfer expression. Since the wired pair can be animated through either parameter, the choice of master parameter is essentially just a convention.

Transfer Expression boxes

Underneath the parameter trees are the transfer expression text boxes. These expressions determine how changes to each parameter affect the other and are usually inverses of one another.

For more information on expression syntax, see the topic "Script Controllers" in the *MAXScript Reference*.

When the parameters are first wired, the default expressions are simple 1-to-1 links between the parameters. These expressions can be edited into any valid script fragment that will yield a result of the correct type for its parameter. For example, if you link a parameter such as height (which contains a float value) to a parameter such as position (which contains a point3 value), your expressions must include conversions that produce the same output value type.

For one-way wires, the Expression box for the controlling parameter is unavailable, since there is no Wire controller assigned to it.

For two-way wiring, both transfer expression text boxes will be enabled. It is possible for the user to supply transfer expressions for the two parameters that are not inverses of one another, but this is of course discouraged, since the relationship of the parameters will be different depending on which one is changed first.

Hierarchies and Kinematics

When animating characters, mechanical assemblies, or complex motion, you can simplify the process by linking objects together to form a *hierarchy (page 3–951)* or chain. In a linked chain, the animation of one member can affect some or all of the others, making it possible to animate a number of objects or bones at once.

The term *kinematics (page 3–960)* describes the movement or animation of the chain. There are two types of kinematics:

- With *forward kinematics (page 3–944)* (FK) , you transform a parent object to move its descendants (its children, their children, etc.).
- With *inverse kinematics (page 3–958)* (IK), you transform a child object to move its ancestors (its parent and so on up the chain). You can also use IK to make an object "stick" to the ground or another surface, while allowing the chain to rotate off the pivot of that object.

Forward kinematics is the most straightforward method for animating hierarchies. Inverse kinematics requires more setup than forward kinematics, but is more intuitive for complex tasks such as character animation .

See also

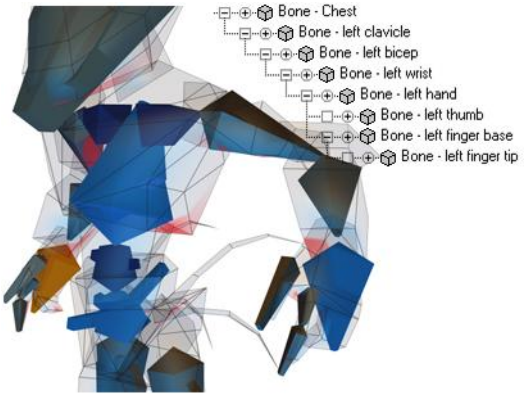
Hierarchies (page 2–416)

Animating with Forward Kinematics (page 2–426)

Inverse Kinematics (IK) (page 2–435)

Hierarchy Panel Commands (page 2–487)

Hierarchies



Bipedal figure with part of its hierarchy diagrammed in the background

One of the most useful tools in producing computer animation is the ability to link objects together to form a *chain*. By linking one object to another, you create a parent-child relationship. Transforms applied to the parent are also transmitted to child objects. A chain is also referred to as a *hierarchy*.

You can find the commands to build and manipulate hierarchies in the following places in the interface:

- The *Bones system (page 1–404)* in the Create panel > Systems category lets you create a hierarchy of bones. You can also create bones by choosing *Bone Tools (page 1–411)* from the *Animation menu (page 3–681)*. You can turn any hierarchy of objects into bones by selecting the hierarchy and turning on Bone Tools dialog > Object Properties rollout > Bone On.
- The *Hierarchy panel (page 3–773)* contains commands to control how links behave.
- The *Motion panel (page 2–463)* contains commands to control how links behave when using an History Dependent (HD) Solver

Common Uses for Hierarchies

- Link a large collection of objects to a single parent so they can be easily animated and transformed by moving, rotating, or scaling the parent.
- Link the target of a camera or light to another object so it tracks the object through the scene.
- Link objects to dummy objects to create complex motions by combining multiple simple motions.
- Link objects to simulate jointed structures to animate characters or mechanical assemblies.

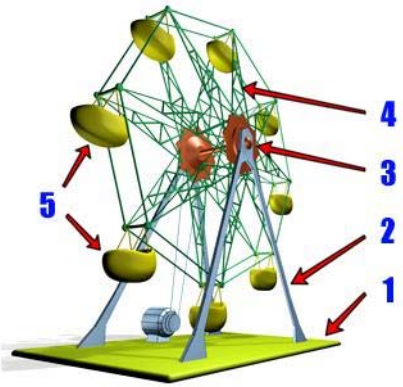
Parts of a Hierarchy

The relationship between objects linked together in a hierarchy is analogous to a family tree.

Parent—Object that controls one or more children. A parent object is often controlled by another superior parent object. In the following figure, objects 1 and 2 are parent objects.

Child—Object controlled by its parent. A child object can also be a parent to other children. In the following figure, objects 2 and 3 (the support and hub) are children of object 1. Objects 5 (the seats) are children of object 4, the Ferris wheel.

Ancestors—Parent and all of the parent's parents of a child object. In the following figure, objects 1 and 2 are ancestors of object 3.



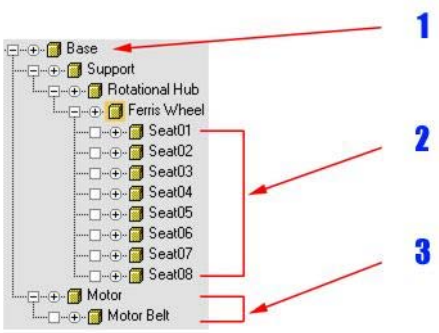
The seats of the Ferris wheel are children of the wheel, which is in turn a child of the base and support objects, as shown in the following hierarchy.

Descendants—Children and all of the children's children of a parent object. In the figures, all the objects are descendants of object 1.

Hierarchy—Collection of all parents and children linked together in a single structure.

Root—Single parent object that is superior to all other objects in the hierarchy. All other objects are descendants of the root object. In the figures, Object 1 is the root.

Subtree—All the descendants of a selected parent. In the figure below, the Rotational Hub, Ferris Wheel, and Seats represent the subtree under the Support object.



1. Root

2. Leaves

3. Subtree

Example of a hierarchical structure

Branch—Path through the hierarchy from a parent to a single descendant. In the figure above, the Support, Rotational Hub, and Ferris Wheel objects comprise a branch from the root to the leaf objects (the seats).

Leaf—Child object that has no children. The lowest object in a branch. In the figure above, the Seat objects are leaf objects.

Link—Connection between a parent and its child. A link transmits position, rotation, and scale information from parent to child.

Pivot—Defines the local center and coordinate system for each object. You can think of a link as the connection between the pivot of a child object and the pivot of its parent.

Linking Strategy

Before you begin linking any but the simplest hierarchy you should take a few minutes to plan your linking strategy. Your choices for the root of the hierarchy and how the branches grow out to the leaf objects will have important effects on the usability of your model.

The strategy behind linking objects into a hierarchy can be reduced to two main principles:

- The hierarchy follows a logical progression from parent to child.
- Parent objects move less than their descendants.

Within these two principles you have almost unlimited flexibility as to how you link your objects. If you think about how you intend to use the hierarchy, and link it with that use in mind, you will rarely have a problem.

Progression from Parent to Child

Progression from parent to child means the links do not erratically jump from object to object. If two objects touch each other they should probably be linked as parent and child. There is nothing to prevent you from linking a body in the order of: Thigh->Foot->Shin->Waist. You would probably regret such a linking strategy later. The effort to figure out how to transform objects linked in such a strange way would be quite difficult. A more logical progression would be Foot->Shin->Thigh->Waist.

Using Multiple Hierarchies

Rather than build a single bone chain from a hip to a toe, you can make one chain from the hip to the ankle, and then a second independent chain from the heel to the toe. You would then link the chains together to form a complete leg assembly.

Because they are linked together, the leg and foot chains could be considered one chain. However, the way you animate them treats each chain separately, allowing fine control over the parts.

With this type of arrangement on leg and foot chains, the foot could be made to stay on the ground while the leg bends. It also allows for independent control of the foot's rotation, pivoting on the heel or toe, which would then cause the knee to bend.

Parents Move Less Than Descendants

Because of the way transforms are inherited from parent to child, small adjustments to a parent object might require you to adjust all of its descendants. The typical approach to linking is to choose as your root object the object that moves the least. Objects close to the root should move very little, and leaf objects should move the most.

This is especially true when you are linking jointed structures like robots or machinery, or intend to use the hierarchy with *inverse kinematics (page 2–435)*.

An exception to this rule occurs when you are using the root object as a handle. All of the descendants of the root are just along for the ride. Consider a tray full of objects traveling on a conveyor belt. All the objects should be children of the tray even though the tray moves much more than any of the other objects.

Choosing the Root of a Hierarchy

You can find the best candidate for the root of your hierarchy by asking the following question:

If I move this object, should all of the other objects in the hierarchy move with it?

- If the answer is almost always, then you are looking at a likely candidate for the root object. Examples of this type of object are a torso, a lamp base, and a tree trunk.
- If the answer is not often, then you are probably looking at a child object. Examples of this type of object are hands, lamp shades, and tree leaves. If you move a character's hand, for example, its torso should not move.

Once you have a few candidates for the root object, you can examine them in greater detail. Use these criteria to determine a good root object for your hierarchy:

- Moving the root object usually has a great effect on all other objects in the hierarchy.
- Conversely, the root object is mostly unaffected by movement of other objects in the hierarchy.
- The root object is rarely animated, and is moved or rotated primarily to place the hierarchy at the correct place in the scene.
- The root object is at or near the hierarchy's virtual center of mass.

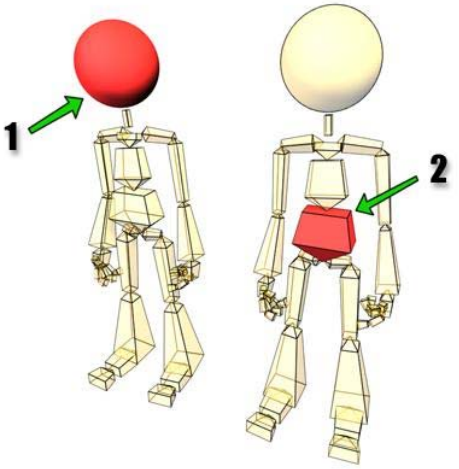
The object that best satisfies these criteria is your root object. You then create your hierarchy with all of the other objects as descendants of that root object.

Linking Objects for Inverse Kinematics

Inverse kinematics (IK) uses the child object as the driving force for the animation. IK is less forgiving and is highly dependent on the linking strategy for performing calculations.

You need to consider two additional principles when linking hierarchies for use with inverse kinematics:

- Links and pivot placement simulate real-world joint locations.
- Choose an object near the structure's center of mass, or center of gravity, as the root of the hierarchy. The center of mass in the real world is the point on an object about which reactions to external forces are applied.



1 and 2 each represent the root of the characters.

Both structures are suitable for forward kinematics.

The structure on the right is best for most inverse kinematics.

The figure above shows two approaches to linking a skeletal structure. Either structure is suitable for working with forward kinematics. The structure on the right, however, is a better choice for working with inverse kinematics.

- The root object is located near the body's center of mass.
- The link order more closely simulates the connections of a real body.

The structure on the left has the arms and torso linked to the neck. The structure on the right links the arms and neck to the torso, a more realistic approach.

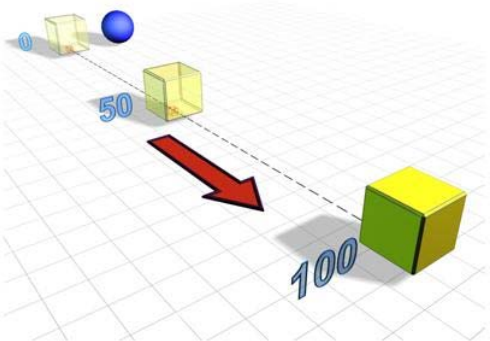
Warning: Be sure to avoid using non-uniform scaling on objects in a hierarchy that will be animated using IK. You will see stretching and skewing if you do. Instead do all non-uniform scaling at the sub-object level, to avoid this problem. Use Reset XForm if you have objects that exhibit this behavior.

Linking Objects After Animation

When you link an object to another, the link relationship between the child and its parent is determined by the position, rotation and scale of the parent and child objects when the link is made.

Imagine linking a stationary sphere to an animated box.

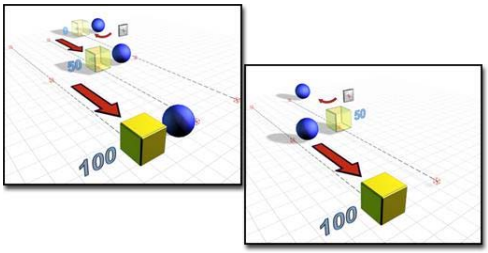
- At frame 0 the box is beside the sphere.
- At frame 50 the box is 20 units away.



Original animation, with ball unlinked and stationary while the box moves.

Linking the sphere to the box causes the sphere to move with the box. The distance between the sphere and the box depends on the frame when the link is made. Linking the sphere on different frames has the following effects:

- Link on frame 0, and the sphere stays next to the box as it moves.
- Link on frame 50, and the sphere stays a distance 20 units away from the box as it moves.



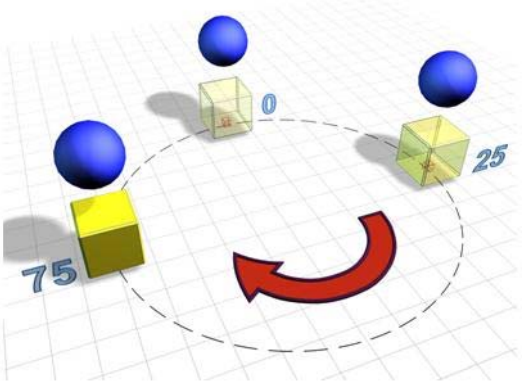
Left: Ball linked at frame 0 follows the box on the side.

Right: Ball linked at frame 50 follows the box 20 units away.

Unlinking Objects After Animation

When you unlink a child, its frame 0 transforms are taken from the transforms of its parent at the frame when the link is removed.

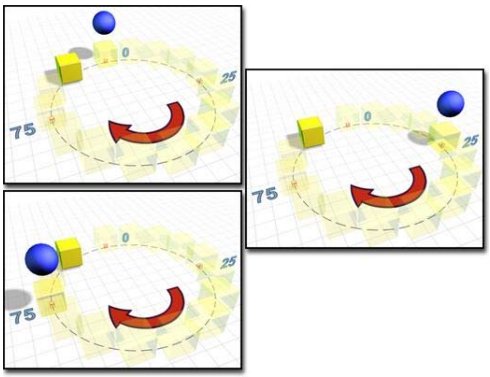
Imagine a sphere linked to a box moving around the face of a clock. The box starts at 12 o'clock and travels all the way around the face over 100 frames. The figure shows a box moving in a circle with a sphere linked above it.



Original animation, with ball linked to follow the animated box.

If you unlink the sphere, it stops following the box. The position of the sphere depends on its position, rotation or scale at the frame on which the link is removed. Unlinking the sphere on different frames has the following effects:

- Unlink on frame 0, and the sphere stays at 12 o'clock.
- Unlink on frame 25, and the sphere stops at 3 o'clock.
- Unlink on frame 75, and the sphere stops at 9 o'clock.



Clockwise from top, position of the sphere unlinked at frame 0, 25 and 75, respectively.

Linking and Unlinking Objects


Use *Select and Link (page 2–422)* and *Unlink Selection (page 2–422)* on the toolbar to make and remove links between objects.

Linking Objects


The general process of creating links is to build the hierarchy from child to parent. You click Select And Link on the toolbar, select one or more objects as children, and then drag the link cursor from the selection to a single parent object. The selected objects become children of the parent object.

Once objects are linked, any transformations applied to the parent are also applied to its children. For example, if you scale the parent to

150%, the size of its children and the distance between the children and the parent are also scaled by 150%.

Unlinking Objects


Click Unlink Selection to remove the link from selected objects to their parents. Any children of the selected object are unaffected.

You can quickly unlink an entire hierarchy by double-clicking the root object to select the object and all of its children. Then click Unlink Selection.

Linking Animated Objects

You should establish links before you begin animating objects. The linkage of objects with Select and Link cannot be animated; the link remains in force throughout the entire animation.

If you want your objects to be linked during one part of the animation but not another, you can a *Link constraint (page 2–403)* to change the linkage at specific frames.

Displaying Links

A complex mesh hierarchy can be displayed with the links visible, or even with the links replacing the mesh objects. To display links, first select the linked objects. On the Display panel > *Link Display rollout (page 1–58)*, turn on Display Links to see the links. You can also turn on Link Replaces Object to see only the links and not the objects.

Select and Link 

Main toolbar > Select and Link

Use the Select and Link button to define the *hierarchical relationship (page 3–951)* between two objects by linking them as child and parent.

You link from the currently selected object (child) to any other object (parent).

You can link an object to a closed group. When you do, the object becomes a child of the group parent rather than any member of the group. The entire group flashes to show that you've linked to the group.

A child inherits the transformations (move, rotate, scale) applied to the parent, but the child's transformations have no effect on the parent. If you want the child not to inherit the transforms, use the *Link Inheritance (Selected) Utility (page 2–435)* or use the controls found in *Link Info (page 2–499)* in the Hierarchy panel.

You can also create hierarchical linkages using *Schematic View (page 3–638)*. Use the Connect button on the Schematic View toolbar to create hierarchical linkages between nodes.

Procedure

To link two objects:

1. 
Click Select and Link.
2. Drag a line from an object (the child) to any other object (the parent).

   Note: You do not need to select the child object first.

Unlink Selection 

Main toolbar > Unlink Selection

Use the Unlink Selection button to remove the *hierarchical relationship (page 3–951)* between two objects.

Unlink Selection detaches a child object from its parent object.

Tip: You can also link and unlink hierarchies in *Schematic View (page 3–640)*.

Procedure

To unlink a child object from a parent object:

1. Select the child object you want to unlink.
2.  Click Unlink Selection.

Adjusting Pivots

You can think of an object's pivot point as representing its local center and local coordinate system. The pivot point of an object is used for a number of purposes:

- As the center for rotation and scaling when the Pivot Point transform center is selected.
- As the default location of a modifier center.
- As the transform offset for linked children.
- As the joint location for IK.

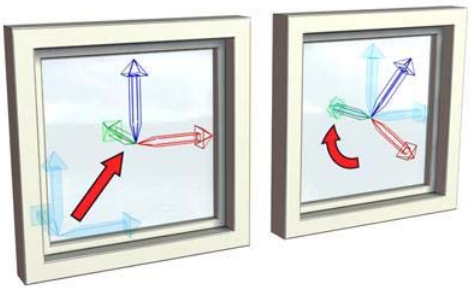
You can adjust pivot points by clicking Pivot on the Hierarchy panel, and then using features on the *Adjust Pivot rollout (page 2–488)*.

The functions under the Adjust Pivot rollout cannot be animated. Adjusting an object's pivot on any frame changes it for the entire animation.

Affecting Pivot Only

When Affect Pivot Only is on, move and rotate transforms are only applied to the pivot of a selected object.

- Moving or rotating the pivot does not affect the object or its children.
- Scaling the pivot scales the object from the pivot center, but its children are unaffected.

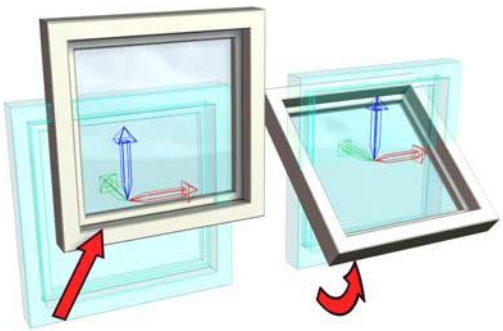


Affect Pivot Only transforms the pivot without moving the object.

Affecting Object Only

When Affect Object Only is on, transforms are applied only to selected objects. Pivots are not affected.

Moving, rotating, or scaling the object does not affect the pivot or its children.



Affect Object Only transforms the object without moving the pivot.

Affecting Hierarchy Only

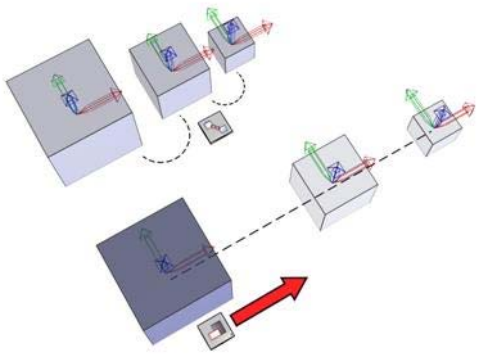
When Affect Hierarchy Only is on, Rotate and Scale transforms are applied only to the links between objects and their children.

- Scaling or rotating an object affects the link offsets of all its descendents without affecting

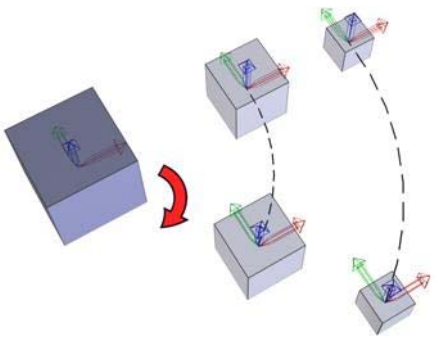
the geometry of the object or its descendents. The descendents shift position because of the scaled or rotated links.

Use this technique to adjust the offset relationship between linked objects.

Use this technique to adjust the offset relationship between linked objects and for adjusting bones to match geometry.



After a hierarchy is created, you can scale the position of the children without changing the individual objects' dimensions.



Rotating the hierarchy does not affect the individual objects' orientation.

Aligning Pivots

Buttons on the Alignment group box of the Adjust Pivot rollout change names based on the state of Affect Object Only and Affect Pivot Only.

Alignment is disabled when Affect Hierarchy Only is active.

Center to Object/Pivot—Moves the object, or pivot, so the pivot is at the center of the object.

Align to Object/Pivot—Rotates the object, or pivot, to align the pivot with the object's original local coordinate system.

Align to World—Rotates the object, or pivot, to align with the world coordinate system.

Resetting the Pivot

Click Reset Pivot to return the pivot point of a selected object to the position and orientation it held when the object was first created.

Reset Pivot has no effect on the object or its children. The state of the Affect Pivot Only and Affect Object Only is ignored.

Viewing and Selecting Hierarchies

There are a number of ways to view a hierarchy structure and select objects in it.

Viewing a Hierarchy

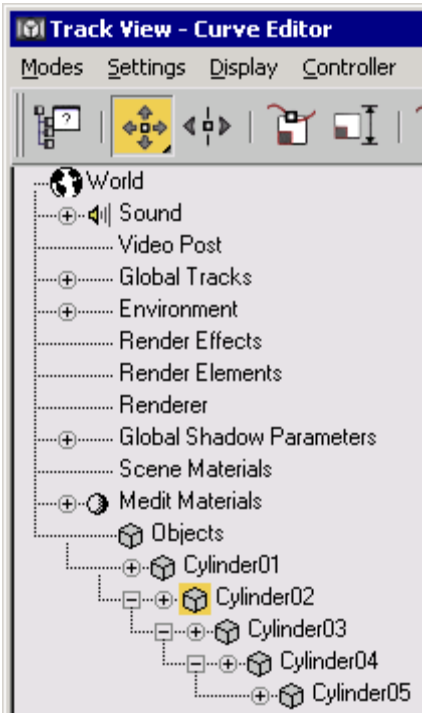
You can use these methods to view the relationships between parents and children in a linked hierarchy.

- The *Select Objects dialog (page 1–78)* appears whenever you use a by-name selection method, such as choosing Edit menu > Select By > Name, by clicking Select By Name on the main toolbar, or by pressing the H key.

  To list objects hierarchically, turn on Display Subtree on the dialog. This indents children below their parent.

- The *Hierarchy list (page 2–512)* at the left side of the *Track View window (page 2–501)* displays all objects using indentation to express hierarchy. Child objects are displayed indented and below their parent. An added advantage of Track View

is that you can control the view by collapsing and expanding branches of the hierarchy.



Displaying the hierarchy in the Controller list

A square icon with a plus indicates a collapsed branch under that object, while a minus indicates an expanded branch. Click a plus icon to expand a branch, or a minus icon to collapse it.

Tip: In complex scenes, use Curve Editor to navigate quickly through the Track View. Simply select the object in the viewport, then right-click and choose Curve Editor. The Track View — Curve Editor will appear with the selected object at the top of the window.

- You can also use *Schematic View (page 3–640)* to view hierarchies. In addition to showing you the structure, Schematic View contains tools for manipulating hierarchies.

Selecting Hierarchy Members: Ancestors and Descendants

Once you have selected one or more objects in a hierarchy, you can select its direct ancestor or descendant with the Page Up and Page Down keys.

- Page Up deselects the object and selects the object's parent.
- Page Down deselects the object and selects all its immediate children, but not all descendants down the chain.

Tip: These navigation commands are particularly useful when setting joint parameters for inverse kinematics.

To select an object and all of its descendants, you can:

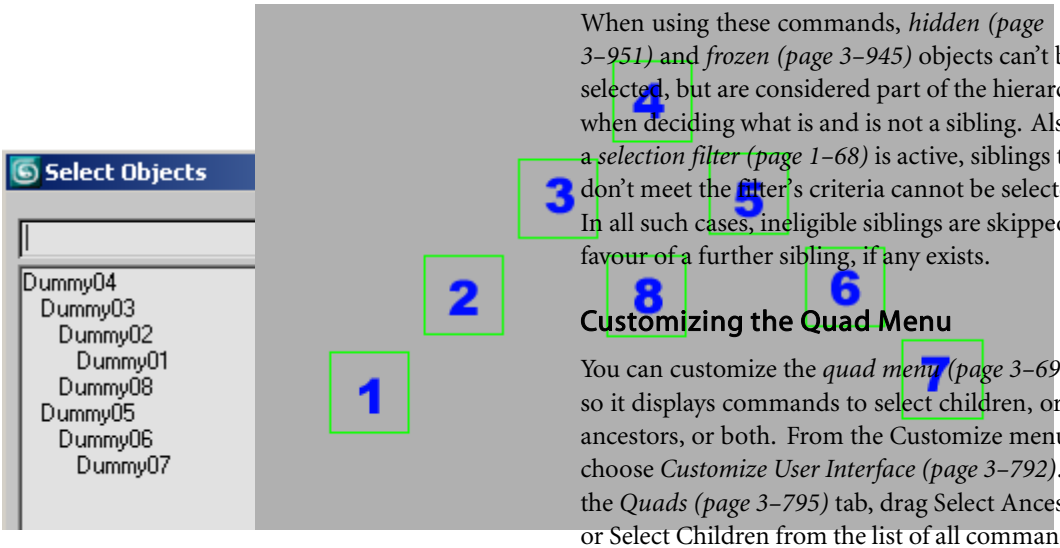
- Double-click the object in a viewport.
- Double-click the object icon in the Track View *Hierarchy list (page 2–512)*.

Selecting Hierarchy Members: Siblings

Available in *Customize User Interface (page 3–792)* are the actions Select Sibling - Next and Select Sibling - Previous, which you can assign as hotkeys, toolbar buttons, etc. We recommend assigning them to the cursor keys Right Arrow and Left Arrow, respectively; by default, those keys are not assigned keyboard shortcuts.

Using one of these commands replaces the current selection with one object only at the same hierarchical level. More precisely, a sibling is defined in this context as an object of equal generational distance from the selected object's *nearest* parent. All objects that fit this definition are siblings, so that in asymmetrical hierarchies, object A can be a sibling of object B, but the reverse is not necessarily the case.

Take, for example, the asymmetrical hierarchy of dummy helper objects shown below:



At first glance it might appear that Dummies 06, 08 and 02 are siblings of each other, but in practice it is more complicated:

- Selecting Dummy06 and activating Select Sibling - Next results in Dummy08 being selected.
- Selecting Dummy06 and activating Select Sibling - Previous results in Dummy02 being selected.
- With Dummy02 selected, activating either Select Sibling command results in Dummy08 being selected, and the selection continues alternating for each subsequent invocation of either command. This is because 3ds Max looks only as far as the next branch above the selected object, resulting in Dummies 02 and 08 being siblings of Dummy06, but not vice-versa.

Defining a sibling in this way has the practical advantage of letting you, for example, cycle through the selection of all finger links on one side of a character without the selection jumping to the other hand. But cycling through siblings with an arm object selected usually results in selecting the opposite arm object.

When using these commands, *hidden (page 3–951)* and *frozen (page 3–945)* objects can't be selected, but are considered part of the hierarchy when deciding what is and is not a sibling. Also, if a *selection filter (page 1–68)* is active, siblings that don't meet the filter's criteria cannot be selected. In all such cases, ineligible siblings are skipped in favour of a further sibling, if any exists.

Customizing the Quad Menu

You can customize the *quad menu (page 3–694)* so it displays commands to select children, or ancestors, or both. From the Customize menu, choose *Customize User Interface (page 3–792)*. On the *Quads (page 3–795)* tab, drag Select Ancestor or Select Children from the list of all commands to the quad menu. Then you can easily select children or parents with a right-click and then a click.

Animating with Forward Kinematics

The default method of manipulating a hierarchy uses a technique called *forward kinematics (page 3–944)*. The basic principles employed by this technique are:

- Hierarchical linking from parent to child
- Placement of pivot points to define the connecting joint between linked objects
- Inheritance of position, rotation, and scale transforms from parent to child

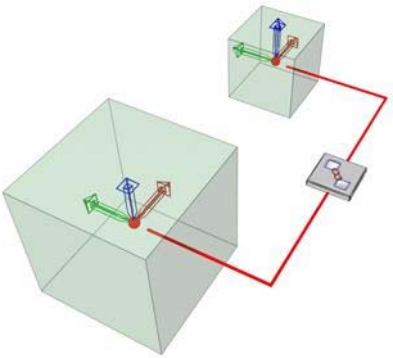
You animate the objects of a hierarchy in much the same way you animate anything else. Turn on the Auto Key button and transform members of the hierarchy at different frames. However, you need

to be aware of a few special issues for animating hierarchies.

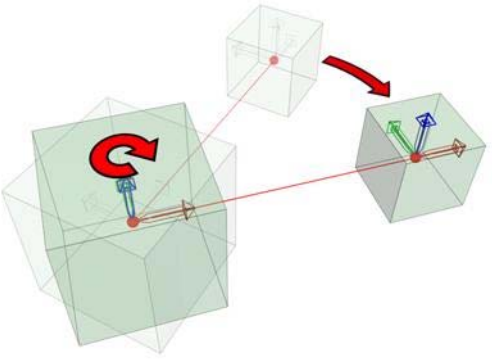
How Links and Pivots Work

Once two objects are linked together, the child object maintains its position, rotation, and scale transforms relative to its parent object. These transforms are measured from the pivot of the parent to the pivot of the child.

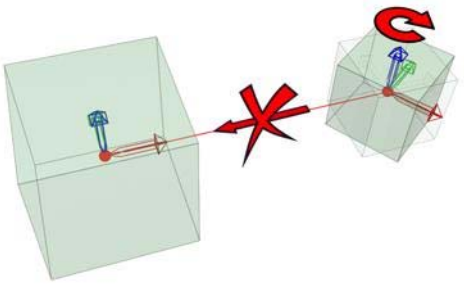
For example, consider the two boxes in the following figure. The larger box is the parent of the smaller. The pivots and link between the boxes are indicated to show how the link works. The link extends from the pivot of the parent and connects to the pivot of the child. You can think of the child's pivot as being the joint between the parent and child.



Parent and child objects linked by their pivot points.



Rotating the parent affects the position and orientation of the child object.



Rotating the child does not affect the parent.

Links act as a one-way conduit to transmit the transforms of a parent object to its child object. If you move, rotate, or scale the parent, the child is moved, rotated, or scaled by the same amount. Because hierarchical links are one-way, moving, rotating, or scaling the child has no effect on its parent.

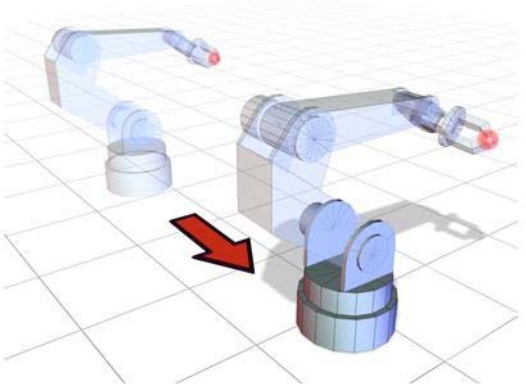
The end result is that transforms applied to a child object are applied in addition to any transforms inherited from the child's parent.

Animating a Parent Object

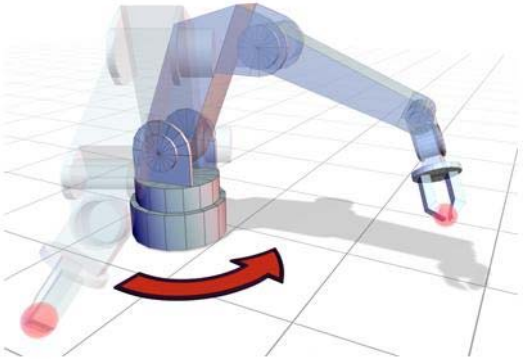
Only transforms are passed from parent to child. Animating a parent object using move, rotate, or

scale animates the parent and the subtree attached to the parent.

Animating a parent's modifiers or creation parameters has no effect on its descendants.



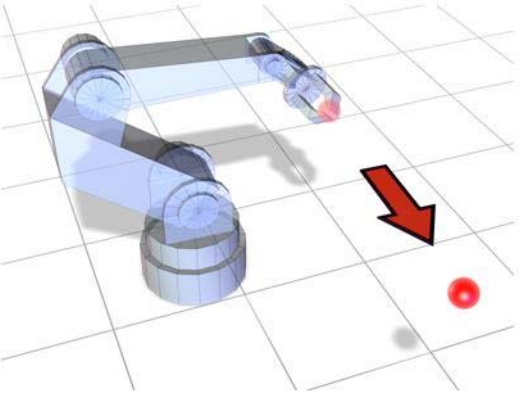
Moving the root parent moves the whole hierarchy.



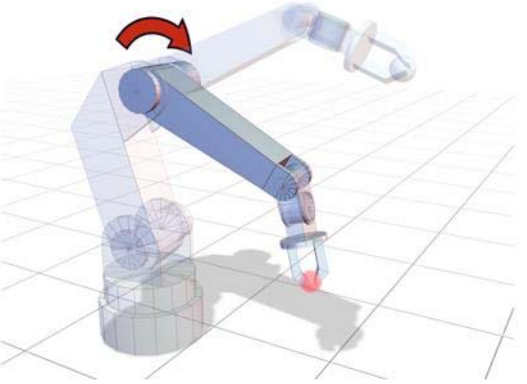
Rotation of a parent object is passed to all the child objects.

Animating a Child Object

With forward kinematics, a child is not constrained by its link to a parent. You can move, rotate, and scale children independent of their parents.



Moving the last child object does not affect any of the previous objects in the hierarchy.

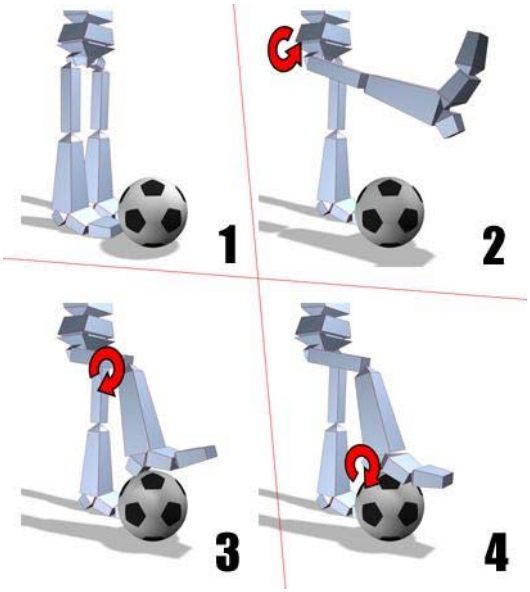


Rotating a child object in the middle of the hierarchy affects all the descendants but none of the parents.

If you want to manipulate parent objects by moving the last child in the chain, use *inverse kinematics (page 2–435)*.

Manipulating the Hierarchy

A child object inherits the transforms of its parent, and the parent inherits the transforms of its ancestors all the way up the hierarchy to the root object. Because forward kinematics employs this method of inheritance, you must position and animate your hierarchies using a top-down method.



Manipulating the hierarchy of a leg.

Consider the linked mannequin in the figure. If you want to position the mannequin's right foot to rest on top of the soccer ball beside it, you perform the following steps:

1. Rotate the right thigh so the entire leg is above the soccer ball.
2. Rotate the right shin so the foot is near the top of the soccer ball.
3. Rotate the right foot so it is parallel with the top.
4. Repeat steps 1 through 3 until the foot is properly placed.

You always start transforming objects at the highest-level parent affected by the motion and work your way down the hierarchy to the last child.

You have considerable control over the exact placement of every object in the hierarchy using forward kinematics. However, the process can become tedious with large and complex hierarchies. In such situations, you might want to use *inverse kinematics (page 2–435)*.

Using Dummy Objects

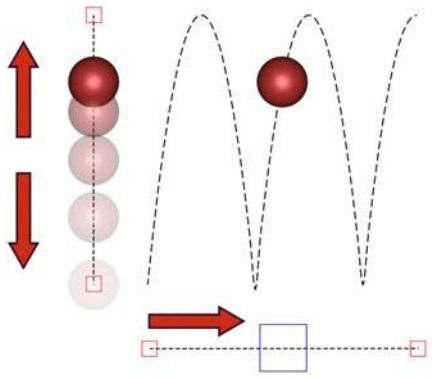
The primary use of *dummy helper objects (page 2–16)* is to assist in creating complex motions and building complex hierarchies. Because dummies are invisible when rendered, they are an excellent choice for offset joints, connectors between objects, and handles for complex hierarchies. Dummies and *Points (page 2–23)* can act as null objects that function as controls for transforming parts of an IK chain.

Using a Dummy to Control Motion

Breaking complex motions into simple components often makes it easier to go back and edit your animations.

Consider a bouncing ball that moves along a path. You could animate the ball by positioning it on many frames, but it would be very difficult for you to go back and adjust the height of the bounce or the path of the ball. You have to edit the motion of the ball on many frames to make even a simple change.

Using a dummy object solves this problem by breaking the motion into simple components. One component is the up and down bounce of the ball. The other is movement on the path.



Combining the bouncing motion of a ball with the forward motion of a dummy results in a moving bouncing ball.

Using a Dummy as a Handle

You might want to move and animate a selection of objects individually but also have the ability to transform them as a single object.

A good example of this is a camera on a tripod. You want to adjust both the camera and its target individually but also want to move them together as a single unit.

Procedures

To create a complex bounce motion using a dummy object:

1. Start with a sphere, then create a dummy object below the sphere, and link the sphere as a child of the dummy.
2. Animate the sphere bouncing up and down above the dummy.
3. Animate the dummy moving.

   The sphere bounces on top of the dummy object as the dummy moves around the scene. You can easily change the height and speed of the bounce by changing the sphere animation. You can change the path through the scene by changing the dummy animation.

To create a camera tripod:

- Create a dummy object below a target camera and link the camera and target as children of the dummy object.

  The camera and the target follow the dummy object. You can quickly position the camera by placing the dummy object and compose your view by adjusting the camera and its target.

Animating Links

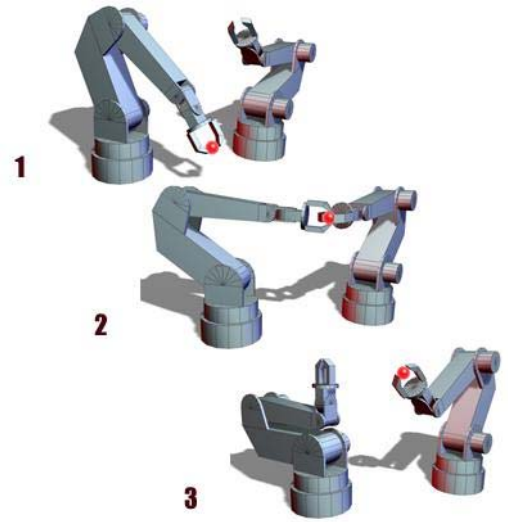
You assign a *Link constraint (page 2–403)* to an object to animate links from one parent to another. You use a link constraint instead of using the regular Select and Link and Unlink Selection buttons on the toolbar. (See *Animation Constraints (page 2–392)*.)

An example of using a link constraint is to pass a ball from one hand to another. Assume that at frame 0 the ball is in the first hand. The hands are animated to meet at frame 50 and then spread apart until frame 100.

To animate the links for the ball:

1. On the Motion panel, assign a Link constraint as the ball's Transform controller. You can also put a Link constraint on the ball from the Animation menu by choosing Constraints > Link Constraint.
2. Move the time slider to frame 0, then on the Motion panel click Add Link, and click the hand holding the ball. The ball will now move along with the hand, as if it were linked to it.
3. Move the time slider to frame 50, where you want the second hand to pick up the ball, click Add Link, then click the second hand. From this frame on, it's as if the ball were linked to the second hand.

When you play back the animation, the ball travels with the first hand until frame 50, where you added the second link, then the ball is passed to the second hand for the rest of the animation.



Robot arms pass a ball from one hand to the other.

Adding and Deleting Links

You add and delete links on the Motion panel. Expand the Link Parameters rollout and click Add Link or Delete Link.

- Click Add Link then click the object that you want to link to as a parent. The frame at which you add the link is the frame at which control is passed. You can change the link frame with the Start Time parameter.
- Click the name of a parent object in the list, and then click Delete Link to remove the link.

Properties of the Link constraint include:

- The Link constraint respects the link inheritance settings applied to the child object.
- The object using a Link constraint is not a true child object. It does not appear in the subtree of any linked parent objects.
- Objects with Link constraint do not participate in IK solutions.

Link to World

You can also link an object to the world using the Link to World button. This will keep the object stationary without the use of a dummy object. Just click Link to World and the world is automatically entered as a Target.

Key Modes

You can choose between three different *key modes*, which determine how keyframes are written on the linked objects as part of the link constraint. These options provide the following:

No Key Mode—No keys are created any of the objects involved. No keys will be visible in the track bar.

Key Nodes—Sets keys for some of the objects. **Child** applies keys to the child object only. **Parents** applies keys to both parents and the child object.

Key Entire Hierarchy—This applies keyframes to the chosen nodes and their entire hierarchies. **Child** keys the chosen object and the nodes in its hierarchy up to the world. **Parents** keys both parents and the child and all three hierarchies up to the world.

Side Effects of the Link Constraint

The Link constraint works to keep a child object from jumping position at the time when the link changes from one parent to another parent.

Considering the previous example, the following should hold true:

- During frames 0 to 50 the ball remains constant relative to the first hand.
- During frames 50 to 100 the ball remains constant relative to the second hand.
- At frame 50, the time when link control changes, the ball does not jump.

If you change the animation of the second hand at frame 75, it affects the position of the hand relative to the ball at the time of the link (frame 50). This change in relative position affects the ball over all frames where it is linked to the second hand. Therefore, as you change the position of the hand at frame 75, the child's position will also change, possibly in a counter-intuitive way. However, when playing back the animation the above three rules will hold true.

Adjusting Object Transforms

You use the features on the *Adjust Transform (page 2–489)* rollout to transform objects after they have been linked without transforming descendents, and to reset an object's transform.

Transforming Parent Objects

You might discover, after linking a number of objects, that you need to move, rotate, or scale a parent object but you do not want to affect the object's descendents. You can transform a parent object without affecting its descendents by clicking Don't Affect Children on the Adjust Transform rollout of the Hierarchy panel.

Resetting an Object's Orientation and Scale

Click the Transform button in the Reset group to rotate an object's pivot to match its parent's local coordinate system. Descendents of the object are not affected.

Resetting an Object's Scale Only

Click the Scale button in the Reset group to set the current scale value as the selected object's base scale value. All following scale transforms are then applied using the base scale value as an absolute local scale of 100%.

Consider a sphere with a radius of 20 units and a linked child object:

1. Use Uniform Scale to scale the sphere to 200%. The sphere and its child become twice as big. Scale Transform Type-In reports an Absolute Local Scale of 200% and the object's creation parameters report a radius of 20 units. The true radius of the sphere is 200% of 20 units, or 40 units.
2. Select the sphere and click Reset Scale. The sphere and its child remain the same size. Here's what has happened:

   The 200% scale has been absorbed by the sphere as its original state. The sphere has a true radius of 40 units, Creation Parameters report a radius of 20 units, and absolute local scale is 100%.

   The sphere's child object accepts a local scale of 200% so it does not change in size.

Resetting the scale of an object can lead to confusion because the object's true size, absolute local scale, and creation parameters no longer match up.

Using the Reset Transform Utility

You can also reset the orientation and scale of an object by clicking *Reset XForm (page 1–438)* on the *Utilities panel (page 3–778)*. Reset XForm takes the rotation and scale transforms of an object and places them in an XForm modifier on the modifier stack.

Consider the same sphere as before with a radius of 20 units and a linked child object:

1. Use Uniform Scale to scale the sphere to 200%. The sphere and its child become twice as big. The Scale Transform Type-In reports an Absolute Local scale of 200% and Creation Parameters report a radius of 20 units. The true radius of the sphere is 200% of 20 units, or 40 units.

2. Select the sphere and click Reset XForm. The sphere remains the same size but its child reverts to its original size and position. Here's what has happened:

   The 200% scale has been placed in an XForm modifier on the sphere's modifier stack. The sphere has a true radius of 40 units, Creation Parameters report a radius of 20 units, and absolute local scale is 100%.

   The sphere's child object now sees only the 100% local scale so it reverts to its original size and position.

Locking Object Transforms

You can lock an object's ability to move, rotate, or scale about any of its local axes by selecting objects and then setting options on the *Locks (page 2–500)* rollout of the Hierarchy panel.

Enabling and disabling Local transform axes is also referred to as setting degrees of freedom (DoF) for an object. If an axis is enabled, an object is free to transform about that Local axis.

The Locks rollout contains three groups: one each for Move, Rotate, and Scale. Each group contains three options, one each for the X, Y, and Z Local axes of the selected objects.

- When turned on, the objects cannot be transformed about the selected Local axes when you directly use one of the transform tools.

  Objects can still be transformed by other means such as being a child of a transformed parent object or being part of an inverse kinematics chain.

- When turned off, objects can be freely transformed about the unlocked Local axes.

Animating Attachment

You assign an *Attachment constraint (page 2–393)* to cause an object to hold a position on the surface of another object. The Attachment constraint is not a hierarchical link, but it has the effect of "linking" an object to the surface of another object as follows:

- Attach a "source" object to the face of a "target" object so that the source object acts as if it's glued to the target object, no matter how the surface of the target object is deformed.
- Animate the Attachment parameters so that the source object moves over the surface of the target object.

Unlike hierarchical linking, which considers only object transforms, an object using an Attachment constraint follows the deformations of another object based on that object's modifiers and space warp bindings.

See *Animation Constraints. (page 2–392)*

Setting Attachment Parameters

You use features on the Attachment Parameters rollout on the Motion panel, to pick a target object and position the source object.

- Click Pick Object, then click the target object to perform the attachment.
- Click Set Position and click or drag on the surface of the target object to place the source object onto the surface.

  If you want to move the source object along the normal of the face (move it below or above the face) click Affect Object Only on the Hierarchy panel and move the object, using Local transform coordinates.

- Turn on Align to Surface if you want to align the world Z axis of the source object with the surface normal of the target object. Even with

Align to Surface selected, you can always rotate the source object to orient it the way you want in relation to the target object.

If you were to align trees on an uneven terrain, you would turn off Align to Surface so that all of the trees grew upright, regardless of the angle of the terrain surface.

Animating Attachment Position

You can move to any frame and click Set Position to animate the source object moving across the surface of the target object. It is not necessary to turn on the Auto Key button, because you are working with an animation constraint.

When you set positions for the source target on multiple frames, its attachment to the target object is only fixed at each keyframe. Frames between keys are interpolated and might not match the target surface. If you need the object to remain on the target surface, try using a Surface Constraint, rather than attachment.

Test your animation, and either adjust the values of the keys, or add intermediate keys to better match the target surface. Too many keys can result in jittery movement of the source object, while too few keys might result in the source object missing the surface of the target object over some frames.

Changing Link Inheritance

Links can transmit transform information from a parent to a child. By default, a child inherits all of the transforms of its parent. To set an object's ability to inherit the move, rotate, and scale transforms of its parent, you use the *Inherit (page 2–500)* rollout of the Hierarchy panel. Use its settings to limit which transforms a child inherits.

The Inherit rollout is available only when you have selected a single object. When multiple objects are selected, the rollout is disabled.

Setting Link Inheritance Options

Set the Inherit options to release the link components between a selected object and its parent for any world axis of Move, Rotate, or Scale.

The options that appear when you expand the Inherit rollout are determined by the transform controllers assigned to the selected object.

The Inherit rollout for objects using standard position, rotation, and scale controllers contains three groups: one each for Move, Rotate, and Scale. Each group contains three options, one each for the X, Y, and Z world axes.

- When turned on, transform information from the parent is passed on to the child for the selected World axes.
- When turned off, transform information for the selected axes is ignored by the child.

Imagine you are animating a Ferris wheel that rotates on the world Y axis. The wheel is the parent and the cars are its children. You want the cars to ignore the Y axis rotation of the wheel. Otherwise, the passengers would all fall out. You use Link Info to turn off the Y axis option on the Rotation group of the Inherit rollout for each car.

Using the Link Inheritance Utility

The *Link Inheritance (page 2–435)* utility works exactly the same as the Inherit rollout in the Link Info group on the Hierarchy panel, except that you can set link inheritance for multiple objects in a selection set.

For example, on a Ferris wheel, you could select all of the cars and use the Link Inheritance utility to turn off Y axis rotation for all the cars at once.

Link Inheritance (Selected) Utility

Utilities panel > Utilities rollout > More button > Utilities dialog > Link Inheritance (Selected)

The Link Inheritance (Selected) utility constrains the links between multiple objects in a selection set for any axis of position, rotation, or scale.

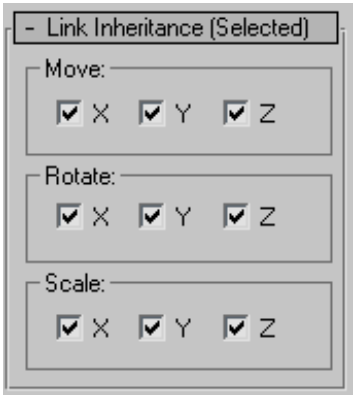
This utility works exactly the same as the *Inherit rollout (page 2–500)* in the Link Info group box in the *Hierarchy panel (page 3–773)*, except that it lets you adjust link inheritance for multiple objects in a selection set, while the functions in the Hierarchy panel work only on single objects.

Procedure

To prevent the inheritance of X rotation in a hierarchy:

- Select one or more objects, and then turn off X in the Rotate group of the Inherit rollout.

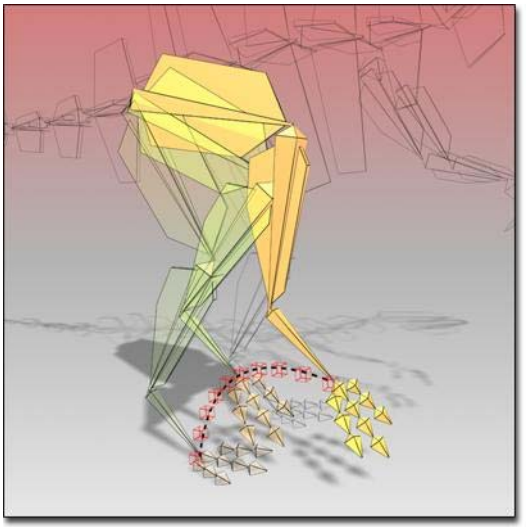
Interface



XYZ—Turn off any axis in the Move, Rotate, or Scale groups to prevent inheritance.

When a box is turned on, transform information from the parent is passed on to the selected objects for the selected axis. When a box is turned off, transform information on that axis is ignored by the selected objects.

Animating with Inverse Kinematics (IK)



Using IK to animate a leg

Inverse kinematics (*IK (page 3–958)*) is a method of animating that reverses the direction of the chain manipulation. Rather than work from the root of the tree, it works from the leaves.

Let's take the example of an arm. To animate an arm using forward kinematics, you rotate the upper arm away from the shoulder, then rotate the forearm, the hand from the wrist and so, on adding rotation keys for each child object.

To animate the arm using inverse kinematics, you move a goal that positions the wrist. The upper and lower arms are rotated by the IK solution which moves the pivot point of the wrist, called an *end effector (page 3–933)*, toward the goal.

In the case of a leg, for example, the foot is constrained to the floor by the goal. If you move the pelvis, the foot stays put since the goal has not moved, and this causes the knees to bend. The entire animation is contained in keyframes for the

goal and the root, without keys being applied to the individual chain objects.

With inverse kinematics you can quickly set up and animate complex motions. The basic procedure involves these tasks:

- Building a model. It could be a jointed structure or many pieces or a single continuous surface.
- Linking the jointed model together and defining pivot points, as described in *Hierarchies (page 2–416)*.

  For a continuous-surface model, create a *Bone (page 1–404)* structure to animate the skin of the character.
- Applying *IK solvers (page 2–440)* to either the jointed hierarchy or to the bones. You will probably create several IK chains throughout the hierarchy, rather than just one. You might also create several independent hierarchies, rather than linking everything together in one large hierarchy.
- Defining *joint behavior (page 2–483)* at the pivot points, setting limits or preferred angles, depending on the type of IK solvers you are using. Here you can set up sliding joints or rotating joints. You might also need to move the root of the hierarchy, as well. You might also want to add control objects such as dummies at this point.
- Animating the goal (in the case of an *HI Solver (page 2–446)* or *IK Limb solver (page 2–472)*) or the end effector (in the case of the *HD Solver (page 2–461)*). This animates all the components of the IK chain.

  You can apply *constraints (page 2–392)* to the goals or control objects or to the root of a chain.

You can externally reference IK chains in your scene. An XRef IK chain behaves the same as a non-XRef chain, except that you cannot retarget its *XRef controller (page 2–383)* once it is in your master scene. For more information, see *XRef Objects (page 3–394)*.

Control Objects to Assist IK

You can link a goal or an end effector to points, splines, or dummy objects that serve as quick controls to translate or rotate the end of the chain. These control objects can be linked together as well, or they can be controlled with constraints. You can also use parameter wiring to build relationships between these control objects.

You can wire control objects to *manipulator helpers (page 2–27)* or to *custom attributes (page 1–129)*, creating an easily accessible interface for your animatable model.

You can add further controls to manipulate the elements in the middle of the chain.

Note: In the *HI Solver (page 2–446)*, the *swivel angle (page 2–449)* has its own manipulator, which can be animated or linked to another target object.

Differences Between Forward and Inverse Kinematics

Forward kinematics uses a top-down method, where you begin by positioning and rotating parent objects and work down the hierarchy positioning and rotating each child object.

Basic principles of forward kinematics include:

- Hierarchical linking from parent to child.
- Pivot points defining joints between objects.
- Children inheriting the transforms of their parents.

These principles are fairly forgiving. As long as everything is linked together and the pivots are located at joint locations, you can successfully animate the structure.

Inverse kinematics (IK) uses a goal-directed method, where you position a goal object and the

program calculates the position and orientation of the end of the chain. The final position of the hierarchy, after all of the calculations have been solved, is called the IK solution. There are a variety of IK solvers that can be applied to a hierarchy.

Inverse kinematics starts with linking and pivot placement as its foundation and then adds the following principles:

- Joints can be limited with specific positional and rotational properties.
- Position and orientation of parent objects are determined by the position and orientation of child objects.

Because of these additions, IK requires greater thought about how you link your objects and place pivots. Where many different solutions for linking objects may be suitable for forward kinematics, there are usually just a few good solutions for any given IK approach.

Inverse kinematics is often easier to use than forward kinematics, and you can quickly create complex motions. If you need to edit those motions later, it can be simpler to revise the animation if you are using IK. It also is the best way to simulate weight in an animation.

See also

Inverse Kinematics Methods (page 2–439)

IK Solvers (page 2–440)

IK Terminology (page 2–437)

Bones System (page 1–404)

Using Objects as Bones (page 1–410)

IK Terminology

Using inverse kinematics requires that you set parameters for a number of IK components. Brief definitions of these components follow; details are provided in other topics.

Note: 3ds Max offers a variety of inverse kinematics systems. There are four kinds of IK solvers, plus applied IK and interactive IK systems that don't use IK solvers. Not all of the following parameters are used by all the systems.

IK Solvers—An IK solver applies an IK solution to a kinematic chain. The kinematic chain is composed of a bones system, or a set of linked objects.

Joints—An IK joint controls how an object transforms with respect to its parent. You specify joint behavior with settings in three categories:

- **Object Pivot Point—**The location of an object's pivot point defines where joint motion is applied.
- **Joint Parameters—**Changing the IK settings in the Hierarchy command panel determines the direction, constraints, and order of how the joint operates.
- **Parent Pivot Point—**The location of an object's parent pivot point defines the origin from which the joint constraints are measured.

The commands you use to place the pivot points for both the object and its parent are described in *Adjusting Pivots (page 2–423)*.

Start and End Joints—The start and end joints define the beginning and end of an IK chain managed by the IK solver. The hierarchy of the chain determines the direction of the chain. The pivot point of the end joint is displayed as the end effector, when end effector display is enabled.

Kinematic Chain—Inverse kinematics calculates the position and orientation of objects in a kinematic chain. The kinematic chain is defined as any part of a hierarchy under IK control. The IK chain starts with a selected node and consists of a start joint and an end joint. The base of the chain is either the root of the entire hierarchy, or an object

that you specify as a terminator for the chain. The kinematic chain is defined when you apply an IK solver to a chain, or when you create a bone chain with an IK solver automatically applied.

Goal—The goal is used by the HI Solver to manipulate the end of the chain. When the goal is animated, the IK solution attempts to match the end effector (pivot point of the last child in the chain) to the goal position. When using an HD Solver, the end effector fulfills the same function as the goal.

End Effector—For any IK solution, you explicitly move a control object. IK calculations then move and rotate all other objects in the kinematic chain to react to the object you move. The object that you move is the goal, in an HI Solver or IK Limb Solver, or in the case of an HD Solver, an *end effector (page 3–933)*.

Preferred Angle—Determines which direction a joint will bend. The preferred angle establishes a base angle between chain elements when an HI Solver is applied. The IK Solution seeks this angle in calculations.

Solver Plane and Swivel Angle Manipulator—A plane can be defined between the start and end joints, which aids in controlling the IK solution. Adjust or animate the solver plane by changing the swivel angle manipulator when in manipulate mode. This determines an up-vector handle for the chain, which can also be animated over time. You can also define a target for the swivel angle manipulator to follow. For more information, see *White Paper: Swivel Angle of the HI IK Solver (page 2–449)*.

Terminators—You can explicitly set the base of an HD (History-Dependent) IK chain by defining one or more objects as terminators. A terminator object stops the IK calculations so that objects higher up the hierarchy are unaffected by the IK solution. Terminator objects are also used to define hierarchies that use multiple HD IK chains.

Terminator objects are not used in HI Solvers or IK Limb Solvers. In these cases the termination is determined by the end joint of the chain.

Bound Objects—Objects in your hierarchy can be bound to the world, or they can be bound to other objects called *follow objects (page 3–943)*.

Binding allows objects in your hierarchy to be influenced by objects that are not part of the hierarchy.

- An object bound to the world will attempt to maintain its current position and orientation.
- An object bound to a follow object will attempt to match the position and orientation of the follow object.
- You can bind to a follow object when you use an HD Solver, or when Applied IK is used. If you want to bind to a follow object when using an HI or IK Limb Solver, apply a position constraint between the goal and any follow object of your choice (usually a point, dummy, spline, or bone).

Enabled IK—(HI Solver) When using the HI Solver, you can animate with either inverse kinematics or forward kinematics (FK). When Enabled is turned on, inverse kinematics controls the chain, and keyframes are placed only on the goal. When Enabled is turned off, forward kinematic rotational keys are placed on all the bones.

IK/FK Snap—(HI Solver) When using the HI Solver, if you animate using both IK and FK, you might create a situation where the goal has moved away from the end of the chain, causing a jump in the animation. Pressing IK/FK Snap will result in the goal being moved back to match the position of the end effector, eliminating the jump in the animation.

IK for FK pose—(HI Solver) When this button is turned on, moving the goal automatically sets rotation keys for the bones. In effect, this lets you use IK manipulation to create your pose, setting FK keys.

IK Methods

Inverse kinematics builds on the concepts of hierarchical linking. To understand how IK works, you must first understand the principles of hierarchical linking and forward kinematics. See *Animating with Forward Kinematics (page 2–426)* and *Inverse Kinematics (IK) (page 2–435)*.

IK Solvers

IK solvers (page 2–440) are specialized controllers that apply IK solutions procedurally across a range of frames.

Four different IK solvers ship with 3ds Max:

- HD (History-Dependent) Solver
- HI (History-Independent) Solver
- IK Limb Solver
- Spline IK Solver

Tip: In order to put an IK solver on a two-bone chain, a third bone is needed. Choose the first bone, then from the Create menu click IK solver, then click the third bone in the viewport.

In general, all IK solvers:

- Work on any hierarchy.
- Work on any hierarchy or bone structure.
- Calculate IK solutions for all frames in real-time as you make changes.
- Allow you to create multiple or overlapping IK chains within a single hierarchy.
- Display active joint axes and joint limits graphically.
- Use a node, goal or an end effector to animate the end of a chain.

An IK solver places keyframes only on the IK goal or end effector, as opposed to forward kinematics (or the non-solver IK methods), which place rotation keyframes on the bones or hierarchical objects themselves.

Other IK Methods

In addition to the IK solvers, 3ds Max provides two non–solver methods of inverse kinematic animation: Interactive IK and Applied IK. These IK methods do not apply an IK solver.

Note: These are older IK methods carried over from the earliest versions of the software. It is generally recommended that IK Solvers be used in most cases rather than Interactive or Applied IK.

- *Interactive IK (page 2–480)*

  Interactive IK gives you the ability to use IK manipulation on hierarchies without applying an IK solver. You animate your IK structure by activating Interactive IK (found in the IK panel of the Hierarchy panel and on the Hierarchy toolbar) and manually animating the position of end effectors. IK solutions are calculated only for the keyframes you set. All other motion is interpolated as set by the object's controllers. Moving the end of the chain simply adds rotational keys to the objects in the chain. The objects can have joint limits assigned, for additional control. The chain can also be terminated using the tools in the *Auto Termination rollout (page 2–499)*.

- *Applied IK (page 2–481)*

  With Applied IK, you animate follow objects and the program calculates the solution on every frame of a specified range. The IK solution is applied as standard transform

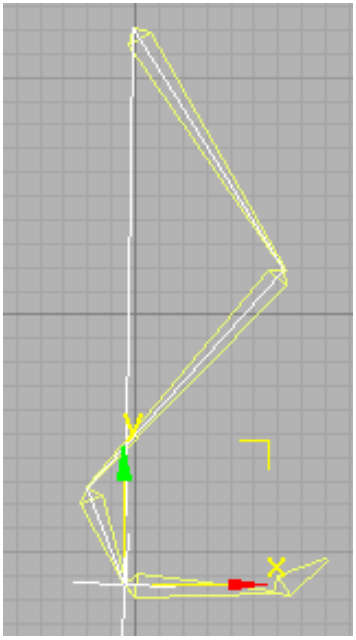
animation keys. Applied IK works with any linked hierarchy of objects except for bones that use the HI Solver or the IK Limb Solver. It can combine forward kinematics with inverse kinematics on the same objects. You can apply it automatically to a range of frames, or interactively to single frames.

Applied IK is fast and accurate, but it creates keys for every object in the kinematic chain, on every frame. The large number of keys can make it difficult to adjust the animation later. Use Applied IK repeatedly to adjust the animation.

IK Solvers

Select an object in a hierarchy where you want IK to start. > Animation menu > IK Solver > Apply an IK solver. > Click the object in the hierarchy where you want the IK chain to end.

An IK solver creates an inverse kinematic solution to rotate and position links in a chain. It applies an IK Controller to govern the transforms of the children in a linkage. You can apply an IK solver to any hierarchy of objects. You apply an IK solver to a hierarchy or part of a hierarchy using commands on the Animation menu. Select an object in the hierarchy, choose an IK solver, and then click another object in the hierarchy to define the end of the IK chain.



Bones system with HI IK solver applied

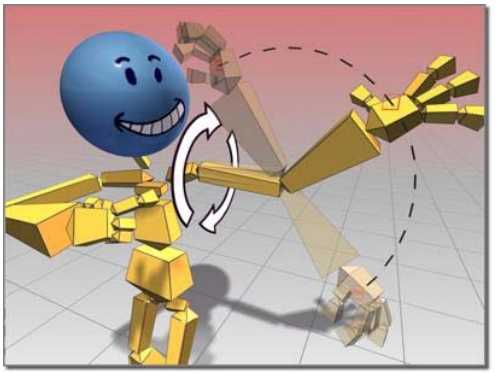
Each type of IK solver has its own behavior and workflow, as well as its own specialized controls and tools that display in the Hierarchy and Motion panels. IK solvers are plug-ins, so programmers can expand the software's capabilities by customizing or writing their own IK solvers.

3ds Max ships with four different IK solvers.

How Does an IK Solver Work?

An IK solver generally operates in this way: an inverse kinematic chain is defined on part of the hierarchy, say from the hip to the heel, or the shoulder to the wrist of a character. At the end of the IK chain is a gizmo, called the goal. The goal can be repositioned or animated over time in a variety of ways, often using linkage, parameter wiring or *constraints (page 2–392)*. No matter how the goal is moved, the IK solver attempts to move the pivot of the last joint in the chain (also called the end effector) to meet the goal. The IK solver rotates the parts of the chain to stretch out and

reposition the end effector to coincide with the goal.



Using an IK solver to animate an arm

Frequently, the end effector is constrained to the ground plane. For example, you might "pin" the toes as the heels lift in a character walk cycle. Then the movement of the root of the chain poses the legs up from the toes.

Four plug-in IK solvers ship with 3ds Max:

- *HI (History-Independent) Solver (page 2–446)*

  The HI Solver is the preferred method for character animation, and for any IK animation in long sequences. With HI Solvers, you can set up multiple chains in a hierarchy. For example, a character's leg might have one chain from hip to ankle, and another from heel to toe.

  Because this solver's algorithm is history-independent, it is fast to use regardless of how many frames of animation are involved. Its speed is the same on frame 2000 as it is on frame 10. It is stable and jitter-free in viewports. This solver creates a goal and an end effector (although the display of end effector is off by default). It uses a *swivel angle (page 2–449)* to adjust the solver plane to position the elbow or the knee. You can display the swivel angle manipulator as a handle in the viewport, and adjust it. HI IK also uses a preferred angle to define a direction for rotation, so the elbow or knees bend correctly.

- *HD (History-Dependent) Solver (page 2–461)*

  The HD Solver is a solver well-suited to use for animating machines, especially ones with sliding parts that require IK animation. It lets you set up joint limits and precedence. It has performance problems on long sequences, so ideally use it on short animation sequences. It is good for animating machines, especially ones with sliding parts.

  Because this solver's algorithm is history dependent, it works best for short animation sequences. The later in the sequence it is solving, the longer it takes to calculate a solution. It allows you to bind the end effector to a follow object, and it uses a system of precedence and damping to define the joint parameters. It allows for sliding joint limits combined with IK animation, unlike the HI IK solver, which only allows for sliding joint limits when using FK movement.

- *IK Limb Solver (page 2–472)*

  The IK Limb solver operates on only two bones in a chain. It is an analytical solver that is fast in viewports, and can animate the arms and legs of a character.

  The IK Limb solver can be used for export to game engines.

  Because this solver's algorithm is history-independent, it is fast to use regardless of how many frames of animation are involved. Its speed is the same on frame 2000 as it is on frame 10. It is stable and jitter-free in viewports. This solver creates a goal and an end effector (although the display of end effector is off by default). It uses a *swivel angle (page 2–449)* to adjust the solver plane to position the elbow or the knee. You can target the swivel angle to an another object to animate it. IK Limb Solver also uses a preferred angle

to define a direction for rotation, so the elbow or knees bend correctly. It also allows you to switch between IK and FK by keyframing IK Enabling, and it has a special IK for FK pose function so you can use IK to set FK keys.

- *Spline IK Solver (page 2–473)*

  The Spline IK solver uses a spline to determine the curvature of a series of bones or other linked objects.

  Vertices on a Spline IK spline are called *nodes*. Like vertices, the nodes can be moved and animated to change the curvature of the spline.

  The number of spline nodes can be fewer than the number of bones. This provides the ease of posing or animating a long multiple-bone structure with just a few nodes, as opposed to animating each bone individually.

  Spline IK provides a more flexible animation system than other IK solvers. Nodes can be moved anywhere in 3D space, so the linked structure can be intricately shaped.

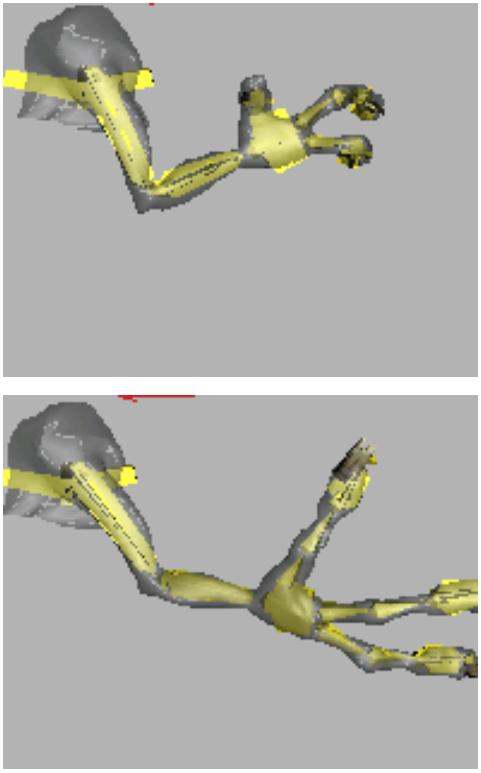
  A helper object is automatically placed at each node when Spline IK is assigned. Each node is linked to its corresponding helper, so a node can be moved by moving the helper. Unlike the HI Solver, the Spline IK system does not use a goal. The positions of nodes in 3D space is the only factor that determines the shape of the linked structure. Rotating or scaling nodes has no effect on the spline or structure.

Note: 3ds Max also provides two other methods of inverse kinematic manipulation of hierarchies, which don't depend on a solver: *Interactive IK (page 2–480)* and *Applied IK (page 2–481)*.

IK with Bones

While you can apply an IK solver to any hierarchy of objects, a system of *Bones (page 1–404)* combined with an IK solver is a good way to animate a character.

A bones system is a jointed, hierarchical linkage of bone objects. Bones are used as an armature on to which objects are linked. If you use the *skin modifier (page 1–791)*, you can "skin" an object to the bones, so the animation of the bones deforms the mesh that models a character. If you have a jointed character, you can use linkage or constraints so the bones animate the mesh.



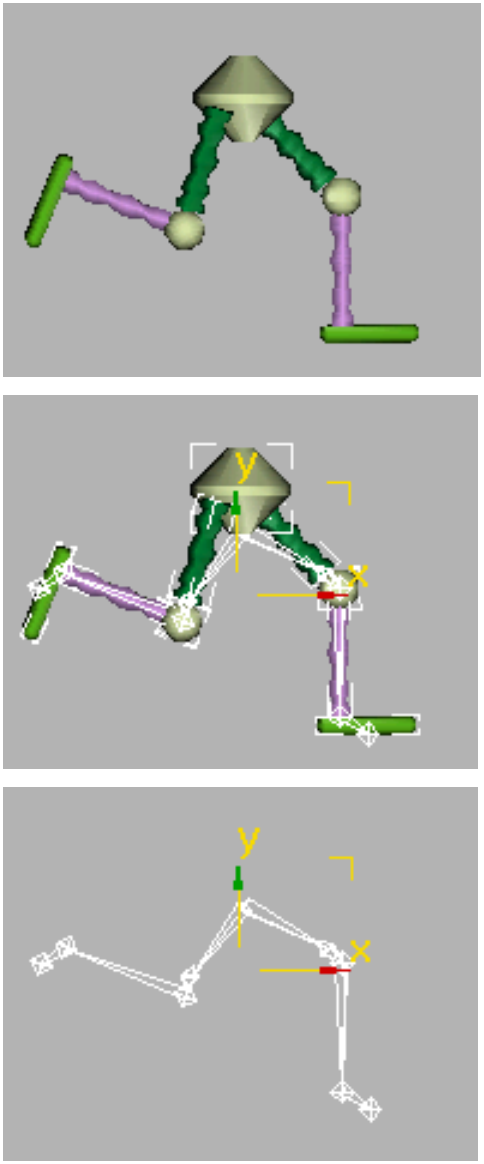
Animating bones with skin causes the skin to stretch or shrink.

Turning Other Objects into Bones

Any object can be turned into a bone object. Right-click any object and select Properties from the *quad menu (page 3–694)*. In the Bone group, choose Bone On/Off. You can then choose Show Links Only to replace the display of the objects with the bones. This can be useful if you have a geometrically intensive hierarchy to animate. The interactive viewport response will be much quicker

when the geometry is hidden and displayed only at links.

You can display any object as a bone object. Select the object and then choose Bone Tools from the Animation menu. This opens the Bone Tools floater. On the Object Properties rollout, turn on Bone On. Then go to the Display panel, and on the Link Display rollout turn on Display Links and Link Replaces Object, which displays the bones instead of the object. This can be useful if you have a geometrically intensive hierarchy to animate. The interactive viewport response is faster when the geometry is hidden and displayed as links only.



Any object hierarchy can be displayed as bones.

Bones can scale, squash and stretch over time. See *Using Objects as Bones (page 1–410)*.

Link Display

You can use Display Links and Link Replaces Object to display the links instead of the object.

These settings are found on the Link Display rollout on the Display panel. This can be useful if you have a geometrically intensive hierarchy to animate. The interactive viewport response is faster when the geometry is hidden and displayed only at links.

Advantages of Animating Bones with IK

It is possible to animate a character's motion through forward kinematics, rotating each limb into position from the shoulder to the fingers, and the hips to the toes. But it's a lot quicker and more realistic to use inverse kinematics to simulate the foot interacting with the ground. And it is a lot easier to control when you need to make changes to the animation. Rather than having keyframes on every bone in the chain, you have to make changes to only one node, to change the animation of the entire chain.

On the other hand, it is common for animators to use IK for the legs and FK for the torso and the arms. FK offers a bit more control for posing the upper body. It is not necessary to use IK for every character animation task. Using the HI IK solver allows you to jump back and forth easily between FK and IK.

How to Apply an IK Solver

You can apply an IK solver when you create a Bones system, or from the Animation menu:

- *Bones creation (page 1–404)*

  When you create bones, turning on the Assign to Children option lets you apply an IK solver immediately. Default=Off.

- You can apply the IK solver to a part of the skeleton using Animation menu > IK Solver. You must select the node where you want the IK chain to start, then choose Animation menu > IK Solvers and select the solver. Then click the node where you want the chain to end. The chain will be created between the pivot points of the two nodes. If you are using an IK Limb Solver, the IK goal will be placed two bones down from the highest node you've selected in the chain.

Where to Adjust the IK Solver

You adjust IK solver settings in the Motion and Hierarchy panels:

- *Motion panel (page 3–774)*

  When the goal of an IK chain is selected, the motion panel displays the rollouts for the individual IK solver.

  HI Solver—Motion panel parameters adjust the *swivel angle (page 2–449)* of the solver plane, to point the knees and elbows. Also, the controls in the IK Solver rollout let you mix periods of IK with FK (forward kinematics) through the Enabled button and the IK button.

  HD Solver—Motion panel parameters assign, remove, and edit the end effector for the currently selected joint. You can parent the end effector to another object, and return the skeleton to an initial pose. Changing IK controller parameters in the Motion panel affects the entire HD IK chain, even when you've selected only a single bone.

  IK Limb Solver—Motion panel parameters for the IK Limb solver are the same as for the HI Solver.

  Spline IK Solver—Motion panel parameters allow you to activate/deactivate the solver, adjust the bones assigned to the start and end joints, adjust start and end twist angles and make twist handle display settings.

- *Hierarchy panel (page 3–773)*

  HI Solver—When a bone is selected, the IK panel displays controls to activate and limit the rotational joints, and set a preferred angle for the joints. The preferred-angle poses of all the joints help control the direction of rotation of

the limbs. When a goal is selected, the IK panel is blank.

HD Solver—Select the end effector of an HD IK chain. In the Hierarchy panel, click IK. The controls that appear affect the HD Solver. You will also find the tools to bind to follow objects, and set precedence and joint limits, damping and spring back.

IK Limb Solver—When a bone is selected, the IK panel displays controls to activate and limit the rotational joints, and set a preferred angle for the joints. The preferred-angle poses of all the joints help control the direction of rotation of the limbs. When a goal is selected, the IK panel is blank. If you select a bone before applying an IK Solver, a different set of sliding and rotational joint parameters will be displayed, but these will be replaced once an IK Solver is applied.

Spline IK Solver—Select the spline of an Spline IK chain. In the Hierarchy panel, click IK. The controls that appear affect the Spline IK Solver. You will also find the tools to bind to follow objects, and set precedence and joint limits, damping and spring back. The controls in the IK panel are similar to the HD Solver.

Procedures

To add an IK solver to a hierarchy or bones system:

1. Create a bones system or any other linked hierarchy of objects.
2. Select a bone or an object where you'd like the IK chain to start.
3. Choose Animation menu > IK Solver, and then choose the IK solver:
   - HI Solver for character animation
   - HD Solver for mechanical assemblages with sliding joints
   - IK Limb Solver for two-bone chains
   - Spline IK Solver for improved control of intricate, multiple-bone structures
4. Click where you want the IK chain to end.

   If you are using the IK Limb Solver, you must apply the IK Solver to control only two bones.

   The IK solver appears in the viewport.

To create a bones hierarchy that uses an IK solver:

1.  Go to the Create panel, choose Systems, and click Bones.
2. On the IK Chain Assignment rollout, choose an IK solver from the list.
3. Turn on Assign To Children.
4. Click and drag in a viewport to create the bones. Right-click to stop bone creation.

   The bones are created with the IK solver already applied.

   Note: If you use the Spline IK Solver, a Spline IK Solver dialog will open where you can make special settings for the spline and helpers used by the Spline IK solver.

To display a hierarchy of objects as bones:

1. Select the hierarchy of objects in the viewport.
2. From the Animation menu > Character submenu, choose Bone Tools.

   This Opens the bone Tools dialog.
3. Expand the Object Properties rollout.
4. In the Bone Properties group, turn on Bone On.
5. On the Display panel, scroll down to Link Display and expand it.
6. On the Link Display rollout, turn on Display Links, and Link Replaces Object.

   The objects disappear and the links are displayed as bones.

History-Independent (HI) IK Solver

The HI (History-Independent) Solver does not rely on IK solutions calculated in previous keyframes in the timeline, so it is just as fast to use at frame 2000 as it is at frame 20.

The HI Solver uses a goal to animate a chain. You animate the goal and the IK solver attempts to move the end effector (the pivot point of the last joint of the chain) to match the position of the goal. Often the goal is parented to other control objects such as points or dummies, splines or bones, and these control objects in turn are wired to viewport or rollout sliders.

The IK solution takes place in a plane, known as the *solver plane*. The angle of the solver plane in world space is controlled by a parameter called the *swivel angle (page 2–449)*. The swivel angle is animatable. You can adjust it directly, or with a manipulator.

The HI Solver allows for the creation of multiple or overlapping chains. This allows you to create multiple goals for additional controls. By linking the goals to points, splines, bones or dummies, you can create simple controls to animate complex chains or hierarchies. You can also use constraints on these goals or control objects, as another animation tool.

Applying an HI Solver

To apply an HI Solver to any part of a hierarchy select the bone or object where you want to the solver to start. Then choose Animation menu > IK Solvers > HI Solver. In the active viewport move your cursor to the bone where you want the chain to end. When you click to select that bone, the goal is drawn at the pivot point of that bone. If you want a goal at the far end of the bone, refine the bone where you want to goal to be placed. An extra bone will be added, and then choosing that bone allows you place the goal at the end.

When you create bones, a small "nub" bone is automatically created at the end of the chain to assist in this process.

Setting Up Multiple Chains

To rig a skeleton for a human leg you could use three chains in one leg, as follows:

- The first chain is created from the hip to the ankle. This chain controls the overall leg motion including bending of the knee.
- The second chain is created from the ankle to the ball of the foot. This chain controls the heel's up and down motion.
- The third chain is created from the ball of the foot to the toe.

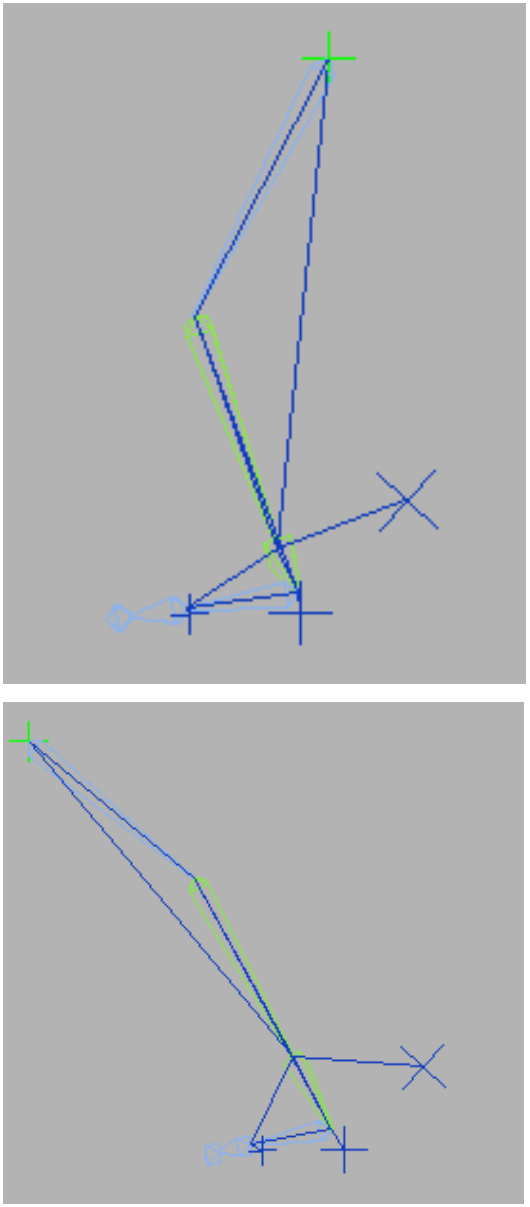
When the three chains work together they help to maintain the foot's position in space. This means it will keep the foot planted on the ground as the character's body moves. All three IK chains in this hip-to-toe setup place goals at key positions in the foot that mimic natural foot behavior. In real life, the toe, ball of the foot, and heel can be planted on the ground or raised.

Each chain has an goal that drives motion on the heel, ball of foot, and toe. Use the IK goals to raise the heel, bend the toe, move and rotate the entire foot, and maintain the foot's position in space.

Overlapping Chains

The IK solver system allows you to create overlapping IK chains in a single hierarchy. In a human leg, for example, you could create a chain running from the hip to the ankle, then a second chain from the knee to the ball of the foot, and a third from the ankle to the toes. Use overlapping chains, when you want to apply goals to sequential bones, but you don't want to refine the bones.

Another good use for overlapping HI IK chains is to keep the goals at the bottom of the hierarchy firmly rooted in place. In the case of the leg example, the ankle, ball and toe of the foot would not move until the upper portion of the hierarchy had reached its full stretch.



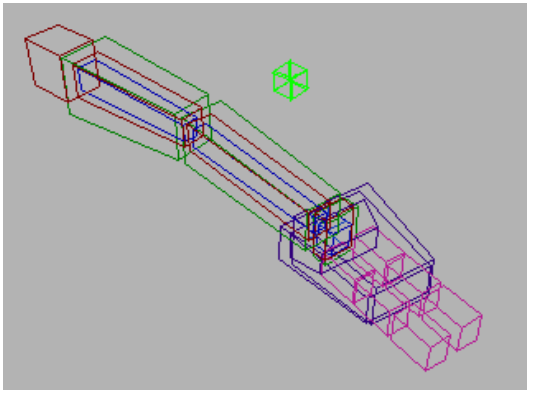
Overlapping chains make the toe stick to the ground

Don't try to overlap chains with different kinds of IK solvers or you might get unpredictable results.

Creating Control Objects

To create control objects to animate the goals, create dummies, points, splines or other objects near the goal, then link the goal to the control object. For example in a human leg, you might have a goal at the ankle, ball of the foot, and the toe. You then create three splines under the foot, one for the heel, one for the ball and one for the toe. Link each goal to each spline, then you can animate the goals using these splines. You could also link the splines together, so that the rotation of the toe is controlled by the movement of the heel.

Besides hierarchical linkage, you can now use the new Constraint system in conjunction with IK solvers. You could apply any of the constraints to create a relationship between the goals or bones and other objects. The goal could be position constrained to a dummy which is moved. Or you could create three bone chains that are all in the same place, and constrain one bone chain to another, then weight the constraints.



Bone chains constrained together

Using control objects gives you something bigger to select in the viewport. You can also use control objects to separate chains, for example in a human arm one chain could end at the wrist and another separate chain be created for the hand and fingers. The control object at the wrist serves as the root node for the hand chain, yet that chain would remain disconnected (hierarchically speaking) from the arm chain.

You can create viewport sliders using the manipulator helper, and then use the viewport sliders to control the transforms of the control objects. Use *wire parameters (page 2–411)* to hook up the sliders with the control objects. You can also create Custom Attributes to add these sliders to the object rollouts.

Mixing Forward Kinematics with Inverse Kinematics

The HI IK solver provides a tool for mixing FK and IK in a single animation track. There is an FK subcontroller beneath the IK controller assigned by this solver.

When the Enabled button is on, the FK subcontroller values are preserved but ignored. When Enabled is turned off, the FK subcontroller values apply. To access the Enabled button, select the goal and go to the Motion panel, then turn off the Enabled button. This will allow you to animate using FK rotations of the bones or hierarchy objects.

IK for FK pose allows one to turn on IK in middle of FK manipulation. When the Enabled button is turned off, and IK for FK Pose is on, then selecting and moving the goal lets you use IK to create the forward kinematic keyframes. Moving the goal poses the skeleton and add rotation keys to all the objects in the chain when the Auto Key button is on.

When working with IK and FK together it is possible to create a situation where the goal has moved away from the end of the chain. Use the IK/FK snap button to reposition the goal, snapping it back to the end of the chain. When AutoSnap is on, the snap happens automatically; when you touch the goal, you don't have to click the IK/FK snap button.

Controlling HI IK Precision

When you are animating with HI Solvers, if you find the animation of the limbs is not smooth, you should try doubling the Iterations in the Solutions group of the *HI Solver Properties rollout (page 2–456)*. You can also try reduce the Thresholds value to smooth the animation.

See also

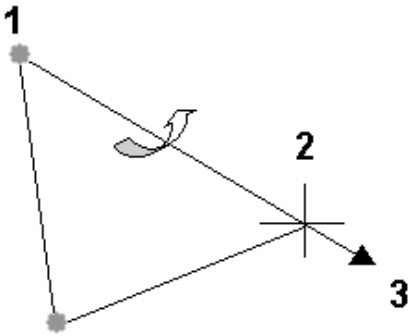
IK Solver Rollout (HI Solver) (page 2–453)

White Paper: Swivel Angle of the HI IK Solver

The fundamental requirement of HI Solver is that the solution be history-independent: the solution has to be based on the goal and other incidental parameters solely at their current states.

Swivel Angle Degree of Freedom

When the positional goal is given for a single chain, there remains an obvious degree of freedom: the rotation about the End Effector Axis (EE Axis). The swivel angle is used to describe this degree of freedom quantitatively.



1. *Start joint*
2. *End effector*
3. *EE axis*

Solver Plane

Let's call the plane passing all the joints the Solver Plane. When joints do not lie on a plane, we will define it to be the plane that (A) passes the Start Joint and End Joint and (B) is closest to the remaining joint in a certain sense.

The Swivel Angle describes the degree of freedom of the Solver Plane and it constrains only the Start Joint.

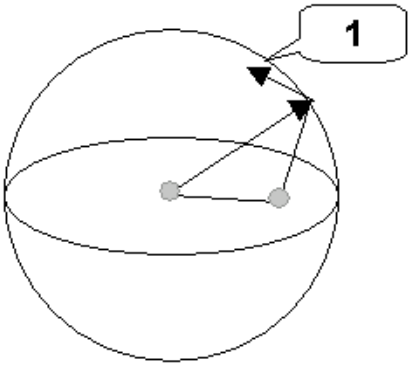
Zero Plane Map

In order to describe the solver plane in terms of a numerical quantity, we have to agree to what 0 means. Given the end-effector position, where is the Zero (Solver) Plane? The Zero Plane Map takes as the argument EE Axis and produces the normal to the zero plane.

The IK system allows individual solver plug-ins to define their own Zero Plane Maps. When not defined, the IK system provides a default one.

The argument to the Zero Plane Map is a unit vector to give the direction of the EE axis. Equivalently, when the EE slides along the EE

axis, the solver plane should be fixed. Therefore, the Zero Plane Map defines a vector field on a sphere. Given a point on the sphere, it produces a tangential unit vector to be interpreted as the normal to the zero plane.



1. Normal to the zero plane

Solver Plane Flipping

It is a mathematical fact that there does not exist a continuous vector field on a sphere. No matter how hard you try, there will always be a point on the sphere where neighboring vectors change dramatically. This is where the solver plane will flip when the end effector axis approaches to it.

This is because, on one hand, the history independent requirement demands us to assign a fixed vector to the singular point. On the other hand, no matter what vector is assigned, it will be dramatically different from some vectors assigned to the neighboring points.

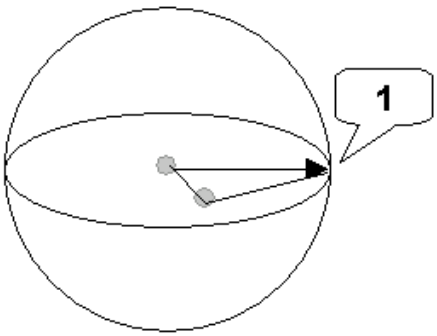
Intrinsic Reference Frame for the Sphere

In order to define the Zero Plane Map, we need to define a reference frame for the sphere. This reference frame is intrinsic to the joint chain itself.

A sphere can be defined by the center, the horizontal plane, and the meridian (zero longitude). The center is assigned to the start joint. The pose when all the joint angles assume preferred angles is particularly important. Let's call it the preferred pose.

We use the solver plane at the preferred pose as the horizontal plane. Since the swivel angle is used to control the start joint, the preferred angles at the start joint are not so intrinsic. It is also reasonable to define the horizontal plane with the solver plane that is derived by assigning zeroes to the start joint and preferred angles to the other joints.

The EE axis defines the meridian. The sphere is now defined as shown in the following figure:



1. EE axis

All the joints assume preferred angles. The Zero Plane Map is to be defined on this sphere.

The API for the plug-in solver to define its own Zero Plane Map in fact takes the EE axis and the normal to the solver plane at the preferred pose:

```
virtual const IKSys::ZeroPlaneMap*

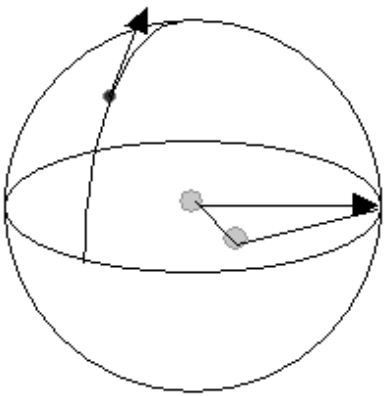
   GetZeroPlaneMap(const Point3& a0,
   const Point3& n0) const
```

where *a0* and *n0* are the EE axis and solver plane at the preferred pose, respectively. Object of ZeroPlaneMap is a function that assigns a plane normal to each point on the sphere.

Default Zero Plane Map

When not provided by plug-in solvers, (the IK Solver itself is implemented as a plug-in solver) the IK system will provide a default one. This map is defined by the following rules:

- **A:** For each point on the equator, the intersection of the horizontal plane and the sphere, the normal vector is defined as the vertical vector, pointing to the same direction as the normal of the solver plane at the preferred pose.
- **B:** For any point on the sphere other than the north or south poles, there is a great circle that passes the point and the north, south poles. This circle hits the equator at two points. One point is closer to the given point. The normal vector at the given point is defined as derived from moving tangentially the normal at the closer point on the equator along the great circle to the point.



Deriving the default normal to the zero plane

Obviously, this method won't extend to the north or south poles. They are the singular points. When the EE axis moves across the poles, the normal will suddenly change direction: it flips from the users' viewpoint.

Normally, the preferred pose is the one when the solver is first assigned. So, the plane on which one lays the joints corresponds to the horizontal plane here. Rule A makes sure that the chain will stay on the plane if one moves the goal on the plane.

Rule B means that, when you move the goal along the great circle vertical to the equator, the chain will stay vertical, except when it passes through the poles, which are the singular points of this map.

Parent Space

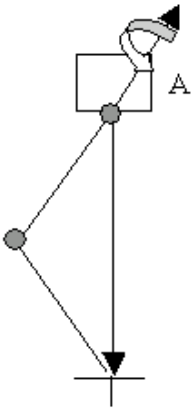
So far, we have described things as if the whole world comprises only IK elements. In practice, the IK chain and goal might sit at points of separate transformation hierarchies. Ultimately, we need to map the position of the end effector that is described in the world to a point on the sphere. Depending how the sphere is mounted relative to the end effector position, the readings of latitude and longitude are different. The parent transformation space that this sphere is to be placed in is called the Swivel Angle Parent Space, or Parent Space when the context is clear.

The parent space has to be invariant with regard to the IK parameters. Right now, we provide two choices:

- **Start Joint—**The Swivel Angle Parent Space is the same as the parent space of the Start Joint.
- **IK Goal—**The Swivel Angle Parent Space is the parent space of the IK Goal.

Example 1

If both the start joint and the goal are rooted directly at the world, the choice of Parent Space does not give rise to any difference. In the following example, the start joint is parented to object A.



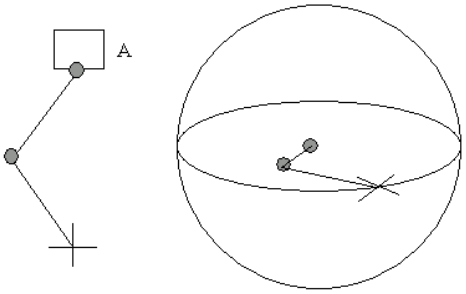
The IK chain is parented (via the start joint) at object A.

Assume this is the pose when the IK solver is assigned. So, this is the preferred pose. The plane on that the joints are laid out is the horizontal plane of the (Zero Plane Map) sphere.

- **A:** Parent Space is Start Joint. In this case, the sphere is parented to A. If A is rotated about the drawn axis, the sphere is rotated together with it. The goal is in a separate transformation hierarchy. It stays in place, and the end effector sticks to it because of the IK solution. Since the (plane) normal is fixed to the sphere, it rotates with A, too. Therefore, the whole chain appears to be rotated together with the parent object.
- **B:** Parent Space is IK Goal. Suppose that the goal is parented to the world. In this case, the sphere is parented to the world and, hence, stays fixed. Since the normal is fixed to the sphere, the chain will appear stationary when A is rotated.

Example 2

In the following example, we look at a case where there exists a rotation in the parent space when the IK solver is assigned.



The parent space of the IK chain contains a rotation when the IK solver is assigned.

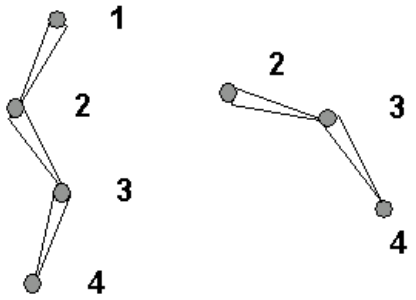
Parent A contains a rotation of 90 degrees. This is an abstraction of the case when a user creates four bones without an IK solver and later assigns an IK solver from Bone2 to Bone4. If we parent the chain directly to world, it would appear as shown in the right figure: the solver plane becomes horizontal.

- **A:** Parent Space is Start Joint. The sphere is mounted after the rotation of A and therefore the "horizontal plane" coincides with the chain plane as shown in the viewport. The singular points are perpendicular to the Start Joint with regard to the plane. When the goal/end effector is moved on the plane, it will never hit the singular points and flip.
- **B:** Parent Space is IK Goal. Again, we assume that the parent space of the goal is the world. The "horizontal plane" of the sphere becomes horizontal, as shown in the right figure. The singular points, the poles, are on the plane that joints are laid out. Therefore, when users move the goal/end effector left to right, or the other around, the end effector will move across the singular point and flip.

A problem of B is that the figure on the right is never shown to the user. They have to envision it in order to understand the flipping.

Example 3

This example describes what happened when Start Joint is reassigned. Suppose we have an IK chain of four bone nodes.



1. Bone01

2. Bone02

3. Bone03

4. Bone04

The Start and End Joints are Bone01 and Bone04, respectively. Suppose the pose shown in the figure is the preferred pose and Bone01 contains a rotation. If we parent Bone02 directly to the world, the hierarchy from Bone02 will appear as in the right figure.

When we reassign Start Joint to Bone02, the Zero Plane Map sphere will be based on the configuration on the right.

- **A:** Parent Space is Start Joint. In this case, the sphere that is fixed to the right configuration is to be parented to Bone01 together. This will be the same as the one based on Bone01 to Bone04. Therefore, we get the same normal and chain stays stationary.
- **B:** Parent Space is IK Goal. In this case, the sphere stays in the world while Bone02 is parented to Bone01. When Bone04 is moved to where it is shown on the left due to the transformation of Bone01, the observer who is fixed to the sphere that is fixed to the world will see that Bone04 moves away from the neutral position defined by the preferred pose. The algorithm then applies the Zero Plane Map to find the normal at the displaced position. It generally is not perpendicular to the viewport. Therefore, we will see a flip once it takes a new Start Joint.

If you delete the solver/goal and assigned a new one from Bone02 to Bone04, you will find that the chain won't flip. Why? Assignment of Start Joint is different from creating a new IK chain/goal. Start Joint is one of many IK parameters. Reassigning it is simply the same as modifying any parameter. The rest parameters are intact. In particular, the Swivel Angle is not changed as a result of this reassignment.

Creating a new IK chain/goal is different. Effort is made to ensure that the joint chain stay fixed by adjusting parameters appropriately. In particular, the Swivel Angle will be set to a value so that the solver plane keeps stationary in the viewport.

HI IK Solver Rollouts (Motion Panel)

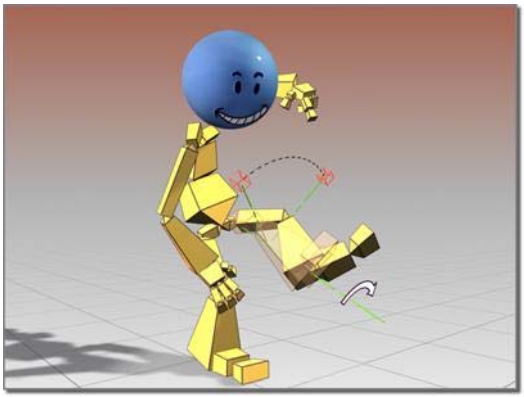
IK Solver Rollout (HI Solver)

Select an HI or IK Limb Solver IK Chain control (the blue cross at the end of the chain.) > Motion panel > Parameters button > IK Solver rollout

The HI Solver is a history-independent solver that doesn't rely on the calculations from previous frames for the IK solution, so it is fast to use no matter the length of your animation. The history-independent solver uses a goal to manipulate the end of the chain. It uses a preferred angle that specifies a preference regarding which direction the link will rotate: positive or negative.

The preferred angle can also be considered as the initial angle; that is, the angle at which the link was rotated at the time the solver was applied.

The IK solution takes place in a plane, known as the *solver plane*. The angle of the solver plane is controlled by a parameter called the *swivel angle*.



Changing the swivel angle

The angle of the solver plane is actually calculated in one of two coordinate systems: the Start Joint Parent space or the IK Goal Parent space. World space is not an explicit option; however, you can easily configure the IK chain to work in world space by choosing the IK Goal Parent option and making sure that the IK goal is unlinked (i.e., has no hierarchical parent). In this case, the IK goal's parent is the world, and therefore, the solver plane will be computed in world space.

The swivel angle is animatable. You can adjust it directly, or with a manipulator. Or you can target the swivel angle to a object and animate that to affect the swivel angle. These parameters are on the *IK Solver Properties rollout (page 2–456)*.

The HI Solver is designed to let the animator jump back and forth between forward and inverse kinematics quickly and conveniently, with tools to automatically enable IK and snap the goal to the end effector.

Procedures

To turn off IK on a chain:

1. Select the goal of a chain with an HI Solver.
2. In the Motion panel IK Solver rollout, turn off Enabled.

   IK is now off, so you can select and rotate any of the objects in the chain. Turning IK on and off is animatable using the Auto Key button.

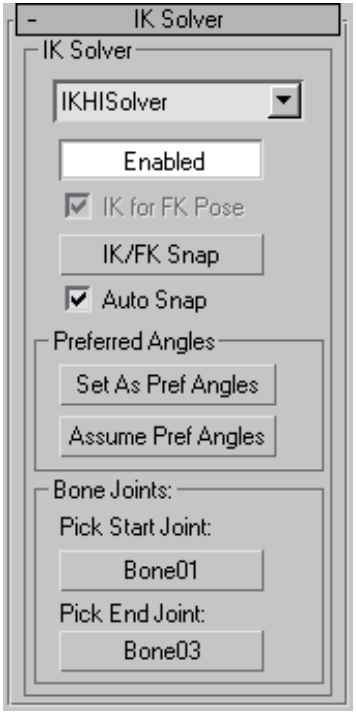
To mix IK and forward kinematics (FK) in a single animation track:

1. Apply an HI IK Solver to a hierarchy or bone system.
2. Turn on the Auto Key button and move the time slider ahead in time.
3. Select the goal of the IK Chain, and open the Motion panel.
4. Select and move the goal to animate the IK chain with Inverse Kinematics.
5. Advance the time slider again.
6. In the IK Solver Rollout, turn off Enabled and IK for FK Pose.
7. Select and rotate the root. The entire hierarchy rotates freely. You are now adding keyframes using forward kinematics.

To match the goal and the end effector positions:

1. If the goal and end effector become separated, click IK/FK Snap and the goal will move to match the end link of the chain.
2. If AutoSnap is turned on, clicking the goal will automatically perform the snap action, matching the goal and end effector positions.

Interface



IK Solver group

The items in this group provide the ability to set the start and end points of the selected HI IK solver chain. There are also controls in this rollout that allow you to use IK manipulation to create forward kinematic rotational keyframes on the hierarchy objects, and there are buttons to align the goal and the end effector.

Solver field—Allows selection of IK Solver. Choose between the HI IK Solver and the IK Limb Solver here. Any HI IK plug-in solver present at startup will appear in this list, as well.

Enabled—Turns IK control of the chain on and off.

The HI IK Controller has an FK subcontroller. When Enabled is selected, the FK subcontroller values are overwritten by the IK controller. When Enabled is turned off, the FK values are used. You can animate Enabled On and Off.

Use this to turn off the chain control by the goal, when you want to do forward rotations.

IK for FK Pose—Lets you turn on IK in middle of FK manipulation. When Enabled is off and IK for FK Pose is on, moving the goal will turn on IK automatically in the middle of an FK manipulation. The result of this is that all the FK subcontrollers receive values from the IK solution. Keys are placed on the hierarchy objects or the bones, not on the goal. When both Enabled and IK for FK Pose are turned off, moving the goal does not affect the chain at all.

IK/FK Snap—Performs an IK snap when in FK mode, and an FK snap in IK mode.

IK Snap—If the goal has moved away from the end of the chain, clicking IK/FK Snap moves the goal to coincide with the end link position.

FK Snap—The values of FK subcontrollers are suppressed by the solution of IK when IK is on (Enabled is on). Their values don't always correspond to the current pose of the chain. If you turn off Enabled, the values of FK subcontrollers will suddenly take over. This can cause the chain to jump. FK Snap, before Enabled is turned on, forces the FK subcontrollers to assume values from the current pose (which is determined by the IK.) This eliminates the jump in the chain's animation.

Auto Snap—When Auto Snap is on, the software automatically applies an IK/FK snap before you turn Enabled on or off. If Auto Snap is not on, you need to click IK/FK Snap before toggling Enabled, otherwise the chain will jump.

Preferred Angles group

Set As Pref Angles—Sets the preferred angle for each bone in the HI IK chain. The current parent-space rotation of each bone is placed into its X, Y, and Z Preferred Angle rotation channels as seen in the

Rotational Joint rollout of the Hierarchy|IK panel. This is useful for establishing a perfect *match frame (page 3–969)* when transitioning between forward and inverse kinematics.

Assume Pref Angles—Copies the X, Y, and Z preferred angle channels of each bone and places them into its FK rotation subcontroller. This essentially performs the inverse operation of the Set As Pref Angles function.

Bone Joints group

Allows you to change the ends of the IK chain.

Pick Start Joint—Defines one end of the IK chain. Select from the viewport or by name (press H).

Pick End Joint—Defines the other end of the IK chain. The direction of the chain is defined by the hierarchy, not by Start and End Joint. Select from the viewport or by name (press H).

Warning: If you change the Start or End Joints so the IK Limb Solver so there are more than two bones between the start and end, the solver will not work. Moving the IK goal does not affect the bones.

Tip: The order of the hierarchy determines the direction of the chain. You cannot reverse the direction of the chain by picking Bone 10 as the start and Bone 1 as the end.

IK Solver Properties Rollout (HI Solver)

Select end effector of an HI IK chain. > Motion panel > Parameters button > IK Solver Properties rollout

Here are the additional controls for the HI IK solver. The IK Solver Plane is controlled here through the Swivel Angle, which can be animated directly or using a target object.

In the Parent Space group you can choose to whether the IK goal or the start joint will be used as the parent space for the swivel angle. If IK Goal is selected (which is the default), then the Swivel Angle of the chain is defined in the parent space of its goal. If Start Joint is chosen, then the Swivel Angle will be relative to the parent space of the start joint. These two options allow a much better control of the swivel angle on a chain with two HI IK solvers. For example, the first solver on the upper part of a chain can use a Swivel Angle in the Parent Space of the Start Joint, while the second IK solver on the lower part of the chain uses the Parent Space of the IK Goal. In this case, the change of the Swivel Angle for one chain won't change the orientation of the other one.

In the Thresholds and Solutions group you can refine the IK solution. The Position Threshold sets a limit in units of how far the goal can be moved from the end effector, the Rotation does the same in angle degrees.

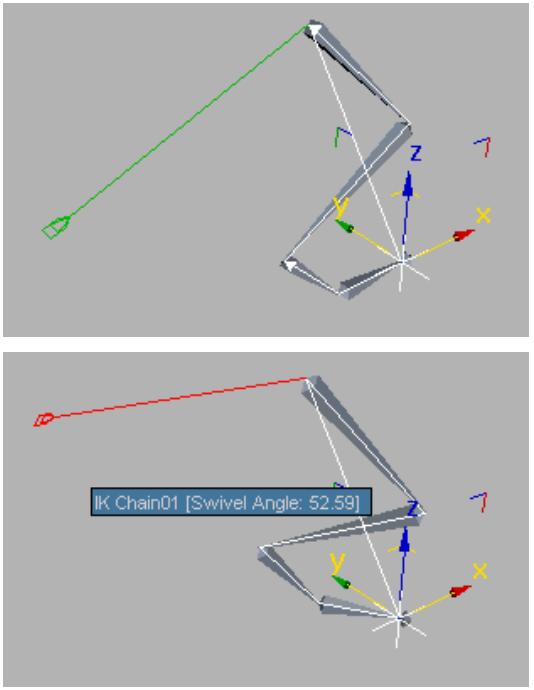
The Solutions group Iterations is the maximum number of attempts the IK solution will perform before giving up.

If your IK animation seems to jitter, try increasing the Iterations to 200. Increasing the Iterations will slow down the computational time, but should produce better results.

Manipulator for the Swivel Angle



You can adjust the swivel angle with the spinner in this rollout, or you can adjust it with a manipulator. To do so, turn on *Select and Manipulate (page 2–15)*. The swivel angle manipulator is displayed as a handle in the viewports.



Swivel angle manipulator.

When you move the mouse over the manipulator, it turns red. At this point, you can drag the manipulator to change the swivel angle. The Auto Key button can be active while you do this.

You can also animate the swivel angle by using a target object. Pick an object for the target, then animating the object will animate the swivel angle.

The swivel angle manipulator is a standard *plane angle manipulator (page 2–29)*. The HI Solver has its own controls for displaying this manipulator. These are on the *IK Display Options rollout (page 2–458)*.

Procedure

To increase the precision of an IK solution:

1. Select the goal in an HI IK chain.
2. In the IK Solver Properties rollout increase the Iterations to 200 in the solutions group.
3. Play or render the animation.
4. You can also adjust the Position and Rotation Thresholds to a smaller value.

Interface



IK Solver Plane group

Swivel Angle—Controls the Solver Plane which determines the direction of the knee or elbow in a human limb. The Swivel Angle can be manipulated in the viewport by turning on Swivel Angle Display in the Display Options rollout, and then turning on Manipulate Mode. When the swivel angle manipulator is displayed in the viewport, you can interactively animate the handle to animate the solver plane.

Pick Target—Lets you select another object to use to animate the Swivel Angle.

Use—Turns the use of the Target on and off.

Parent Space—Defines what space the Swivel Angle is relative to. Use this on chains with multiple solvers, so the swivel angle manipulations on one chain will not affect the second. For more

information, see *White Paper: Swivel Angle of the HI IK Solver (page 2–449)*.

IK Goal—Defines the Swivel Angle parent space relative to the IK Goal.

Start Joint—Defines the Swivel Angle Parent Space relative to the Start Joint.

Thresholds group

Use to define the tolerances the system uses for its calculations.

Position—Sets a limit in units of how far the goal can be moved from the end effector. Keep these numbers low.

Rotation—Sets a limit in degrees of how far the goal can be rotated away from the end effector rotation. Keep these numbers low for best results.

Note that adjusting the Rotation threshold on the HI Solver and IK Limb Solver has no effect. It is up to the plug-in solver's discretion as to whether or not they will support a rotation component in the IK solution.

Solutions group

Use to control the precision of the animation. Increase the iterations when the IK solution produces an animation that appears rough or jittered.

Iterations—The number of attempts an IK solution will make to find the best match between the goal and the end effector positions. Increase this number if you are seeing jumpy animation.

IK Display Options Rollout (HI Solver)

Create a bone chain > Add an HI Solver. > Motion panel

Here are the controls to turn on and off a variety of gizmos in the viewport. The Goal and End effectors, the vector handle of the swivel angle manipulator, and the IK solver display itself, are all located in this rollout.

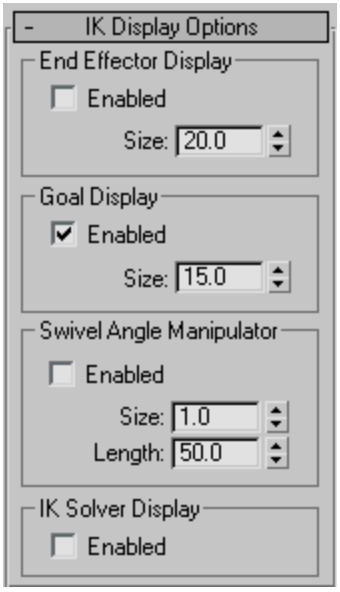
Procedure

To see multiple IK solvers in a single chain:

1. In a chain with multiple IK solvers, select a goal in the viewport.
2. In the IK Display Options rollout, turn on IK Solver Display Enabled.
3. Select a second goal and repeat the process.

   Each goal's display remains on whether it is selected or not.

Interface



End Effector Display group

Controls the appearance of the end effectors in the IK chain. Off by default.

Enabled—Turns the end effectors display on or off.

Size—Controls the size of the end effector gizmo in the viewport.

Goal Display group

Controls the appearance of the goals in the IK chain. On by default.

Enabled—Turns the goals display on or off.

Size—Controls the size of the goal gizmo in the viewport.

Swivel Angle Manipulator group

Controls the display of the swivel angle manipulator in the IK chain. Default = on. Turn this on, then turn on Manipulate mode to see the swivel angle manipulator.

Enabled—Turns the swivel angle manipulators on or off.

Size—Controls the size of the manipulator's handle in viewports.

Length—Controls the length of the manipulator in viewports.

IK Solver Display group

Controls the appearance of the IK Solver display, the line drawn between the start and end joints. Turn this on when you have multiple chains that you want to see at the same time.

Enabled—Turns on or off the IK Solver display.

Sliding and Rotational Joints (HI Solver)

Create a bone chain. > Apply an HI Solver > Select any bone in the chain. > Open the Hierarchy panel. > IK button

Here are the controls for setting the joint limits when using the HI Solver. You can turn X, Y, or Z axes on or off, limit them, and set those limits here. This is also where the preferred angle is defined, and adjusted. The preferred angle defines the direction a chain will bend, which angle a chain will tend to rotate toward.

Note: Sliding joints are implemented in the HI Solver, but they do not take part in an IK solution. They do, however, take part during FK interactive manipulation.

Use the HD Solver if you need to animate sliding joints with inverse kinematics.

Procedures

Example: To set the direction an IK chain will bend (preferred angle):

To understand setting the preferred angle it helps to use a simple example.

1. Go to Create panel > Systems and turn on Bones.
2. Turn on Assign to Children and make sure the HI Solver is displayed in the Solver name field.
3. In the Top viewport, create a bone chain of four bones in a straight line .
4. Turn on the Auto Key button, and move the time slider to frame 50.
5. Select and move the goal so it is close to the root node.
6. Select any bone in the chain.
7. On the Hierarchy panel > IK panel, open the Rotational Joint rollout.
8. There are three axes. Find the one with the preferred angle that is not zero.
9. Change the preferred angle. If it is negative make it a positive one.
10. Play the animation, see how changing the preferred angle can redefine the direction of the rotation.

To set rotational joint limits on a hierarchy of objects or a bone chain:

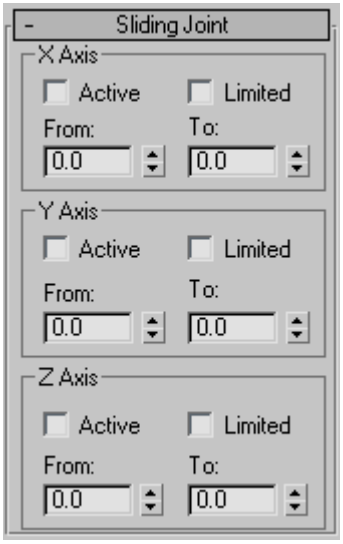
1. Create a bone chain or a hierarchy of objects.

2. Select any object in the chain.
3. On the Hierarchy panel > IK panel, open the Rotational Joint rollout.
4. If you want to limit an axis, turn on Active and Limited, and set the values of From and To. The hierarchy will rotate in the viewport.

   Tip: Pressing the From and To labels instantly rotates the bone to those angles.

Interface

Sliding Joints rollout



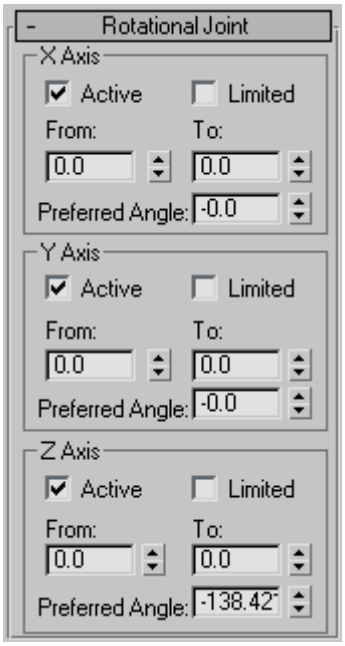
X, Y, Z Axis groups

Active—Activates an axis (X/Y/Z). Allows the selected object to slide on the activated axis.

Limited—Limits the range of motion allowed on an active axis. Use in conjunction with the From and To spinners. Most joints are limited in their range of motion along an active axis. For example, a piston slides only within the length of its cylinder.

From and To spinners—Determine for position limits. Use in conjunction with the Limited function. Clicking the labels From and To moves the object to its limit position. Use this to check the sliding limits on an object.

Rotational Joints rollout



X, Y, Z Axis groups

Active—Activates an axis (X/Y/Z). Allows the selected object to rotate about the activated axis.

Limited—Limits the range of rotation allowed on an active axis. Use in conjunction with the From and To spinners. Most joints are limited in their range of rotation along an active axis. For example, a wheel rotating around an axle could be limited to rotate around only one axis.

From and To Spinners—Determine for rotation limits. Use in conjunction with the Limited function. Clicking the labels From and To rotates the object to its limit rotation. Use this to check the rotation limits on an object.

Preferred Angle— Used by the HI Solver to set the initial pose for internal calculations. Determines the direction of rotation preferred by a bone.

History-Dependent (HD) IK

Select a node of hierarchy or bone chain. > Animation menu > IK Solvers > HD Solver > Click a second node of chain.

Animating with the HD Solver lets you use sliding joints combined with inverse kinematics. It has controls for spring back, damping, and precedence not found in the HI Solver. It also has quick tools for viewing the initial state of the IK chain. Use it for animations of machines and other assemblies. Since this is history dependent, performance is slower at the end of long animations. For lengthy scenes, use HI Solvers, if possible.

Animating the End Effectors

You animate a hierarchy of objects or bone structure using HD Solvers by animating special end effectors located at the bone joints. There are two types of end effectors: Position and Rotation. They are displayed as three intersecting blue lines at the joint.

When you select and transform a joint that carries either end effector, only the end effector itself is transformed. The objects or bones in the chain then use IK to calculate the IK solution.

Animating these end effectors is much like animating with Interactive IK except interpolation between keyframes uses correct IK solutions.

Linking End Effectors to a Parent

You can link HD IK end effectors to a parent to achieve results similar to using bound follow objects with Interactive and Applied IK. To link a HD IK end effector, you click Link on the Motion panel.

There are two reasons to link an end effector to a parent:

- You can rescale your entire hierarchy. If you animate a robot, and then later decide you want to scale the entire robot and its animation, link the robot's root object to a dummy, link all of the robot's end effectors to the same dummy, and then scale the dummy to scale all of the hierarchical objects as well as the animation of the end effectors.
- You can use linked end effectors for things like feet or hands. End-effectors are automatically bound to the world, so when you move a root object in your hierarchy, the end effectors remain behind. This is good for keeping feet on the ground, but not very good for hands, where you want them to move along with the character. To have hands follow the root object, link the end effectors of the hands to the root object.

Linking Bones to Follow Objects

You can also use follow objects instead of end effectors and bind bones to any number of follow objects using Interactive IK.

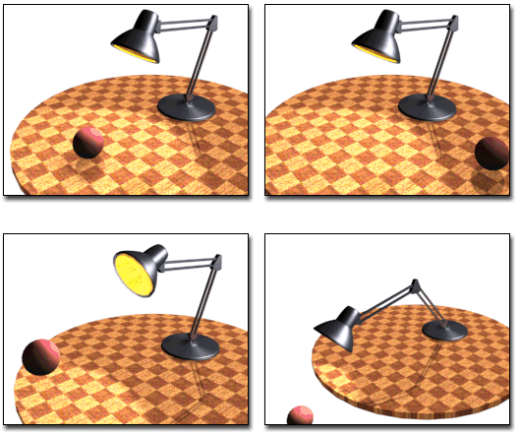
If you use follow objects with HD IK, you do not need to click Apply IK. The IK solver takes care of everything automatically.

See also

Inverse Kinematics Rollout (page 2–497)

Binding Objects (HD Solver)

Use the *binding functions (page 2–492)* in the HD Solver to cause an object to attempt to maintain its position and orientation relative to the world or to another object. A bound object will resist motion in the kinematic chain until the chain is stretched near its limits and the bound object must move to complete the IK solution.



In the top row, the ball is independent of the lamp. In the bottom row, the lamp head is bound to the ball's position.

Setting Binding Options

You can bind an object by position, orientation, or both. Once an object is bound you can set binding axes, weight, and whether the binding is absolute or relative.

Bind Position—Causes an object to attempt to maintain its location or to match the location of a follow object.

Bind Orientation—Causes an object to attempt to maintain its rotational orientation or to match the orientation of a follow object.

Relative—Changes how an object reacts to a bound follow object. See the following discussion about binding an object to a follow object.

Axis and Weight—Set these options to control which axes are affected by the binding and the binding's influence over other bindings in the kinematic chain.

Binding an Object to the World

You can bind an object to the world if you want the object to hold its position and orientation as long as possible during IK operations.

If you are animating a walking figure you want one foot to remain in place while you position the other foot. Without binding, the whole hierarchy would have a tendency to slide around when you attempt to position a foot. Binding both feet to the world ensures that the unselected foot stays in place while you position the other foot.

Binding an Object to a Follow Object

You can bind a selected object in your hierarchy to any other object that is not a descendant of the selected object. This other object is called the follow object.

The behavior of the bound object varies depending on the state of the Relative buttons next to the bind options.

- When Relative is inactive, the bound object tries to match the exact position and orientation of its pivot point to the position and orientation of the follow object's pivot point.
- When Relative is active, the bound object mimics any changes in position or rotation of the follow object but does not try to match it exactly.

You might want to animate a figure that always points to another object. Bind the hand of the figure to the other object with Relative active. Turn IK on, and as you move the object, the hand and arm of the figure move to point at it.

Unbinding Objects

If you decide that you do not want an object to be bound you can either turn off the Bind Orientation and Bind Position boxes, or unbind the object from its follow object.

- If you turn off the binding options, you disable the effect of binding without removing the binding to a follow object. This method is handy if you decide to reposition a follow object after you have bound an object to it.

- If you unbind an object, it permanently removes the binding to a follow object.

Controlling IK Precision (HD Solver)

Even with carefully defined joint precedence and joint parameters there are many valid IK solutions for any placement of an end effector. You set position and rotation thresholds to control IK precision and solution speed.

- The threshold and iteration settings are found on the Motion panel on the IK Controller Parameters Rollout, when the end effector is selected.
- The threshold and iteration settings for the HD solver are also found on the Inverse Kinematics tab of the Customize > Preferences dialog.

Position Threshold—Sets how close the end effector has to get to the follow object or cursor position to be considered a valid solution. The value expresses a distance in the current unit system. Small values increase accuracy but take longer to solve.

Rotation Threshold—Sets how accurately the end effector has to match the orientation of the follow object to be considered a valid solution. The value expresses a rotation angle in degrees. Small values increase accuracy but take longer to solve.

Iterations—Sets the maximum number of times 3ds Max repeats the IK calculations to find a valid solution. A high Iterations value increases the chance that the software can calculate a valid solution, but takes longer to complete. The following rules apply:

- Calculations stop as soon as a valid solution is found, even if the maximum number of iterations have not been performed.
- The last solution calculated is used if the maximum number of iterations is reached, regardless of whether the solution satisfies the position and rotation thresholds.

Setting Applied and Interactive Thresholds

You set the threshold and iterations settings as a trade-off between accuracy and speed. You have two groups of settings so you can individually tune the behavior of applied IK and interactive IK.

- Use interactive IK if you want fast, real-time response. Your interactive IK preferences should be set for speed.
- Use applied IK if you want the solution to match the follow objects very closely. Your Applied IK preferences should be set for accuracy.

The default applied IK and interactive IK settings are shown in the following table:

Thresholds	Applied IK	Interactive IK
Position	0.100	1.000
Rotation	0.100	1.000
Iterations	50	10

See also

HD Solver Motion Panel Rollouts (page 2–463)

HD Solver Motion Panel Rollouts

Select end effector of an HD chain > Motion panel > Parameters button > IK Controller Parameters rollout

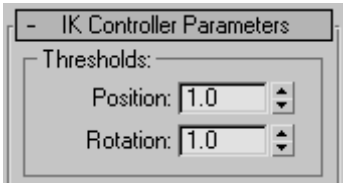
To adjust the parameters for all bones or hierarchically linked objects in the chain, you select a single bone or object, and then adjust the parameters in Motion panel > Parameters to affect all the bones or objects in the hierarchy.

Procedures

See *History-Dependent (HD) IK Solver (page 2–461)*.

Interface

Thresholds group

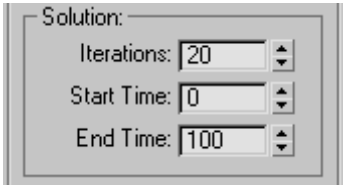


These settings are the same as those in Customize menu > Preferences > Inverse Kinematics. They set the threshold for the changes in position and rotation between the end effector and the last link of the chain.

Position—Specifies, in units, the "slop" factor between the end effector and its associate object. This is the allowable distance between the end effector and its object, given the current IK solution.

Rotation—Specifies the allowable number of degrees of rotational error between a rotational end effector and its associate object.

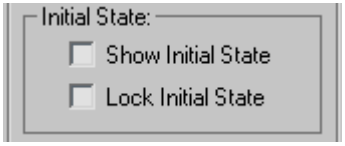
Solution group



Iterations—Specifies the maximum number of iterations allowed to solve the IK solution. The actual number of iterations used depends on the difficulty of the problem. This value simply puts a cap on it.

Start Time and End Time—Species the range of frames in which IK is solved. If the Start Time is set to 20 and the End Time is set to 80 in a 100-frame animation, and then the end effector is animated, the animation of the chain will only be correctly solved within the Start and End range.

Initial State group



When you first create a hierarchical chain, the initial position of the bones or objects is the *initial state*. The options here let you display, alter, or lock the initial state. When both of the options here are turned off (the default), transforming the end effector is different from transforming the other joints in the chain. When you select and transform a joint assigned an end effector, you are actually transforming the end effector, and an IK solution transforms the joints. This alters the position of the bones or objects in the viewport, but will not alter their initial state. On the other hand, when you select and transform any of the joints without end effectors, you alter their initial state.

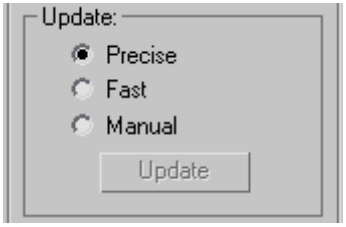
Show Initial State—Shuts off the real-time IK solution. All objects in the chain are moved to their initial positions and orientations prior to any changes caused by IK calculations. The positions of the end effectors, however, are not affected.

Turn on Show Initial State to view the initial state, or adjust the initial state of bones or objects using end effectors. This is useful when you set joint limits using the IK controls in the Hierarchy panel: the joints jump to their limits while you set them. When Show Initial State is off (the default), you can select and alter the initial state of bones or objects without end effectors, but when you select a bone or object with an end effector, the end effector is transformed, leading to an IK solution; the initial state of the associate object is not changed. By selecting this, you can transform the initial state of the associated object.

Lock Initial State—Locks all bones or objects in the chain from direct transformations. You can, however, transform the end effector of any joint

to generate the IK solution without altering the initial state. Activate this when you want to avoid accidentally selecting and transforming joints that are not assigned end effectors. When Show Initial State is selected, Lock Initial State is ignored.

Update group



The options in this group determine how the IK calculation is solved during the transformation of the end effector.

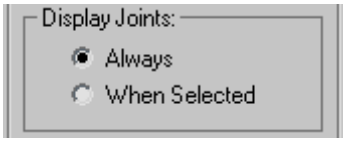
Precise—Solves the entire chain precisely for all frames from the Start time to the current time. When you release the mouse button, the chain will not move. Since all frames must be solved from Start to current time, the solution takes longer if the current frame is 1000 rather than 100.

Fast—Solves the chain only for the current frame while the mouse is moved. Releasing the mouse button solves for all frames. Although this is faster than Precise, sometimes the current-frame solution differs slightly from the all-frame solution. As a result, when you release the mouse, the objects in the chain might shift a bit.

Manual—With this option, the IK problem is not solved until you click Update.

Update—With Manual on, click to solve the IK solution.

Display Joints group



The options here affect the display of the joint axes and limits. When a rotational or sliding joint is active (Active is selected in the Sliding or Rotational Joints rollouts in Hierarchy > IK panels), an orange "rod" appears representing the active axis. When the joint is also limited, a pair of small orange squares appears, representing the From and To limits of the joint. If it's a sliding joint, the squares are positioned on the axis rod. If it's a rotational joint, the squares are at the ends of an orange arc.

Always—Displays the axis rod and joint limits for all joints in the chain at all times.

When Selected—Displays the axis rod and joint limits only on selected joints.

Tip: The joint icons can be difficult to see when the bones are linked to mesh objects. When animating a bone-based hierarchy, you can hide all of the objects and display and animate only the bones, which makes the joint icons more visible. This is easily done by hiding geometry by category in the Display branch.

End Effectors group



With the options in this group, you can add or delete Position or Rotation end effectors for any joint. Options affect only the currently selected joint. An *end effector* is the object that the IK chain follows to come to a solution. There are two types of end effectors: Position and Rotation. Both are

displayed as three intersecting blue lines at the joint. When you select a joint that carries either or both end effectors and perform a transformation matching the type of end effector, only the end effector itself is transformed. The objects in the chain then use IK to follow and create the IK solution.

When you transform a joint that does not have an end effector (or perform a type of transform that doesn't match the assigned end effector), you transform the joint itself and alter its initial state.

Position—Creates or deletes Position end effectors. If the node already has an end effector, only the Delete button is available.

Create—Creates a Position end effector for the selected node.

Delete—Removes the Position end effector from the selected node.

Rotation—Works like the Position buttons, except Rotation end effectors are created rather than Position end effectors. When you want to solve for IK rotation, add a Rotation end effector, and then use the Rotate tool to rotate the end effector.

End Effector Parent—Links an end effector to another object in the scene. Since end effectors are not, in themselves, objects, you won't see this linkage in Track View, but the linked end effector will inherit the transformations of its parent.

Note: Apart from the root object, you can't link an end effector to an object in the hierarchy because it would create an endless loop. The position of all objects in the hierarchy depends on the position of the end effector, so the end effector can't depend on the position of an object in the hierarchy.

End Effector (Display)—Displays the name of the selected parent object.

Link—Makes a selected object a parent of the currently selected link.

Unlink—Unlinks the currently selected end effector from its parent.

Remove IK group



Delete Joint—Deletes any selection of bones or hierarchical objects. Doesn't delete the IK; only the selected hierarchical members.

Note: Using the DELETE key will delete the entire hierarchy.

Remove IK Chain—Deletes the IK Solver from the hierarchy. Doesn't delete the objects or bones, only the IK controls.

Position—Displays the Key Info parameter rollouts specific for a Position end effector. This option is not active if no Position end effector is assigned. (The Key Info parameters are the same as those used for any Position controller.)

Rotation—Displays the parameters for an assigned Rotation end effector. This option is not active if no Rotation effector is assigned. (The Key Info parameters are the same as those used for any Rotation controller.)

HD Joint Settings

You control a joint's resistance to motion, and its tendency to return to its original position, by setting Ease, Damping, and Spring Back options.



Moving telescopes with and without damping

Easing a Joint at Its Limits

An organic joint, or a worn mechanical joint, moves freely in the middle of its range of motion but moves less freely at the extremes of its range. Use Ease to cause a joint to resist motion as it approaches its From and To limits.

For example, your forearm might move freely in the middle of its range of motion, but it resists movement when you try to squeeze it against your upper arm or extend it all the way out. Ease simulates this effect.

Damping Joint Action

As a joint corrodes, dries out, or is put under a heavy load, it resists motion along its active axes. Damping simulates the natural effect of joint friction or inertia. Enter a value greater than zero in the Damping field to apply resistance over a joint's full range of motion.

As damping increases a joint resists motion and other joints are required to move more. A damping value of 1.0 means there is extreme resistance and a joint will not move on that axis.

For example, a telescope with no damping at all allows each cylinder to move to its maximum limit before the next cylinder moves. If the cylinders have damping values assigned, then each cylinder causes its parent to begin moving before it reaches full extension.

Setting a Joint to Spring Back

When a joint resists motion, it also has a tendency to return toward its at-rest position. You simulate this by setting Spring Back tension in the joints. As the joint moves further from its rest position, an increasingly larger force pulls the joint back, like a spring.

When you set Spring Tension higher, the spring pulls harder as the joint moves farther away from its rest position. Very high settings can turn the joint into a limit, because you can reach the point where the spring is too strong to allow the joint to move any farther.

See also

Setting Joint Precedence (HD Solver) (page 2–467)

Using Default Joint Precedence (HD Solver) (page 2–468)

Choosing Child-to-Parent Precedence (HD Solver) (page 2–469)

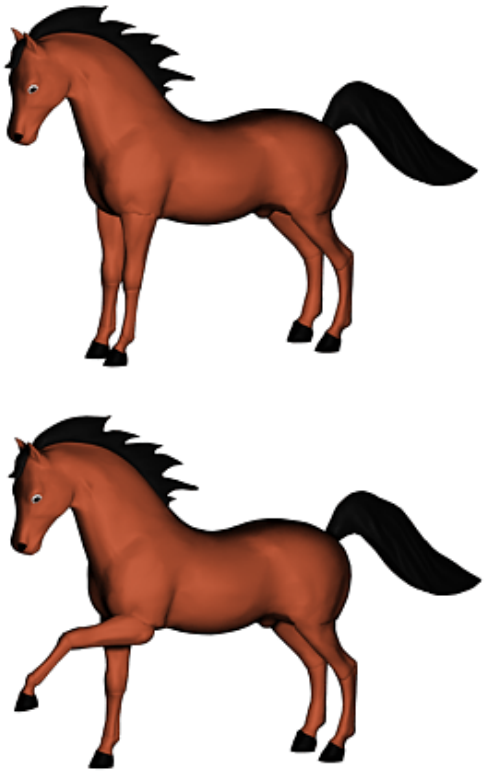
Choosing Parent-to-Child Precedence (HD Solver) (page 2–470)

Setting Precedence Manually (HD Solver) (page 2–471)

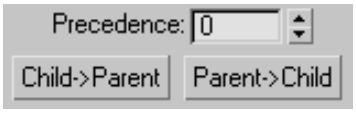
Terminating Chains (HD Solver) (page 2–471)

Setting Joint Precedence (HD Solver)

You set Joint Precedence to control the order in which joint calculations are applied to the kinematic chain. When an IK solution is calculated, the result is dependent on the order of calculation for each joint.



The hoof (a child) has precedence over the leg (its parent).



For any given position of an end effector, there are many possible IK solutions.

The three joint precedence controls in the Object Parameters rollout are:

Precedence—Sets joint precedence manually.

Child->Parent—Automatically sets joint precedence to decrease in value from child to parent.

Parent->Child—Automatically sets joint precedence to decrease in value from parent to child.

Determining Order of Calculations

You control the HD IK solution by setting joint precedence to determine which joints contribute the most to the HD IK solution and which joints contribute the least.

- Joints with high precedence values are calculated first, and contribute more motion to the HD IK solution.
- Joints with low precedence values are calculated last and contribute the least motion to the HD IK solution.
- Joints with equal precedence values are calculated by order in the hierarchy. Joints closer to the end effector are calculated first and joints closer to the root are calculated last.

Of all of the joint parameters, Joint Precedence is the most subtle.

See also

Using Default Joint Precedence (HD Solver) (page 2–468)

Choosing Child-to-Parent Precedence (HD Solver) (page 2–469)

Choosing Parent-to-Child Precedence (HD Solver) (page 2–470)

Setting Precedence Manually (HD Solver) (page 2–471)

Using Default Joint Precedence (HD Solver)

The default joint precedence is suitable for many HD IK solutions. It assumes that joints closest to where a force is applied (the end effector) will move more than joints farther from the force.

The default joint precedence is a value of 0 for every joint: *all joints have the same precedence value*. This default starts calculations with the end

effector and progresses up the IK chain until the base object is reached.

Using methods described in the following topics you can change the default precedence values. You can also set your kinematic chain back to its default precedence.

To reset the default joint precedence:

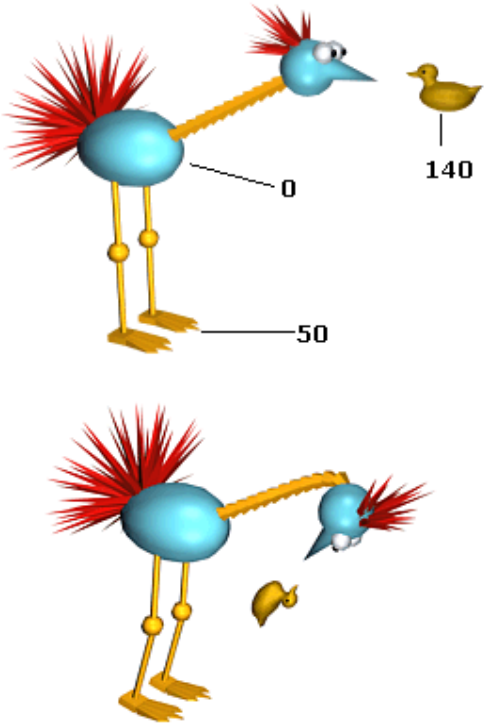
1. Select all objects in the kinematic chain.
2. On the Hierarchy panel, expand the Object Parameters rollout.
3. Set Precedence to **0**.

   Note: Default joint precedence occurs whenever all joints in the kinematic chain have the same precedence value. Assigning a value of 100 to all objects in the kinematic chain is exactly the same as assigning a value of 0.

Choosing Child-to-Parent Precedence (HD Solver)

Child–>Parent precedence causes joints closest to where a force is applied (the end effector) to move more than joints that are farther away from the force. Like the default precedence, Child->Parent precedence produces a natural result that is suitable for many IK solutions.

The difference is that Child–>Parent precedence assigns unique values to each object in the kinematic chain (default precedence uses the value of 0 for every object). Child–>Parent precedence is more flexible if you want to go back and manually change precedence values.



The first figure shows values for child-to-parent precedence.

Assigning Child->Parent Precedence

Clicking Child–>Parent sets joint precedence based on a child having a higher precedence than its parent. The values are calculated by setting the root of the entire hierarchy to a precedence value of 0 and each child to a value equal to 10 times its depth from the root.

You almost always assign Child->Parent precedence to an entire kinematic chain. The chain for the structure in the figure uses the body as the root object and the duck as the end effector.

Each object receives a value equal to its depth from the root of the hierarchy times 10. Note that precedence value calculation is based on the root of the hierarchy whether or not the root is selected as part of the kinematic chain.

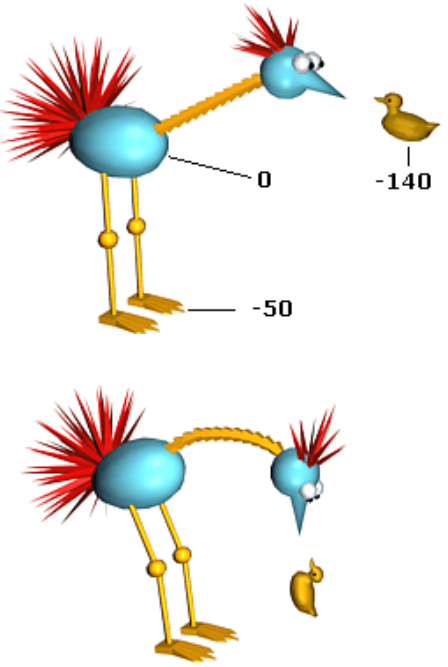
You might want to assign Child->Parent precedence to just a single object in the kinematic chain. For example, if you have been manually changing precedence values, you might want to set an object to its original Child->Parent value. The value assigned to the object is equal to its depth from the root of the hierarchy times 10.

Choosing Parent-to-Child Precedence (HD Solver)

Parent->Child precedence causes joints closest to where a force is applied (the end effector) to move less than joints farther away from the force. This is the opposite of Child->Parent precedence. Parent->Child precedence assigns the highest precedence to the base object and the lowest precedence to the end effector.

You might use this type of precedence for a number of special cases. For example:

- When joints near the end effector are less flexible than joints near the base. You use Parent->Child precedence in conjunction with limits and damping. For more information about the latter two, see *Setting Joint Resistance and Spring Back (page 2–466)*.
- When animating a kinematic chain where the real-world object is moved by applying force to the base of the chain. Imagine an animal's tail. Using IK you animate the tail by moving the tip of the tail. However, a real animal moves its tail using muscles at the tail's base. The real-world motion is applied from the base to the tip in a Parent->Child order.



The figure shows values for parent-to-child precedence.

Assigning Parent->Child Precedence

Clicking Parent->Child sets joint precedence based on a child having a lower precedence than its parent. The values are calculated by setting the root of the entire hierarchy to a precedence of 0 and each child to a value equal to -10 times its depth from the root.

You almost always assign Parent->Child precedence to an entire kinematic chain. The chain for the structure in the above figure uses the body as the root object and the duck as the end effector. Using Parent->Child precedence, the feet are assigned a precedence of -50 and the duck is assigned a precedence of -140.

You might want to assign Parent->Child precedence to a single object in the kinematic chain. The value assigned to the object is equal to its depth from the root of the hierarchy times -10.

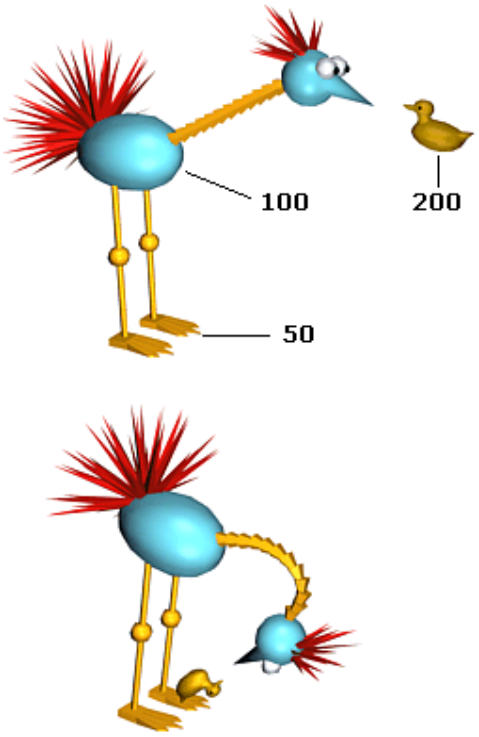
Setting Precedence Manually (HD Solver)

Some models and animated motions don't fit neatly into a Child->Parent or Parent->Child precedence. In such situations you can manually assign precedence values to any object in the IK chain on a joint-by-joint basis. For example:

- Animating models with a combination of light, flexible joints and heavy, resisting joints. Imagine a model of heavy iron balls linked together with lengths of chain. Setting the precedence values of the chains higher than the precedence values of the iron balls simulates the balls' inertial resistance to motion.
- Animating a motion where certain joints must move before other joints. Imagine a golfer's arm where the elbow should remain locked while swinging a golf club. You could accomplish this by setting the precedence of the elbow lower than the precedence of the wrist and shoulder.

High precedence values are calculated before low precedence values. Precedence values that are equal are calculated in Child->Parent order.

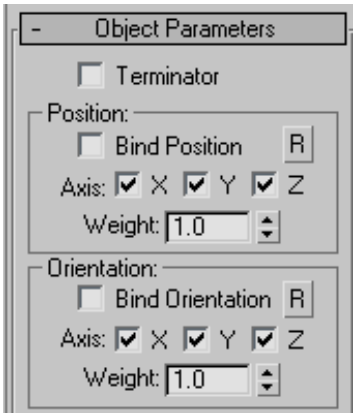
The precedence calculations only consider the relative IK values. This means that an IK chain of three objects with precedence values of 0, 30, and 200 would have the same solution if the precedence were changed to 1, 2, and 3.



The figures show precedence values that were assigned manually: the chain for the structure uses the body as the root object and the duck as the end effector.

Terminating Chains (HD Solver)

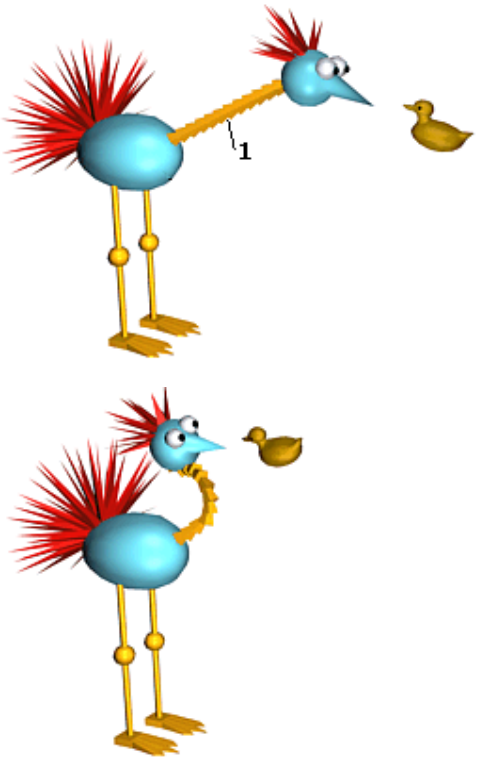
The kinematic chain is automatically defined using the object you select as the end effector and working up the hierarchical tree to use the root object as the base of the kinematic chain. Sometimes you might not want the kinematic chain to go all the way to the root of the hierarchy. This is especially true when you are animating multi-limbed or branching structures, such as an octopus or a tree.



Terminator option in the Object Parameters rollout

You use the Terminator option in the Object Parameters rollout to stop calculation of the kinematic chain before it reaches the root object of the hierarchy. A terminator object stops calculation at the terminator's child object; the terminator itself is not affected by the IK solution. This gives you very precise control over the behavior of the kinematic chain.

For example, look at the bird in the figure. The bird is linked with its body as the root of the hierarchy. If you use inverse kinematics to move the duck, all objects from the duck to the body will be affected by the IK solution. However, if you define a neck bone as a terminator object, then only the objects from the duck to the neck vertebrae are affected.



With a terminator in the neck, the body of the bird is not affected.

IK Limb Solver

The IK Limb solver is specifically meant for animating the limbs of human characters; for example, the hip to the ankle, or the shoulder to the wrist. Each IK Limb solver affects only two bones in a chain, but multiple solvers can be applied to different parts of the same chain. It is an analytical solver that is very fast and accurate in viewports.

To use the IK Limb solver, a bones system must have at least three bones in the chain. The goal is placed at the pivot point of the bone that is two bones away from the first selected bone.

The IK Limb solver works not only with bone hierarchies, but with any linked hierarchy that has at least three elements, and is set up to model a human limb. The additional requirements are:

- The first joint is "spherical." That is, it has three degrees of freedom.
- The second joint is "revolute," a robotics term that means it is based on a pin and has one degree of freedom.

The IK Limb solver uses the same controls as the HI Solver, so it allows for mixing periods of forward and inverse kinematics in the same animation period. It does not use the HD Solver methods of damping, precedence, and setting joint limits, instead it has a preferred angle parameter, swivel plane and IK/FK Enabling.

The IK Limb solver can be exported directly to a game engine.

Procedures

To apply a IK limb solver:

1. Create a chain that has three bones.
2. Select the root of the chain.
3. From the Animation menu, choose IK Solvers > IK Limb Solver.
4. In the viewport, move the mouse. You will see a dotted line attached to the cursor as you move the mouse.
5. Click the third bone in the chain, or any bone after the third bone.

   The IK Limb solver displays on the bone chain. The IK Limb solver only affects two bones in the chain, but you need to select the third bone to put it on the other two bones.

To animate an IK limb solver chain:

1. Apply the IK Limb solver to a two bone chain, with joint limits as described above.
2. Select the goal.
3. Move or rotate the goal.

Interface

The Interface for the IK Limb solver is identical to that of the HI Solver.


Select the goal at the end of the chain, and open the Motion panel. Here you will see the controls for the IK Limb solver. (The fourth is available from the Hierarchy > IK panel when a bone is selected.)

See also

IK Solver Rollout (HI Solver) (page 2–453)

IK Solver Properties Rollout (HI Solver) (page 2–456)

IK Display Options Rollout (HI Solver) (page 2–458)

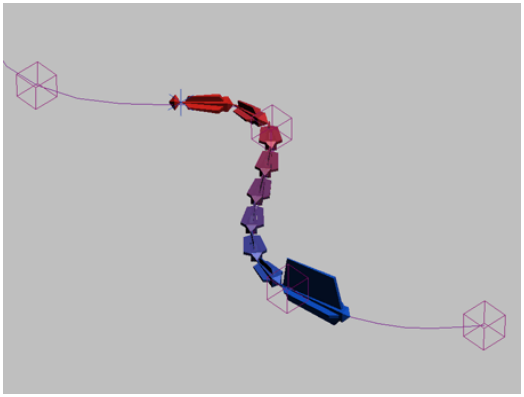
Sliding and Rotational Joints (HI Solver) (page 2–459)

Spline IK

Animation menu > IK Solvers > Spline IK Solver

Create panel > Systems > Bones > Choose SplineIKSolver from the IK Solver list.

The Spline IK solver uses a spline to determine the curvature of a series of bones or other linked objects.



You can move and animate the spline vertices to change the curvature of the spline. Usually, a helper is placed at each vertex to assist in animating the spline. The spline curvature is then passed on to the entire linked structure. The bones themselves do not change shape.

Normally the number of spline vertices and bones are the same, but you can use fewer vertices for easier posing and animating a long multiple-bone structure with just a few nodes, as opposed to animating each bone individually.

Spline IK provides a more flexible animation system than other IK solvers. You can position vertices/helpers anywhere in 3D space, so the linked structure can assume any shape you want to give it.

A helper object is automatically placed at each vertex when Spline IK is assigned. Each vertex is linked to its corresponding helper, so a vertex can be moved by moving the helper.

Unlike the HI Solver, the Spline IK system does not use a goal. The positions of helpers/vertices in 3D space is the only factor that determines the shape of the linked structure.

Note: Moving the Spline IK helpers to curl the spline can sometimes cause bones to rotate or flip unexpectedly. If this happens, you can try choosing a different upnode object with the IK chain's Pick Upnode option in the IK Solver Properties rollout (see *Spline IK Solver Rollouts (page 2–478)*), or use a different IK method for your hierarchy.

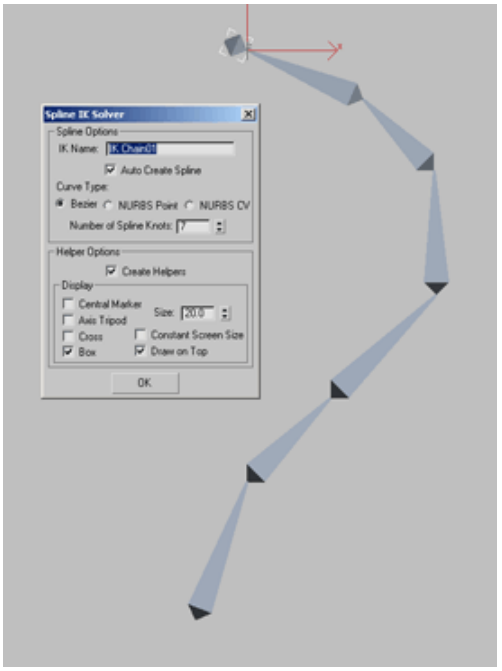
Applying a Spline IK Solver

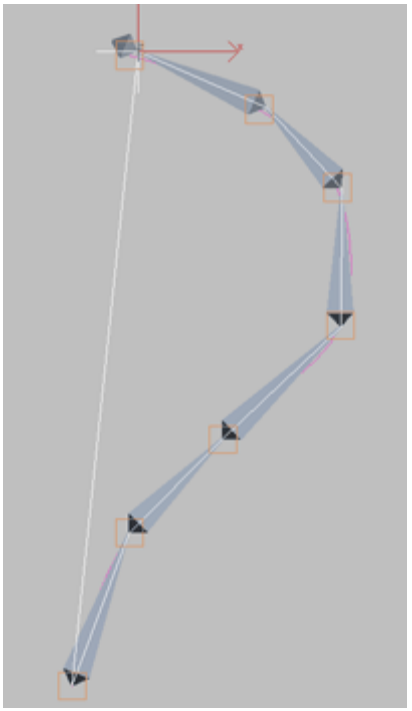
The Spline IK solver can be applied at the time bones are created, or after the bone structure has already been made.

To apply a spline IK solver when bones are created:

1. Choose Create panel > Systems > Bones.
2. On the IK Chain Assignment rollout, set IK Solver to SplineIKSolver, and turn on both check boxes: Assign To Children and Assign To Root.

   By default, Assign To Root is on automatically when you turn on Assign To Children.
3. Draw the bone structure as usual. When you right-click to end the bone creation process, the *Spline IK Solver dialog (page 2–477)* appears. Set parameters and click OK.

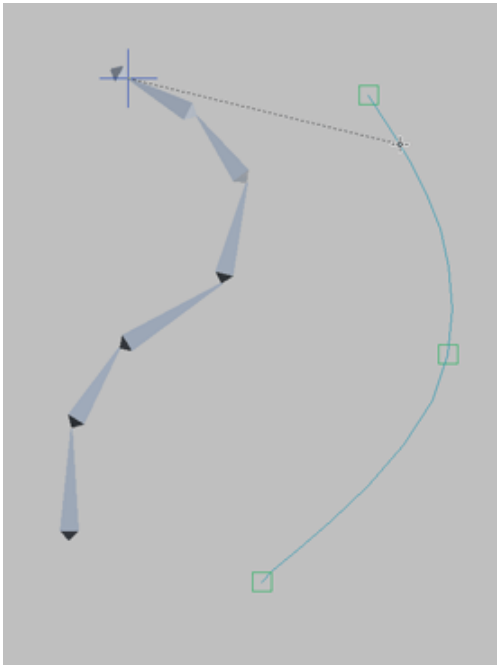




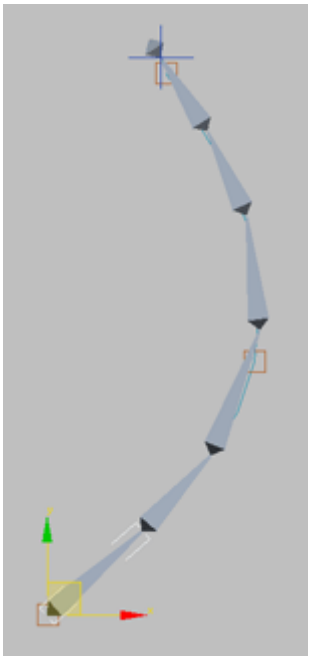
This procedure automatically draws the spline based on your dialog selections and sets up the Spline IK system to work with the spline. The software automatically assigns a Path constraint to the root bone to constrain it to the helper/vertex at that end of the spline.

To apply a Spline IK solver to an existing bone structure:

1. Create a bone structure without an IK chain.
2. Draw a spline or NURBS curve to be used by the bones. The curve can have any length and shape, and does not have to match the length or shape of the bone structure.
3. Select the bone or object where you want to the solver to start.
4. Choose Animation menu > IK Solvers > Spline IK Solver. In the viewport, move the cursor to the bone or object where you want the chain to end and click that bone. Then move the cursor to the spline and click it.



The bone structure jumps to the spline and takes its shape, and a helper is created on the spline at each vertex. A Path constraint is automatically assigned to the root bone to constrain it to the helper/vertex at that end of the spline.



To specify a spline after the spline IK solver is applied:

You can also apply a Spline IK solver to an existing bone structure without selecting a spline, then choose the spline later.

1. Create a bone structure without an IK chain.
2. Select the bone or object where you want to the solver to start.
3. Choose Animation menu > IK Solvers > Spline IK Solver. In the viewport, move the cursor to the bone or object where you want the chain to end and click that bone. Then right-click in the viewport to end the IK solver creation without choosing a spline.
4. Create a spline or NURBS curve to be used with the Spline IK system. It is recommended that you apply the *Spline IK Control modifier (page 1–830)* to it and then click Create Helpers to place helpers at each vertex, making the spline easier to control.
5. Select the IK chain cross hairs and access the Motion panel. On the Spline IK Solver rollout, click Pick Shape and click the spline.

   With this method, the position constraint is not automatically assigned to the root bone, so it must be assigned manually.
6. Select the root bone. Choose Animation menu > Constraints > Path Constraint, and then click the spline.

   This moves the bone structure to the spline, if they're apart, and creates a Position List controller for the bone with the Path Constraint as the second (active) constraint.

   Note: This also animates the bones along the path. If you don't want this, delete the second animation key.

Working with the Spline IK Solver

To use the Spline IK Solver, move the helpers to shape the spline. Do not move the spline. The shape of the spline is determined by the positions of the vertices, so moving the spline itself has no effect on the bone structure. If the spline is moved, it will snap back to the helpers the next time a helper is moved. However, it is recommended that you not move the spline at all.

Once you have finished setting up the spline IK solver, it can be helpful to freeze the spline to avoid accidentally moving it.

Moving helpers changes the shape of the spline. To twist the spline, select the spline IK chain object and go to the Motion panel to change twist angles in the IK Solver Properties rollout.

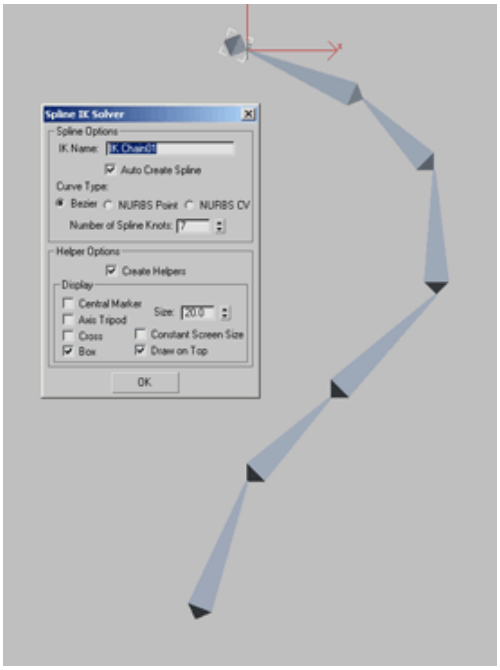
See also

Spline IK Solver Rollouts (page 2–478)

Spline IK Solver Dialog

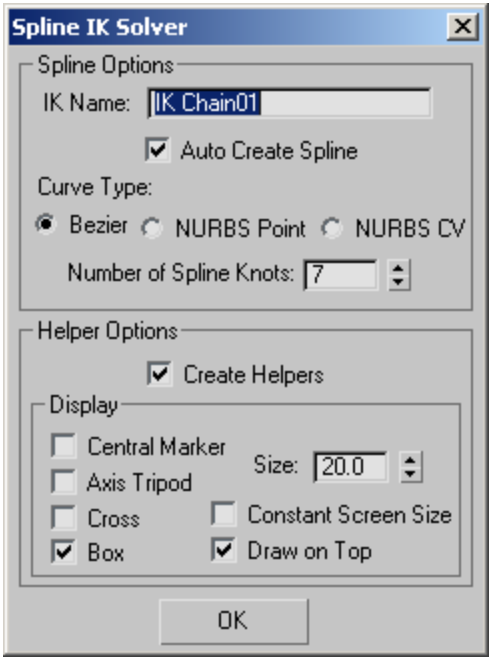
Create menu > Systems > Bones > Choose SplineIKSolver as IK Solver and turn on Assign To Children. > Create bone structure.

The Spline IK Solver dialog appears after bone creation if the Spline IK Solver has been selected as the IK solver and Assign To Children is turned on.



This dialog sets parameters for the spline that will be created to control the bone structure. The spline, in turn, is shaped by the positions of helper objects at each *knot* (control point).

Interface



Spline Options group

IK Name—Sets the name of the IK chain object. This object appears as a set of blue cross hairs at the end of the chain after you click OK.

Auto Create Spline—Automatically creates the spline used for the Spline IK solver based on the parameters on this dialog. When off, no spline is created and a spline must be manually created and assigned to the spline IK solver in order to use the solver. See *Animating with the Spline IK Solver (page 2–473)* for a description of this procedure. Default=on.

Curve Type—Choose the type of curve to serve as the Spline IK control:

- **Bezier—**Creates a Bezier spline, similar to a spline created with the Create > Shapes > Line tool. This is the default curve type.
- **NURBS Point—**Creates a NURBS Point Curve, where all knots (control nodes) lie on the curve.

- **NURBS CV**—Creates a NURBS CV Curve, where knots (control nodes) do not necessarily lie on the curve.

Number of Spline Knots—Vertices on the spline, or control points and CVs on the NURBS curve. You alter the spline or curve by moving and animating the knots. The number of knots sets the degree of refinement on the curve as well as the number of points that can be selected and moved to change the shape of the curve.

The default value is the number of bones in the Spline IK chain. Because one purpose of a Spline IK solver is to reduce the number of objects that require animation, you might want to use a lower value, such as half the number of bones.

Helper Options group

Create Helpers—When on, the software places one or more *Point helper (page 2–23)* objects at each knot in the spline or curve, making it easier to move and animate the knots.

You can enable more than one type of helper. When you move one helper, the others at the same location also move, as does the curve knot.

Center Marker—Places a small X-shaped Point helper at each knot.

Axis Tripod—Places a small tripod-axis-shaped Point helper at each knot.

Cross—Places a small cross-shaped Point helper at each knot.

Box—Places a small box-shaped Point helper at each knot.

Size—Sets the size for helpers.

Constant Screen Size—Keeps the sizes of helpers constant regardless of the zoom extent of the viewports.

Draw On Top—Displays the helpers on top of all other objects in the scene for improved visibility in busy scenes.

Tip: To change the display of helpers after creation, select each helper and change selections on the Modify panel.

See also

Spline IK (page 2–473)

Spline IK Solver Rollouts

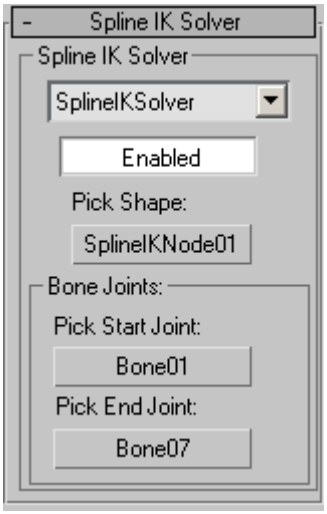
Select a Spline IK Solver control (the blue cross at the end of the chain) > Motion panel > Parameters button

When a Spline IK chain is selected, the options on the Motion panel can be used to change the starting and ending bones, and twist angles for the entire linked structure.

To create a Spline IK Solver, see *Animating with the Spline IK Solver (page 2–473)*.

Interface

Spline IK Solver rollout



The options on this rollout allow you to pick a new spline for the spline IK chain, or pick new start and end joints for the chain.

Spline IK Solver group

Spline IK Solver—Displays the name of the solver. The only available solver is Spline IK Solver.

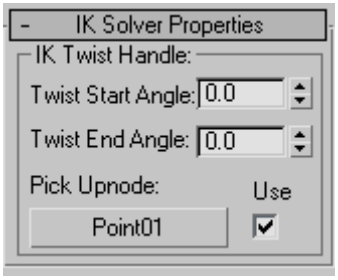
Enabled—Turns control of the solver on and off.

Bone Joints group

Pick Start Joint—Picks the start joint for the Spline IK Solver and displays the object name.

Pick End Joint—Picks the end joint for the Spline IK Solver and displays the object name.

IK Solver Properties rollout



The options on this rollout set the start and end twist angles for the spline IK chain. The twist manipulators on each end of the chain show the degree of twist.

Twist Start Angle—Sets the twist angle for the entire chain, rotating all bones in the chain.

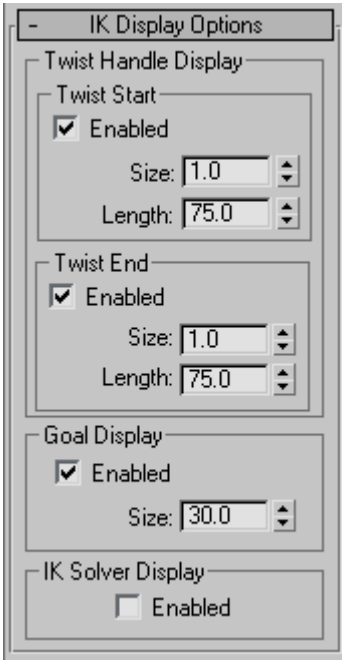
Twist End Angle—Sets the twist angle for the end joint relative to the start joint, creating a gradual twist over the entire chain.

Pick Upnode—When Spline IK is assigned to a bone chain, the IK system attempts to determine the "up" direction for bones, which is the direction in which the bones' back fins should point. The Spline IK Solver might interpret this direction incorrectly, causing bones to flip to one side or the other. Picking an upnode causes the "up" direction to point along the picked object's local Z axis. By default, the upnode is set to the first helper/knot on the spline, which usually causes the "up" direction to be interpreted correctly.

Use—Use the picked node to specify the "up" direction for the root bone.

IK Display Options rollout

The options on this rollout set the display of twist manipulators and the IK chain object.



Twist Start group

Enabled—Turns on the display of the start twist manipulator. Default=on.

Size—Sets the size of the start twist manipulator. Default=1.0.

Length—Sets the length of the start twist manipulator. Default=75.0

Twist End group

Enabled—Turns on the display of the end twist manipulator. Default=on.

Size—Sets the size of the end twist manipulator. Default=1.0.

Length—Sets the length of the end twist manipulator. Default=75.0.

Goal Display group

Enabled—Turns on the display of the IK goal. Default=on.

Size—Sets the size of the IK goal. Default=30.0.

IK Solver Display group

Enabled—Turns on the display of the IK chain object. Default=off.

See also

Animating with the Spline IK Solver (page 2–473)

Interactive and Applied IK

Animating with Interactive IK

Select the end of a hierarchy > Hierarchy panel > IK > Interactive IK.

With Interactive IK and the Auto Key button turned on, you position your model on keyframes, and the IK solution is interpolated between those keyframes. Because the IK solution accounts for multiple objects and the joints between objects, the interpolated animation of an IK object is usually different from the animation of objects without IK.

Tip: Interactive IK is a feature carried over from the earliest version of the software. It is recommended that you explore the IK Solver methods first, and only use Interactive IK if you find the IK Solvers will not satisfy your needs.

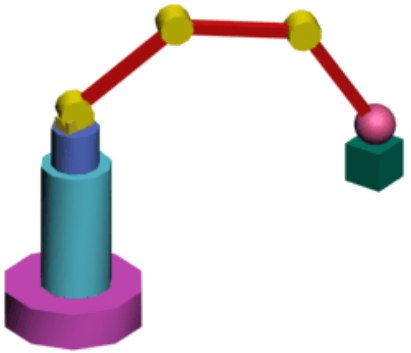
Example of Interactive IK

The following example demonstrates how an interactive IK animation works and how it might differ from what you expect. The figure shows an IK structure with its end effector resting on top of a box. The box moves in a straight line over 100 frames.

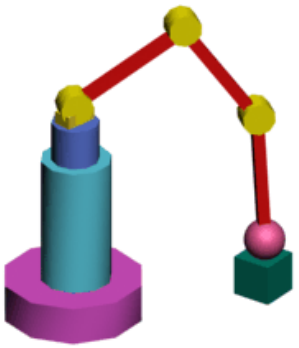
Turn the Auto Key button on and then in the Hierarchy panel > IK tab, turn on Interactive IK. Move the end effector of the IK structure to rest on top of the box at frame 100. The interpolated animation of the end effector follows a natural looking curved path.

You might have expected the IK bone structure to follow the same path as the box. However, the IK solution is only calculated at the keyframes. The positions and rotations of all the objects in the kinematic chain are interpolated between the keyframes to produce the curved result.

To make the end effector closely follow the path of the box, you would have to add more keyframes. You can also use binding and standard applied IK. For information about making the end effector precisely follow the motion of the box, see *Animating with Applied IK (page 2–481)*.



IK structure



Results of animating the box and IK structure

Behavior of Objects in IK Mode

As you move and rotate objects using interactive IK you notice that some objects might not be able to move or rotate about all axes. This is because the objects are constrained by the joint parameters you have set. If the joint parameters specify that motion cannot occur in a certain axis, the end of the chain will not move.

Behavior of Root Objects in IK Mode

An option on the *Inverse Kinematics panel (page 3–830)* of the Preference Settings dialog enables the special case of moving and rotating root objects when IK mode is on. The name of this control is Always Transform Children Of The World.

While you are transforming an end effector using IK mode, the settings for the root object's joint parameters are used and the root object does not move with respect to the World. If you select the root object, its joint parameters are released and you can transform the root object.

If you decide you want root objects to always use their joint parameters, you can turn off the Always Transform Children Of The World option.

Single, unlinked objects are hierarchies of one. An unlinked object is its own root and is also a child of the World. Turning Always Transform Children of the World will prevent you from transforming single objects in IK mode.

Animating with Applied IK

Select the end of a hierarchy > Hierarchy panel > IK > Set object parameters such as Bind to Follow Object > Apply IK.

Applied IK requires you bind one or more parts of your IK structure to animated follow objects. Once bound, you can select any object in your kinematic chain and click Apply IK. This calculates the IK solution for each frame of the animation and places transform keys for every object in the IK chain.

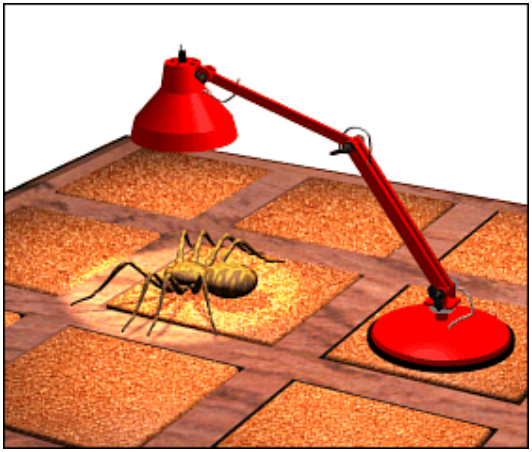
The Apply IK method of animation works best when you want objects in the kinematic chain to match the motions of other objects exactly.

Tip: Applied IK is a feature carried over from the earliest version of the software. It is recommended that you explore the IK Solver methods first, and only use Applied IK if you find the IK Solvers won't satisfy your needs.

Example of Applied IK



IK structure



Animating the lamp with applied IK

The illustration demonstrates how an applied IK animation works. Just as in the example in *Animating with Interactive IK (page 2–480)*, the figure shows an IK lamp structure with its head positioned over a spider. To maintain the offset distance between the lamp head and spider, the actual end effector is an invisible Dummy helper object, placed within the spider and linked to the lamp head as a child. The spider moves in a straight line over 100 frames.

To bind the end effector to the spider, click Bind, and then drag a line between the two. Once the end effector is bound to the spider, it will try to match the location of its pivot point with the pivot point of the spider.

Select the end effector and click Apply IK. The software matches the end effector with the box and calculates the IK solution for every frame. Playback of the animation shows that the end effector perfectly follows the box.

Note: In the above example, the end effector is an invisible Dummy helper object, which is used to maintain the offset distance between the lamp head and the spider.

Clearing Keys from Previous Animation

If you have animated any members of the IK chain interactively, or run Apply IK previously, the existing animation keys will affect the new IK solution. Sometimes that is exactly what you want. You can use manual animation to subtly nudge the IK solution toward a particular result. However, it is more likely that you want to remove old keys in order to begin the IK calculations with a clean slate.

The Clear Keys option in the Inverse Kinematics rollout controls whether or not old move and rotate animation keys are removed.

Constraining the IK Solution to Specific Frames

You use the Start and End fields in the Inverse Kinematics rollout to set the range of frames used to calculate the applied IK solution. Using these fields, you can restrict IK solutions to specific frames and solve for different solutions in different time segments.

You can set the Start and End fields to include frames outside of the active time segment.

You can also constrain an IK solution to frames where a key exists for an end effector. This is useful if you want to animate a hierarchy using end

effectors but do not want keys generated on every frame.

Turn on Apply Only To Keys to constrain the IK solution to frames with end effector keys.

Watching Progress of the IK Solution

Normally, Apply IK calculates all frames before updating the viewports. To watch the progress of Apply IK frame by frame, turn on Update Viewports on the Inverse Kinematics rollout.

Updating the viewports greatly slows down the Apply IK process but it can help you troubleshoot complex animations.

Joint Controls

Setting Joint Parameters

You set whether a joint behaves as a hinge, a drawer slide, or another type of joint by setting *joint parameters (page 2–459)* for each object in the kinematic chain. Joints control the rotation and position of an object with respect to its parent.

Any object has a maximum of two joint-type rollouts: One rollout contains settings to control the object's position, and the other controls the object's rotation. There can be many different types of positional and rotational joints. Which joint parameters are available is determined by the type of IK solver assigned to an object. HI Solvers, for instance, are controlled with a preferred angle setting found in the Rotational Joint parameters. HD Solvers have additional parameters for spring back, precedence, and damping, not found in the HI Solver.

Any hierarchy of object or bone systems can have its joint limits defined. Select all the objects, and turn on their bone or link display. Select the bone or link and open the Hierarchy panel > IK tab. Scroll down to the Sliding and Rotational Joints. There, you can activate axes and set their individual limits.

Note: Different IK solvers use different joint limits. When using a bone system, add the IK solvers first, then set the joint limits second.

Common Joint Types

The most common joint types are Rotational and Sliding joints. Other common joint types are Path and Surface joints. Each joint type displays its own set of joint parameters.

Rotational joints—Control rotation of objects using many of the standard rotation controllers. The parameters for rotational joints set an object's ability to rotate about a given axis.

For many IK structures, consider using an *Euler XYZ controller (page 2–318)*. Quaternion-based controllers tend to freeze if an object is moved beyond rotational joint limits before turning on IK.

Sliding joints—Control the position of objects using many of the standard position controllers. The parameters for sliding joints control whether an object can move along a given axis.

Surface joints—Control the position of objects using *Surface constraint (page 2–396)*. The parameters control how the object travels along its assigned surface.

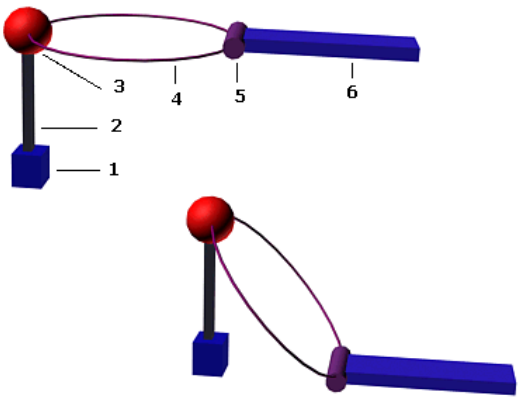
Path joints—Control the positional motion of objects using *Path constraint (page 2–398)*. The parameters for path joints control how far an object can travel along its assigned path.

Using Joint Parameters

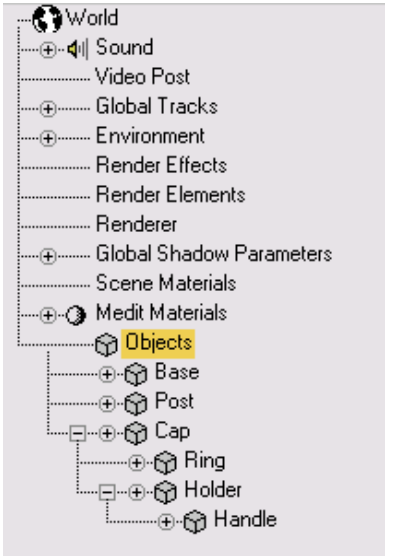
The way the children behave toward one another is governed by the inheritance of the transforms

down the chain. Setting *joint limits* on individual child objects can influence this inheritance. If three children in a row all have their rotational axes made inactive, they will not be able to rotate, and so a piece of the chain will appear stiff. Or if only one axis is made active as a sliding joint, the component can separate in space from the chain.

When you are using an object with a path constraint in an IK chain you might want the path to appear as if it is part of the IK chain. You achieve this effect by linking the object using the path constraint and the path to the same parent. The path object should have no children, and other objects in the IK chain should be linked to the object using the path constraint.

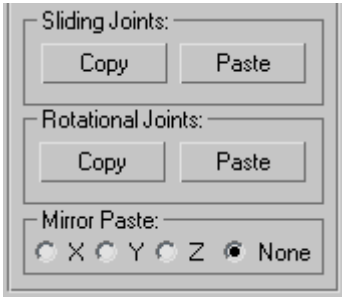


1. *Base*
2. *Post*
3. *Cap*
4. *Ring*
5. *Holder*
6. *Handle*



In the figure, the bead uses a path constraint to hold it to the ring. The IK chain from parent to child is Post->Cap->Bead->Handle. The ring is a child of the Cap but is not part of the IK chain.

Copying and Pasting Joint Parameters



You can copy and paste at the bottom of the Object Parameters rollout in the Hierarchy panel to copy and paste a complete set of joint parameters from one object to another.

There are separate copy and paste functions for sliding joints and rotational joints. Each joint type saves copied parameters in separate clipboards.

Tip: Choose one of the Mirror Paste options if you want to mirror the joint parameter settings as you

paste them. This is very useful when pasting from one side of an object to another, such as left arm joints to a right arm.

You can also copy joint settings from a non-IK controller to an IK controller, but you can't copy from an IK controller to a non-IK controller.

Activating Joint Axes

You set whether an object can move or rotate about a given axis by using the Active check box in the joint rollouts. Joints have a maximum of six possible axes: three for rotation and three for position. You constrain the motion of a joint by setting which axes are active.

- A joint with all axes active can move and rotate freely, independent of its parent.
- A joint with all axes inactive is locked to its parent and cannot move independently.
- The setting of IK joint axes overrides Inherit and Lock settings on the Link Info rollout.

Understanding Joint Axis Orientation

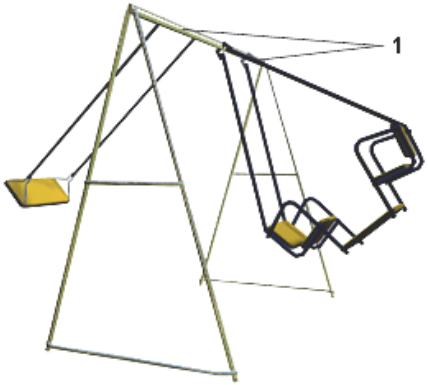
IK joint axes for an object are defined by the Local axes of the object's parent. That means if you activate the X axis of an object's rotational joint parameters, the object rotates about its parent's X axis, not its own X axis.

You might have a problem when you set joint parameters for an object whose local coordinate system is oriented 90 degrees from its parent's coordinate system. In such a case, the rotation angle about one axis becomes indeterminate. The result is that one axis will often cause rotation about one of the other two axes.

The following techniques will make setting joint parameters easier and more successful:

- Assemble IK hierarchies with objects aligned along the World axes.
- Consider using the *Adjust Pivot (page 2–488)* functions to align your object pivots with either the World axes or the root object of the hierarchy.
- Set the transform managers to Parent coordinate system and Use Pivot Point Center, while setting joint parameters. This helps you see the orientation of the joint axes by displaying the parent axis icon at the selected object's pivot point.
- When activating joint axes, drag the From spinner up and down. This causes the object to move or rotate about the active axis and is a quick check that you chose the right axis.

Activating Rotational Joints

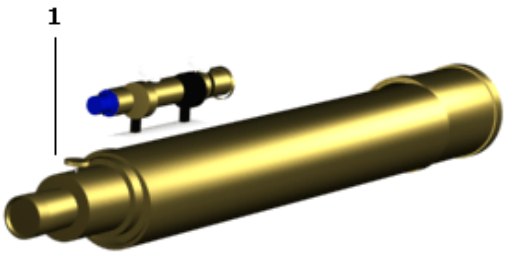


1. Rotational links

When you turn on Active for one of the X, Y, Z axes of a rotational joint, the object can rotate about that axis of its parent's coordinate system.

Joints that rotate about multiple axes are very common. A ball joint, like your shoulder, is a rotational joint active about all three axes. A pin joint, like your elbow, is a rotational joint active on a single axis.

Activating Sliding Joints



1. Sliding axis

When you turn on Active for one of the X, Y, Z axes of a sliding joint the object can move along that axis of its parent's coordinate system.

Most sliding joints are active only along a single axis. A telescope is an example of a sliding joint active on a single axis. You rarely see a sliding joint active along all three axes.

If a sliding joint is active along all three axes, it moves independently from its parent. It's almost as if there is no joint connection at all.

Tip: Use the HD IK Solver when you need to use sliding joints with IK.

Path and Surface Joints



Circular path

When you turn on Active for Path or Surface joints you are setting whether the object can move along the assigned path or surface.

A house key on a ring is an example of an active path joint.

See *Path Constraint (page 2–398)*.

Limiting Joint Action

Most joints are limited in their range of motion along an active axis. For example: a hinge might open only to 120 degrees; a piston slides only within the length of its cylinder. To limit the range of motion allowed on an active axis, you turn on Limited and set values for the From and To fields on the joint rollouts.

You can also view joint limits in the viewport by pressing and holding the mouse button on the From or To label of a limit field. The object will move or rotate to the limit value until you release the mouse button.

Limiting Rotational Joints

Limits for a rotational joint define how far the object can rotate about its parent's axes. The values in the From and To fields represent the rotation angle about the active axis measured from 0 degrees on the parent object.

For example, an elbow joint rotates the forearm with respect to the upper arm. In the figure the limits on X axis rotation are from 0 to 135 degrees. The Y, Z axes are inactive because an elbow joint rotates about a single axis.

Limiting Sliding Joints

Limits for a sliding joint define how far a joint can move along its parent's axes. The values in the From and To fields represent a distance for movement along the active axis measured from

the pivot point of the parent to the pivot point of the selected object.

For example, a sliding joint on a piston moves the piston in and out of the cylinder. In the figure, the limits on Z axis movement are from 10 to 90. This prevents the piston from hitting the bottom or moving past the end of the cylinder.

Limiting Path and Surface Joints

Limits for path and surface joints define how far along the path or surface an object can move. The values in the From and To fields represent a percentage of the total distance measured along the path or surface.

For example, a path joint for a house key moves the key along a key ring. Setting the limits on the path joint at 5% to 95% prevents the key from traveling along the ring where the fob is attached.

See also

Sliding and Rotational Joints (HI Solver) (page 2–459)

Hierarchy Panel Commands

Once you have set up a hierarchy using the *Select and Link command (page 2–422)* or a system such as *Bones (page 1–404)*, you can manage it using the Hierarchy panel. This command panel has three tabs:

Pivot (page 2–487)

IK (page 2–491)

Link Info (page 2–499)

You use the Pivot tab to adjust the pivot points of objects in the hierarchy. You use the IK tab to manage the behavior of inverse kinematics (IK). You use the Link Info tab to apply locks or inheritance to movement within the hierarchy.

See also

Hierarchies and Kinematics (page 2–416)

Inverse Kinematics (IK) (page 2–435)

IK Terminology (page 2–437)

Animating with Interactive IK (page 2–480)

Animating with Applied IK (page 2–481)

History-Dependent (HD) IK Solver (page 2–461)

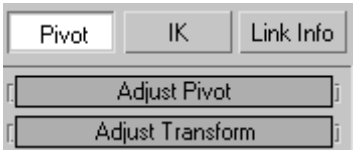
History-Independent (HI) IK Solver (page 2–446)

IK Limb Solver (page 2–472)

Spline IK (page 2–473)

Pivot

Make a selection in the viewport. > Hierarchy panel > Pivot button



All objects have a pivot point. You can think of the pivot point as representing an object's local center and local coordinate system. The pivot point of an object is used for the following:

- Functions as the center for rotation and scaling when you select the Pivot Point transform center.
- Sets the default location of a modifier center.
- Defines the transform relationship for the object's linked children.

- Defines the joint location for *inverse kinematics (IK) (page 2–435)*.

You can adjust the position and orientation of an object's pivot point at any time using the buttons in the Adjust Pivot rollout in the Hierarchy panel. Adjusting an object's pivot has no effect on any children linked to that object.

Note: You cannot animate the functions under the Adjust Pivot rollout. Adjusting an object's pivot on any frame changes it for the entire animation. Use caution when adjusting the pivot of an animated object. Adjusting the pivot on one frame might produce unexpected results in another part of the animation. You should adjust the pivot point to the location you want, before you start to animate, if possible.

Tip: Use the *Link constraint (page 2–403)* to simulate the effect of an animated pivot.

The Pivot panel has two rollouts:

Adjust Pivot Rollout (page 2–488)

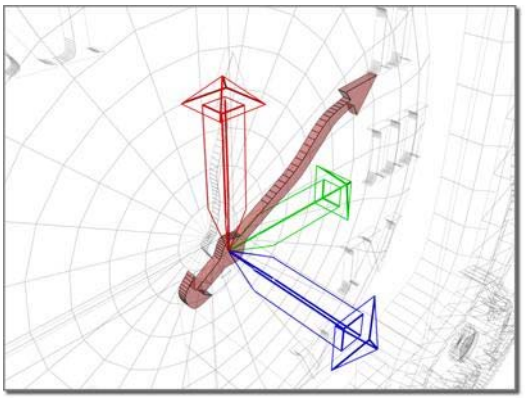
Adjust Transform Rollout (page 2–489)

Adjust Pivot Rollout

Make a selection in the viewport. > Hierarchy panel > Pivot button > Adjust Pivot rollout

You can adjust the position and orientation of an object's *pivot point (page 3–995)* at any time using the buttons in the Adjust Pivot rollout. Adjusting an object's pivot has no effect on any children linked to that object.

You cannot animate the functions under the Adjust Pivot rollout. Adjusting an object's pivot on any frame changes it for the entire animation.



Pivot point sets hand to the center of the clock face.

Procedure

To reposition an object's pivot point:

1. Select an object and then turn on Adjust Pivot rollout > Affect Pivot Only.
2. Move or rotate the pivot.

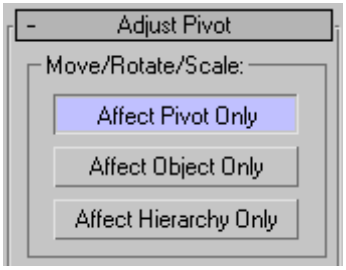
   You can also use *Align (page 1–462)*, *Normal Align (page 1–465)*, and *Align to View (page 1–468)* on the toolbar to align the pivot.

   See *Adjust Transform Rollout (page 2–489)* for related information.

Interface

Move/Rotate/Scale group

Each of the buttons in the Move/Rotate/Scale group box highlights when active. This determines which part of the object is affected by the three buttons in the Alignment area, as well as the Transform and Align commands on the Main toolbar.



Affect Pivot Only—Affects only the pivot point of the selected objects.

Note: A Scale transform has no effect on the pivot.

Affect Object Only—Affects only the selected objects (and not the pivot point).

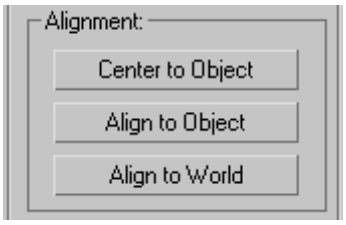
Affect Hierarchy Only—Affects only the Rotation and Scale tools. It applies the rotation or scale to the hierarchy by rotating or scaling the *position* of the pivot point without rotating or scaling the pivot point itself.

Note: You can use this on hierarchies of 3ds Max objects, but don't use it on Bones systems. To reposition, rotate or change the size of bones in a chain, see *Bones (page 1–404)*.

Note: It's important to remember that the Align, Normal Align, and Align to View functions are all affected by the state of Affect Pivot Only, Affect Object Only, and Affect Hierarchy Only. Snap mode allows you to snap the pivot to its own object, or to any other object in the scene.

Alignment group

The effect of these buttons depends on whether you chose Affect Pivot Only or Affect Object Only. They don't apply to Affect Hierarchy Only.



If you chose Affect Pivot Only, the buttons work as follows:

Center to Object—Moves the pivot to the center of its object.

Align to Object—Rotates the pivot to align with the object's transformation matrix axes.

Align to World—Rotates the pivot to align with the world coordinate axes.

If you chose Affect Object Only, the buttons work as follows:

Center to Pivot—Moves the center of the object to its pivot location.

Align to Pivot—Rotates the object to align its transformation matrix axes with the pivot.

Align to World—Rotates the object to align its transformation matrix axes with the world coordinate axes.

Pivot group



Reset Pivot—Resets the pivot point to the position and orientation it held when the object was first created. This is not affected by the state of the Affect Pivot Only and Affect Object Only buttons.

Adjust Transform Rollout

Make a selection. > Hierarchy panel > Pivot button > Adjust Transform rollout

You can transform an object and its pivot without affecting its children using the buttons in Adjust Transform rollout. Adjusting an object's transforms has no effect on any children linked to that object.

Procedure

To scale a parent without scaling the children:

1. Select the parent object in a hierarchy.
2. On the Hierarchy panel click the Pivot button if it's not already on, then on the Adjust Transform rollout, click Don't Affect Children.
3. Scale the parent object.

   The children will remain unaffected.

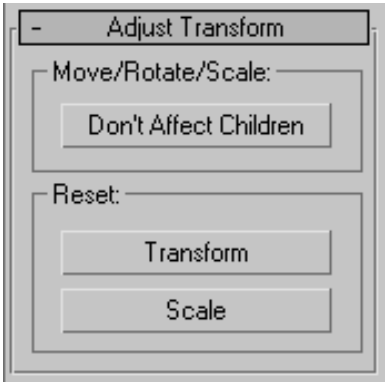
   Tip: Never use non-uniform scale at the object level for objects in a hierarchical chain. Always go to the sub-object level when you do a non-uniform scale on hierarchically linked objects.

To reinitialize the scale of an object:

1. Select an object.
2. In the Adjust Transform rollout > Reset group, click Scale.

   The XYZ values are re-initialized to 100%. This operation will prevent Non Uniform Scale inheritance if a child object is linked to this object.

Interface



Move/Rotate/Scale group

Don't Affect Children—Limits transforms to the selected object and its axis, not to its children.

This is very useful when working with bones and other hierarchies. If you need to make a particular bone or object longer or shorter, select the object and then turn on Don't Affect Children. You can readjust it and any parent objects in the hierarchy will stretch or shrink to compensate to the new dimension of the selected bone.

Reset group

Transform—Resets the orientation of an objects local axis coordinates to align with the world coordinate system, regardless of the current orientation of the object.

This does not affect descendants.

When applied to *closed* groups, it aligns the dummy node of the parent, but keeps the children in the same place. It reorients and resizes the dummy node to correctly bind the children of the group.

Scale—Resets the scale values in the transformation matrix to reflect the new scale of the object. There is no visual change in the object.

You can use the Scale to correct Non Uniform Scale inheritance problems. If Non Uniform Scale is inherited by a child object in a hierarchy, it can result in an undesirable deformation in the child object. To correct for this, use Reset: Scale in the Adjust Transform rollout on an object before linking.

You can view the scale values of the transformation matrix of a selected object. On the toolbar, turn on Scale, then right-click any of the three toolbar Transform buttons (Move, Rotate, or Scale).

There is also a *Reset XForm utility (page 1–438)* to do the same tasks.

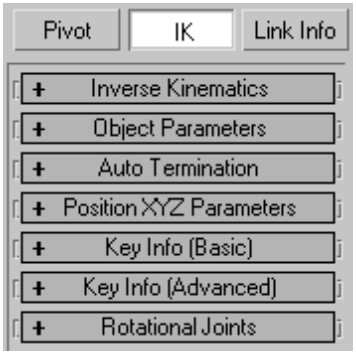
See also

Adjust Pivot Rollout (page 2–488)

IK

Make a selection. > Hierarchy panel > IK button

The IK rollouts contain controls for interactive IK and the HD IK solver.



The rollouts for interactive IK are:

Inverse Kinematics Rollout (page 2–497)

Auto Termination Rollout (Interactive IK) (page 2–499)

The rollouts for applied IK are:

Inverse Kinematics Rollout (page 2–497)

Object Parameters Rollout (HD Solver) (page 2–491)

The rollouts for the HD IK solver are:

Inverse Kinematics Rollout (page 2–497)

Object Parameters Rollout (HD Solver) (page 2–491)

Sliding and Rotational Joints Rollouts (HD Solver) (page 2–496)

Note: The rollouts for the HI IK and IK Limb solvers are not available while IK is chosen in the Hierarchy panel. Most controls for these solvers are found on the Motion panel.

HD IK Solver Rollouts

Object Parameters Rollout (HD Solver)

Make a selection. > Hierarchy panel > IK > Object Parameters rollout

The Object Parameters rollout lets you set IK parameters for an entire hierarchical chain. This topic covers the Terminator check box; other parameters on the Object Parameters rollout are covered in the following topics:

Position/Orientation/Bind to Follow Object (HD Solver) (page 2–492)

Precedence (HD Solver) (page 2–494)

Copying, Pasting, and Mirroring Joint Parameters (HD Solver) (page 2–495)

Sliding and Rotational Joints Rollouts (HD Solver) (page 2–496)

Procedure

To create a terminator in any hierarchy or HD IK chain:

1. Select an object in any hierarchy or HD IK chain that you want to be the terminator.
2. In the Object Parameters rollout, turn on Terminator.

   The terminated objects will not move when you use Interactive, Applied, or HD IK animation.

Interface



Terminator—Sets the base of an IK chain by defining one or more selected objects as terminators.

Turn on Terminator to stop calculation of the kinematic chain before it reaches the root object of the hierarchy. A terminator object stops calculation at the terminator's child object; the terminator itself is not affected by the IK solution. This gives you very precise control over the behavior of the kinematic chain.

Position/Orientation/Bind to Follow Object (HD Solver)

Make a selection with an HD Solver applied. > Hierarchy panel > IK > Object Parameters rollout > Position/Orientation/Bind to Follow Object groups

Use the controls in these group boxes to bind objects in a HD IK chain to the world or to follow objects. Use the various options to modify the behavior of the bindings.

Use weighted and axis-specific binding controls to create different weights for each binding, and to specify which axes the bound follow object influences. Use the R (Relative) buttons to establish a relative offset with Bind Position and Bind Orientation.

Binding an Object to the World

Binds an object to the world if you want the object to hold its position and orientation as long as possible during IK operations.

Binding an Object to a Follow Object

Binds a selected object in your IK chain to any other object that is not a descendant of the selected object. This other object is called the follow object.

Procedures

To bind an object to the world:

1. Select the object to be bound.
2. Click IK in the Hierarchy panel and expand the Object Parameters rollout.
3. Select one or both of the Bind boxes.
   - Turn on Bind Position to cause the object to attempt to maintain its current location.
   - Turn on Bind Orientation to cause the object to attempt to maintain its current orientation.

To bind an object to a follow object:

1. Turn on Bind and select the object in a kinematic chain to be bound.

   Moving your mouse over an acceptable target will change the cursor to a push-pin icon.
2. Drag to the follow object and release.
3. Turn on Bind Orientation if you want the bound object to match the rotation of the follow object. Bind Position is automatically selected.

   Turn off an axis (Axis: XYZ). The specified axis is no longer influenced by the follow object or the New IK Position end effector.

   Use Weight to set the relative influence of multiple follow objects or end effectors, and thus their priority in solving the IK solution. The higher the relative Weight value, the higher the priority.

   Note: As you drag, a dotted line is drawn from the pivot point of the selected object to the cursor. When the cursor is over a valid follow object, it changes to a push-pin cursor. Release to set the follow object. The name of the follow object will appear in the text area.

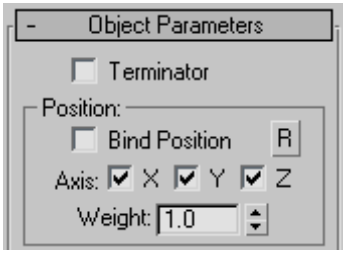
To unbind an object:

- Select the object to unbind, and then click Unbind.

  In the text area, the name of the follow object is replaced with the word "None."

Interface

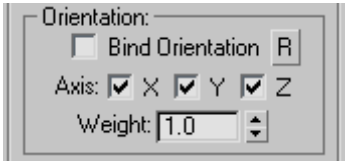
Position group



Bind Position—Binds the selected object in the IK chain to the world (attempts to maintain its location), or to a follow object, if one has been assigned. If a follow object has been assigned, then the translation of the follow object affects the IK solution.

Note: This check box has no effect on the HD IK Solver Position end effectors, which are always bound to their assigned joints.

Orientation group



Bind Orientation—Binds the selected object in the hierarchy to the world (attempts to maintain its orientation), or to a follow object, if one has been assigned. If a follow object has been assigned, then the *rotation* of the follow object affects the IK solution.

Note: This check box has no effect on the HD IK Solver Rotation end effectors, which are always bound to their assigned joints.

R—Establishes a relative position offset or rotation offset between follow object and end effector.

This button has no effect on the HD IK Solver Position end effectors. They are created on top of the joint to which they're assigned, and are automatically absolute.

If you move the joint away from the end effector, and want to reset the end effector to an absolute position, you can delete and then recreate the end effector.

Axis X/Y/Z—If one of the axes is turned off, the specified axis is no longer influenced by the follow object or the HD IK Solver Position end effector.

For example, if you turn off the X axis in the Position group, movement of the follow object (or end effector) along the X axis has no effect on the IK solution, but movement along the Y or Z axis has.

Note: These do *not* constrain the movement of the follow object or end effector.

Weight—Sets the influence of the follow object (or end effector) on its assigned object, and on the rest of the chain. A setting of 0 turns off the binding. Using this value you can set the relative influence of multiple follow objects or end effectors and their priority in solving the IK solution. The higher the relative Weight value, the higher the priority.

The Weight settings are relative; there's no point in using them if there's only one follow object or end effector in your IK hierarchy.

However, if you have a single HD IK chain with both Position and Rotation end effectors on a single joint, you can give them different weights to give priority to either the position or rotation solution.

You can adjust the Weight for multiple joints. Select two or more objects in the hierarchy, and the Weight value represents the commonality of the selection set.

Note: Differences in the Weight values only have an effect when there is no solution that satisfies two or more opposing end effectors in the chain. In

this case, the end effector with the greatest weight "wins."

Bind To Follow Object group



Controls to bind and unbind an object in an inverse kinematic chain to a follow object.

(Label)—Displays the name of the selected follow object. Displays the word "None" if no follow object is set.

Bind—Binds an object in an inverse kinematic chain to a follow object.

Unbind—Unbinds the selected object in a HD IK chain from its follow object.

Precedence (HD Solver)

Make a selection with an HD IK Solver applied. > Hierarchy panel > IK > Object Parameters rollout > Precedence

You can use the three joint precedence controls to alter how the IK Solution is calculated.

The default joint precedence is 0. This is suitable for many IK solutions. It assumes that joints closest to where a force is applied (the end effector) will move more than joints farther from the force.

Child->Parent—values are calculated by assuming the root of the entire hierarchy has a precedence value of 0 and each child has a value equal to 10 times its depth from the root. In a four-object hierarchy starting with the root, the values would be 0, 10, 20, and 30.

Parent->Child—values are calculated by assuming the root of the entire hierarchy has a precedence value of 0 and each child has a value equal to 10 times its depth from the root. In a four-object hierarchy starting with the root, the values would be 0, -10, -20, and -30.

- Default joint precedence occurs whenever all joints in the kinematic chain have the same precedence value. Assigning a value of 100 to all objects in the kinematic chain is exactly the same as assigning a value of 0.
- An IK chain of three objects with precedence values of 0, 30, and 200 would have the same solution if the precedence were changed to be 1, 2, and 3.

Procedures

To assign a precedence value to an object manually:

1. Select an object in an HD IK Chain.
2. Open the Object Parameters rollout in the Hierarchy panel.
3. Enter a value in the Precedence field.

To assign Child->Parent precedence to all objects in a kinematic chain:

1. Select all objects in an HD IK chain.
2. Open the Object Parameters rollout in the Hierarchy panel.
3. Click Child>Parent.

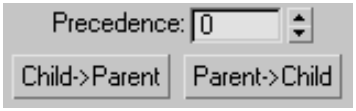
To assign Parent->Child precedence to all objects in a kinematic chain:

1. Select all objects in the kinematic chain.
2. Open the Object Parameters rollout in the Hierarchy panel.
3. Click Parent->Child.

To reset the default joint precedence:

1. Select all objects in the kinematic chain.
2. Open the Object Parameters rollout in the Hierarchy panel.
3. Set the value in the Precedence field to 0.

Interface



Three controls set precedence:

Precedence—Manually assigns precedence values to any object in the IK chain. High precedence values are calculated before low precedence values. Precedence values that are equal are calculated in Child>Parent order.

Child->Parent—Automatically sets joint precedence to decrease in value from child to parent. Causes joints closest to where a force is applied (the end effector) to move more than joints farther away from the force.

You almost always assign Child>Parent precedence to an entire kinematic chain.

Child>Parent behaves like the default precedence settings but is more flexible if you want to go back and manually change the settings.

Parent->Child—Automatically sets joint precedence to decrease in value from parent to child. Causes joints closest to where a force is applied (the end effector) to move less than joints farther away from the force.

You almost always assign Child>Parent precedence to an entire kinematic chain.

Copying, Pasting, and Mirroring Joint Parameters (HD Solver)

Make a selection. > Hierarchy panel > IK > Object Parameters rollout > Copying, Pasting, and Mirroring joint parameters

Use the Copy and Paste buttons in the Object Parameters rollout to copy and paste the joint settings between IK controllers. You can also copy joint settings from a non-IK controller to an IK controller, but you can't copy from an IK controller to a non-IK controller. With an HD IK solver hierarchy, you can paste to a multiple selection of joints.

Use Mirror Paste to mirror the IK joint settings about the X, Y, or Z axis during a Paste operation. It consists of four buttons: one each for the X, Y, and Z axes, and a None option that prevents mirroring altogether.

You can also mirror the IK joints using the Mirror tool on the main toolbar.

See also

Sliding and Rotational Joints (page 2–496)

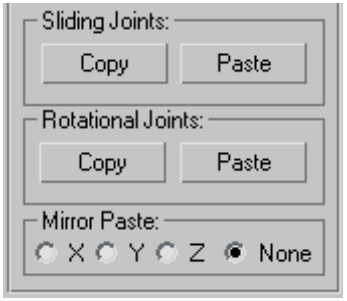
Procedure

To copy and paste joint parameters:

1. Select an object in an inverse kinematic chain.
2. Set sliding or rotational joint parameters.
3. Click Copy in either the Sliding Joints or Rotational Joints group on the Object Parameters rollout.
4. Select a different object in the inverse kinematic chain.
5. Select an axis to mirror in the Mirror Paste group.
6. Click paste in either the Sliding Joints or Rotational Joints group on the Object Parameters rollout.

   Both Sliding Joints and Rotational Joints maintain separate clipboards. The copied joint parameters are saved in the clipboard until you replace them with new copied parameters.

Interface



Sliding Joints group

Use these buttons to copy sliding joint parameters from one object to another. These buttons are not available for Path joints.

Rotational Joints group

Use these buttons for copying rotational joint parameters from one object to another.

Mirror Paste group

Use to mirror the IK joint settings about the X, Y, or Z axis during a Paste operation.

Sliding and Rotational Joints Rollouts (HD Solver)

Make a selection. > Hierarchy panel > IK > Sliding and Rotational Joints rollouts

In inverse kinematics, joints operate by allowing motion on one or more axes and restricting motion on the remaining axes. You set whether an object can move (slide) on or rotate about a given axis with the Active check box in the joint rollouts. Joints have a maximum of six possible axes: three for position and three for rotation. You constrain the motion of a joint by setting which axes are active.

The setting of IK joint axes overrides any Inherit and Lock settings in *Link Info (page 2–499)*.

Procedures

To activate or deactivate an axis:

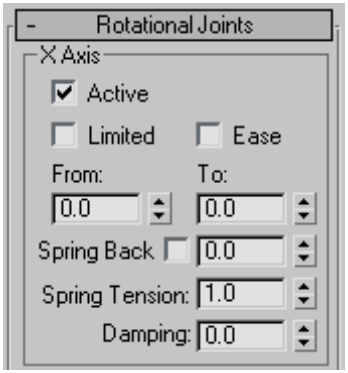
1. Select an object.
2. On the Hierarchy panel, click IK.
3. Expand the rollout for a joint type.
4. Turn on Active for any available axis.
   - When on, the axis is active and the object can move or rotate about that axis.
   - When off, the axis is inactive and the object cannot move or rotate about that axis.

To limit a joint:

1. Select an object.
2. Click IK in the Hierarchy panel.
3. Turn on Limited for any active axis. If an axis is inactive the limits settings are ignored.
   - When on, the axis is limited and the object can move or rotate within the range set by the From and To fields.
   - When off, the axis is not limited and the object moves or rotates freely about that axis. If you turn on Limited, specify limit values in the From and To fields. When setting limits, the object transforms to the position or rotation defined by the From or To field. When you exit the field or release the spinner, the object returns to its original position. You can also view joint limits in the viewport by pressing and holding the mouse button on the From or To label of a limit field. The object will move or rotate to the limit value until you release the mouse button.

Interface

Sliding/Rotational Joints rollouts



The Sliding Joints rollout and Rotational Joints rollout contain similar controls for position and rotation, respectively. Also, each rollout contains identical group boxes for control of the X, Y, and Z axes.

X/Y/Z Axis groups

Active—Activates an axis (X/Y/Z). Allows the selected object to slide on or rotate about the activated axis.

Limited—Limits the range of motion or rotation allowed on an active axis. Use in conjunction with the From and To spinners. Most joints are limited in their range of motion along an active axis. For example, a piston slides only within the length of its cylinder.

Ease—Causes a joint to resist motion as it approaches its From and To limits. Simulates an organic joint, or worn mechanical joint, moving or rotating freely in the middle of its range of motion but moving less freely at the extremes of its range.

From and To Spinners—Determine for position and rotation limits. Use in conjunction with the Limited function.

Spring Back—Activates Spring Back. Each joint has a rest position. As the joint moves further from the rest position, an increasingly larger force pulls the joint back to its rest position, like a spring.

Spring Back (spinner)—Sets the rest position for the joint. For rotational joints, this is the orientation of the joint in degrees; for sliding joints, it's the position in units. Adjusting this is similar to adjusting the From/To spinners. You see the orientation/position while performing the adjustment, but when you release the spinner the object returns to its previous state.

If you're using the HD Solver, you must turn on Show Initial State (Motion panel > IK Controller Parameters rollout) to see the effect on the geometry during the adjustment .

Spring Tension—Sets the strength of the "spring." Higher values cause the spring to pull harder as the joint moves further from its rest position. A setting of 0 turns off the spring; very high settings can turn the joint into a limit because you can reach the point where the spring is too strong to allow the joint to move past a certain point.

Damping—Applies resistance over a joint's full range of motion or rotation. Simulates the natural effect of joint friction or inertia. As a joint corrodes, dries out, or is put under a heavy load, it resists motion along its active axes.

Interactive and Applied IK Rollouts

Inverse Kinematics Rollout

Make a selection. > Hierarchy panel > IK > Inverse Kinematics rollout

The Inverse Kinematics rollout displays different controls based on IK Solvers applied to the selected

hierarchy. IK Solvers are applied to the hierarchies using the Animation menu.

When an HD IK Solver is applied to an IK chain the Inverse Kinematics rollout displays the controls described below.

The Inverse Kinematics rollout provides controls for interactive and applied IK, as well as the controls for the HD Solver (history dependent). Use Apply IK to calculate an IK solution and generate Transform keys (move, rotate) for all objects in an IK chain (the IK chain must include a follow object). By default, keys are created at every frame.

- Applied IK requires that one or more parts of your IK structure be bound to animated follow objects. Once bound, you can select any object in your kinematic chain and click the Apply IK button.
- The Apply IK method of animation works best when you want objects in the kinematic chain to match the motions of other objects exactly. An example of this would be a mechanical connection that should always appear to be attached to the follow object.

Procedures

To use interactive IK:

1. Build a hierarchy of objects.

   If you want to use Bones, be sure Assign to Children is not active when you create the Bones.
2. Select the end of the chain.
3. Go to the Hierarchy panel > IK button.
4. Scroll up to the Inverse Kinematics Rollout
5. Click Interactive IK.
6. Advance the time slider and move the end of the chain.
7. Turn off Interactive IK when you want to do forward kinematics.

To use applied IK:

1. Add an HD Solver to your hierarchy.
2. Bind objects in your IK structure to follow objects.
3. Animate the follow objects.
4. Select any object in the IK structure.
5. Click IK in the Hierarchy panel and expand the Inverse Kinematics rollout.
6. Click Apply IK.

   Tip: Be sure your Start Frame and End frames match the length of your animation.

Interface



Interactive IK— Allows for IK manipulation of hierarchies without applying an IK Solver or using a follow object.

Apply IK—Calculates the IK solution for each frame of the animation and creates transform keys for every object in the IK chain. A bar graph appears on the prompt line to indicate progress of the calculations.

Apply Only To Keys—Solves the IK solution for keyframes that already exist for one of the end effectors.

Update Viewports—Views the progress of Apply IK frame by frame in the viewports.

Clear Keys—Removes all move and rotate keys from the selected IK chain before applying IK.

Start/End—Sets the range of frames to calculate the applied IK solution. The default settings for Apply IK calculates the IK solution for every frame of the active time segment.

Auto Termination Rollout (Interactive IK)

Make a selection. > Hierarchy panel > IK button > Auto Termination rollout

The Auto Termination controls temporarily assign terminators a specific number of links up the hierarchical chain from the selected object. This only works with Interactive IK; it does not work with applied IK or with IK solvers.

Procedure

To use auto termination:

1. In the Auto Termination rollout, turn on Auto Termination.
2. Enter a value in the # of Links Up field.

   Turn on the IK button on the toolbar.
3. Select any object in an IK chain to move or rotate.

Interface



Interactive IK Auto Termination—Turns on the auto termination feature.

# of Links Up—Specifies how far up the chain the termination is applied.

For example, if you set this to 5, when you move any object in the hierarchy, the object that's five links up the chain from the object you're adjusting acts as a terminator. If you select a different object in the hierarchy, termination is switched to whichever object is five links up the chain from the newly selected object.

A setting of 1 simply locks the hierarchy, since it terminates the joint just beyond the currently selected object.

Link Info

Make a selection. > Hierarchy panel > Link Info button

This part of the Hierarchy panel contains two rollouts. The Locks rollout has controls to restrict the movement of objects in a particular axis. The Inherit rollout has controls to limit the transforms that a child inherits from its parent object.

Interface



The Link Info panel has two rollouts:

Locks rollout (page 2–500): Locks an object's ability to Move, Rotate, or Scale about any of its local axes.

Inherit rollout (page 2–500): Limits what transforms a child inherits.

Locks Rollout

Make a selection. > Hierarchy panel > Link Info button > Locks rollout

The Locks rollout contains controls that prevent transforms along particular axes.

See also

Inherit Rollout (page 2–500)

Procedure

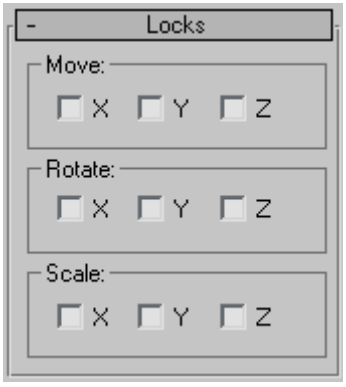
To lock all move transforms:

1. Select an object in the viewports.
2. In the Locks rollout, turn on Move group > X, Y, and Z.

   Only rotate and scale transforms will work.

   The object does not need to be part of a hierarchy.

Interface



X, Y, Z—Turn on any axis in the Move, Rotate, or Scale group box to lock the axis. For example, if Rotate > X and Y are turned on, you'll be able to rotate the object only around the Z axis. All locks are relative to an object's local coordinate system.

Inherit Rollout

Make a selection. > Hierarchy panel > Link Info button > Inherit rollout

The Inherit rollout constrains the links between a selected object and its parent for any axis of position, rotation, or scale.

See also

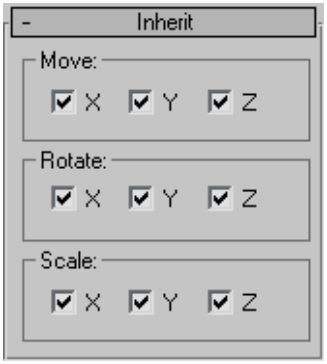
Locks Rollout (page 2–500)

Procedure

To prevent inheritance of X rotation in a hierarchy:

- Select an object, then turn off X in the Rotate group on the Inherit rollout.

Interface



X, Y, Z—Turn off any axis in the Move, Rotate, or Scale group boxes to prevent inheritance.

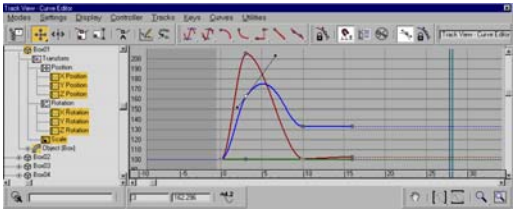
When you turn on an axis, transform information passes from the parent to the child for that axis. When you turn off an axis, transform information on that axis is ignored by the child.

Track View

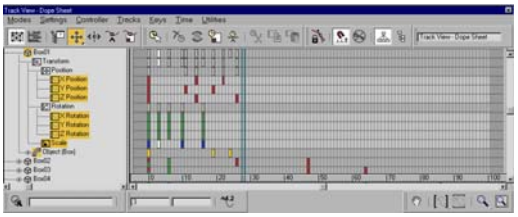
Main toolbar > Curve Editor (Open)
Graph Editors menu > New Track View
Graph Editors menu > Track View - Curve Editor
Graph Editors menu > Track View - Dope Sheet
Graph Editors menu > Saved Track View > Choose a saved Track View.
Right-click active viewport. > Transform quadrant > Curve Editor or Dope Sheet
Viewport right-click menu > Views > Track > Choose New or a saved Track View.

With Track View, you can view and edit all the keys that you create. You can also assign animation controllers to interpolate or control all the keys and parameters for the objects in your scene.

Track View uses two different modes, Curve Editor and Dope Sheet. Curve Editor mode lets you display the animation as function curves. Dope Sheet mode displays the animation as a spreadsheet of keys and ranges. Keys are color-coded for easy identification. Some of the functions in Track View, such as moving and deleting keys, are also available on the track bar near the time slider, which can be expanded to show curves as well. You can dock the Curve Editor and Dope Sheet windows beneath the viewports at the bottom of the interface, or use them as floating windows. Track View layouts can be named and stored in the Track View buffer and reused. Track View layouts are stored with the MAX file.



Track View — Curve Editor



Track View — Dope Sheet (Edit Keys)

Typical Uses for Track View

Track View can perform a variety of scene management and animation control tasks. Use Track View to:

- Display a list of objects in your scene and their parameters.
- Change key values.
- Change key timing.
- Change controller ranges (see procedure).
- Change interpolation between keys.
- Edit ranges of multiple keys.
- Edit blocks of time.
- Add sound to your scene.
- Create and manage notes about the scene.
- Change the behavior of the animation outside the range of keys.
- Change controllers for animated parameters
- Select objects, vertices, and hierarchies.
- Navigate the modifier stack in the Modify panel by clicking the modifier items in the Track View Hierarchy.

  Note: Tracks are created for animated vertices in Track View. A Bezier Point3 controller is the default vertex interpolation controller.

Procedures

To change the frames in which a controller takes effect:

When you apply a controller or constraint to an object's motion, the frame range over which controller takes effect is determined by the current *active time segment (page 3–904)*. If you then change the active time segment or the animation length, the duration of the controller's influence doesn't change. Sometimes applying a controller (such as Path Constraint) automatically sets keys that you can use to change this range. But others, such as Noise controllers, don't set keys. In such cases, follow this procedure:

1. Select the object, and then right-click it and choose Curve Editor from the menu.
2. Expand the object hierarchy to find the track or tracks to adjust.
3. From the Modes menu, choose Dope Sheet.
4.  On the Dope Sheet, click the Edit Ranges button.
5. Adjust the range duration by dragging its endpoints, or its position in the animation by dragging between endpoints.

   For more information, see *Dope Sheet (page 2–507)*.

To select keys in Track View (either mode), do any of the following:

- Click the key to select an individual key.
- Drag a selection rectangle around keys to select multiple keys.
- Hold down the Ctrl key and click to create discontinuous multiple key selections.

Note: If you are in Dope Sheet – Edit Ranges mode, you can use Select Time to select multiple keys.

To delete keys in Track View (either mode):

1. Select keys on the curve, or on the dope sheet.
2. Press the Delete key on the keyboard to delete the selected keys.

To force Track View to always display on a second monitor:

If you are running a dual-monitor setup, you can force Track View to display on the right-hand monitor by editing a script.

1. Right-click the Curve Editor (Open) icon in the main toolbar, and then choose Edit Macro Script.

   The MAXScript script that opens the Track View – Function Curve Editor appears.
2. Locate the line that reads:

   if (trackviews.open "Track View - Curve Editor" layoutName:"Function Curve Layout") == true then
3. Replace that line with this one:

   max_window = getMaxWindowSize() --get Desktop size if (trackviews.open "Track View - Curve Editor" layoutName:"Function Curve Layout" pos:[max_window.x/2,0] height:max_window.y width:(max_window.x/2)) == true then
4. Save the script and restart 3ds Max.

This should open the Track View *in a new session* over the right half of the desktop. Assuming that a dual-monitor setup reports twice the width, this will force the Track View on the second monitor. On a single monitor, it opens it over the right half of the monitor. Of course, you could enter your own numbers like pos:[1024,0] height:768 width:1024 in case you are running two monitors at 1024x768.

To open Track View in a viewport:

- Right-click a viewport label, and then from the menu that opens choose Views > Track > New, or click the name of a saved Track View.

See also

Track View Menu Bar (page 2–521)

Curve Editor (page 2–507)

Dope Sheet (page 2–507)

Track View Workspace (page 2–504)

Time Ruler (page 2–510)

Filters (page 2–541)

Assign Controller (page 2–546)

Copy Controller (page 2–544)

Paste Controller (page 2–545)

Make Controller Unique (page 2–550)

Parameter Curve Out-of-Range Types (page 2–551)

Add Note Track (page 2–552)

Remove Note Track (page 2–553)

Edit Keys (page 2–554)

Track View Utilities (page 2–561)

Select Time (page 2–566)

Edit Ranges (page 2–573)

Curve Editor (page 2–578)

Status Bar/View Controls (page 2–588)

Track View Hierarchy (page 2–512)

Animation Controllers (page 2–307)

Working with Track View

Track View is the tool that you use to see a data-driven view of your scene and animation.

The standard views, such as Top, Front, User, and Camera, are geometry-driven views. These views display the surfaces of your objects and show them changing over time.

Track View displays the values and timing that produce the geometry and motion you see in standard views. Using Track View, you have very precise control over every aspect of your scene.

Track View has two modes: Curve Editor and Dope Sheet. The Curve Editor displays your animation as keys on function curves; editing the tangency of the keys allows control of the in-betweens. Dope Sheet displays your animation as keys and ranges on a box grid and allows you to adjust the timing of your motions.

Track View keys and curves can also be displayed in the track bar. The same Key properties dialogs available on Track View can also be found on the Motion panel as well.

For details about using Track View and editing animation, see the following:

- *Track View Edit Window (page 2–504)*
- *Track View Menu Bar (page 2–521)*
- *Curve Editor Toolbars (page 2–535)*
- *Dope Sheet (page 2–507)*
- *Dope Sheet Toolbars (page 2–538)*

Understanding Track View Concepts

The left side of Track View, called the Controller Window, presents a Hierarchy list of everything in your scene. Every object and environment effect appears in the list, along with its associated animatable parameters. Choose items from this list to apply changes to the animation values. Expand or collapse the list using manual navigation, or allow the Auto expand to determine the display in this window.

The right side of Track View, called the Key window, charts the changes applied to parameters over time. Any change you make to one of these parameters when the Auto Key button is on, appears as a key in the right side of Track View. Select keys to apply changes to one or more specific keys.

This section covers some of the editing functions you can use with Track View:

Copying and Pasting Items (page 2–575)

Copying and Pasting Objects (page 2–576)

Making Instance and Reference Controllers and Objects Unique (page 2–577)

Track View Workspace 

Main toolbar > Curve Editor (Open)
Graph Editors menu > Track View - Curve Editor
Graph Editors menu > New Track View

The Track View Key window displays function curves and keys when in Curve Editor mode. When in Dope Sheet mode, the tracks can be displayed as keys or ranges. You select and change animation values and timing in these windows. The Key window also indicates the active time segment. Time within the active time segment is highlighted with a light gray background. The Track View Key window is sometimes referred to as the Key window.

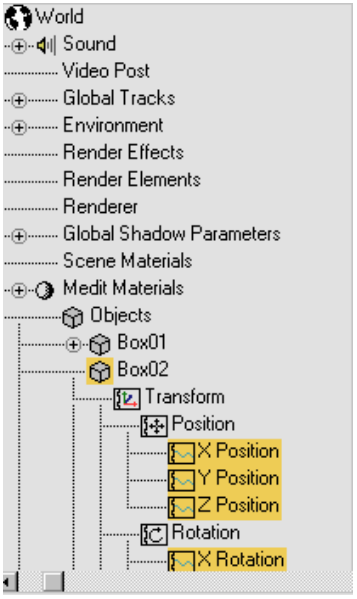
Components of the Key window include a time slider, a time ruler, and a scale origin slider. The Track View time slider indicates the current frame and is synchronized with the viewport time slider. A time ruler at the bottom of the window can be raised to measure keys against time. A scale origin indicator (a horizontal orange line at 0) can be moved during scale value operations as a reference point for scaling.

Interface

The two main sections of the Track View workspace consist of the Key window and the Controller window.

Controller Window

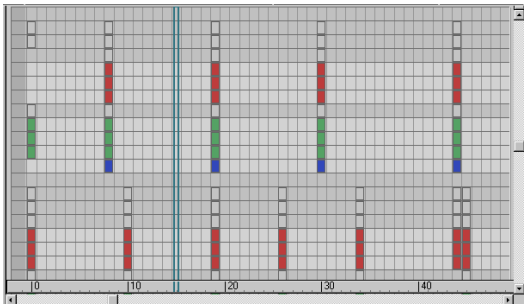
The Controller window displays object names and controller tracks, and determines which curves and tracks are available for display and editing. Hierarchy items in the Controller window can be expanded and rearranged as necessary using the Hierarchy list right-click menu. The navigation tools can also be found in the Track View Options menu. Default behavior is to only show selected object tracks. Use Manual Navigation mode to collapse or expand tracks individually, or press Alt+right-click to display an alternate menu to collapse and expand tracks.



Controller window

Key window

The Key window displays the keys as either curves or tracks. The tracks can be displayed as a box graph of keys or range bars.



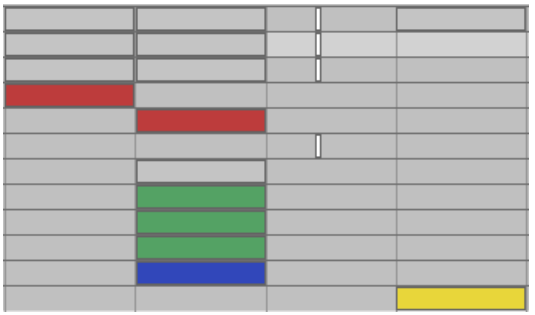
Keys displayed as box graph (Dope Sheet - Edit Keys mode)

Key Creation

Keys are created using a variety of methods. Keys can be created by turning on Auto Key, moving the time slider, and then transforming the object or adjusting its parameters. Keys can also be created by right-clicking the viewport time slider to access the Create Key dialog. Keys can be created directly in Track View using Add Keys. Finally, keys can be created by turning on Set Key mode, moving to a desired frame, posing the object, then clicking Set Key.

Key Display

Keys are displayed as points on the function curves, or as boxes on the Dope Sheet. Keys on the Dope Sheet are color-coded for easy identification. When there are multiple tracks keyed at one frame, the boxes appear with bands of color to indicate the shared key types. Key coloration is also used to show soft selection of keys. Subframe keys (keyframes that fall between frames) are indicated as narrow rectangles within the boxes.

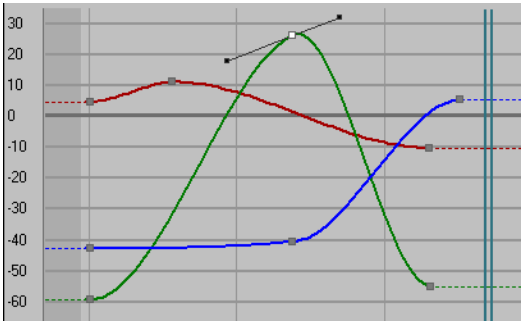


Colored keys with subframe display

Keys are also displayed on the track bar below the viewport.

Keys displayed on the function curves have tangency types. The tangency buttons found on the Key Tangents toolbar can be used to change the function curve keys. Use Custom tangency to show editable curve handles. Use Step tangency to freeze motion or create classic storyboard pose-to-pose blocking.

Function curves can also be displayed below the track bar.

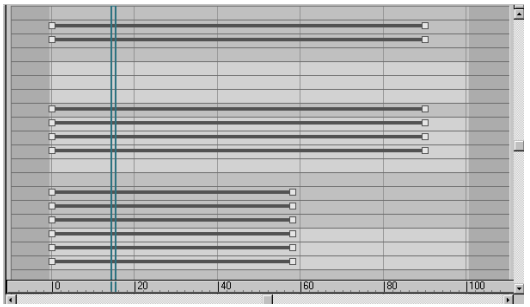


Custom Tangency handles

Range Bars

In the Dope Sheet - Edit Ranges mode (when animation keys have been created), range bars display to indicate the range of time the animation occurs. Tools specific to working with ranges (position ranges and recouple ranges) can be found on the Ranges toolbar (off by default).

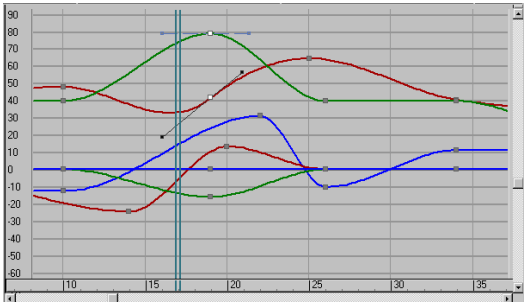
Right-click the toolbar, choose Show Toolbars, and then select Range-Track View to access these tools.



Keys displayed as range bars (Dope Sheet - Edit Ranges mode)

Function Curves

Function Curves display the values of keys, and the interpolated values between keys, as a curve. These curves express how a parameter varies over time. Only animation tracks can display function curves. You can edit the curves using tangency handles on the keys to change the shape of the curve.

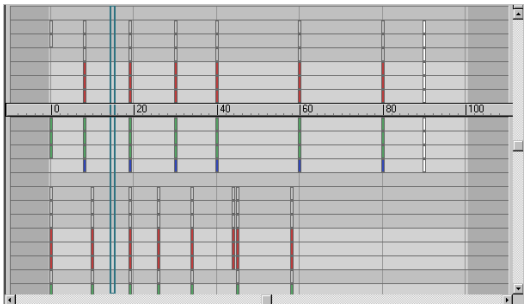


Keys displayed as function curves (Curve Editor)

Multiple curves can be viewed simultaneously by selecting tracks in the Controller window. This is especially useful when using Multiplier or Ease Curves. You can adjust the multiplier or ease curve point tangencies and watch the final result in the controller curve at the same time.

Time Ruler

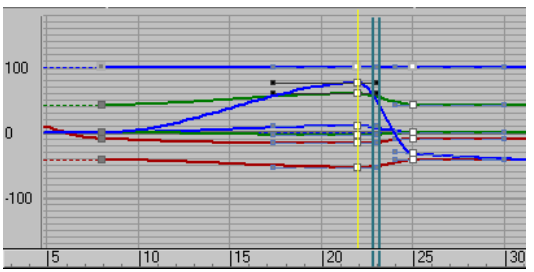
The time ruler measures time. Markings on the ruler reflect the settings in the Time Configuration dialog. Move the time ruler up to the keys for more accurate key placement.



Time ruler

Track View Time Slider

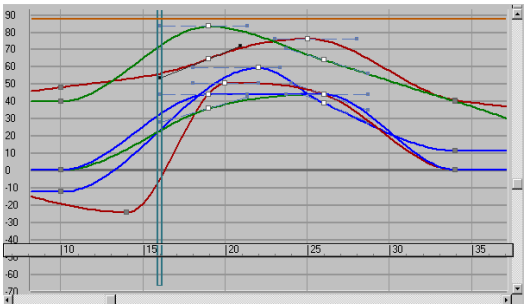
The current time is indicated by the Track View time slider. This is displayed as a set of blue vertical lines synchronized with the position of the viewport time slider. You can move the Track View time slider by dragging it in the Key window. Moving either time slider updates the animation in the viewports. The blue time slider also serves as a scale origin point when scaling keys in time.



Scale keys using time slider

Scale Values Origin Line

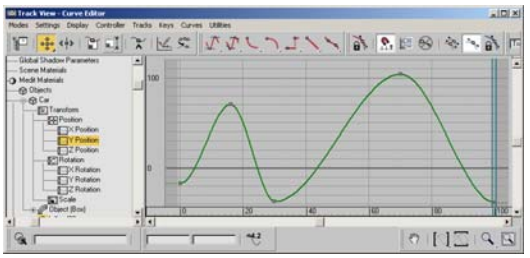
When you choose to scale key values (that is, scaling in space rather than time), a horizontal orange line appears at the 0 value on the vertical graph axis. This orange line is a scale values origin indicator, which you can move verticall to use as a reference point for scaling values.



Horizontal line at the top of the illustration is the moveable scale origin line.

Curve Editor

Menu bar > Graph Editors > Curve Editor
Menu bar > Graph Editors > New Track View > Modes > Curve Editor
Main toolbar > Curve Editor (Open)
Right-click the selected object in the viewport. > Curve Editor



The Track View - Curve Editor is a Track View mode that allows you to work with motion expressed as function curves on a graph. It lets you visualize the interpolation of the motion, the object transformations that the software creates between the keyframes. You can easily see and control the motion and animation of the objects in the scene using tangent handles on the keys found on the curves.

The Curve Editor interface consists of a menu bar, a toolbar, a Controller window, and a Key window. There is also a time ruler, and navigation and status tools at the bottom of the interface.

You can loop or cycle your animation beyond its range by adding Parameter Curve Out-Of-Range Types from the Curve Editor, as well as by adding Multiplier or Ease Curve onto other animated tracks for added control.

Tip:  You can also click Show Curves in the track bar to display function curves.

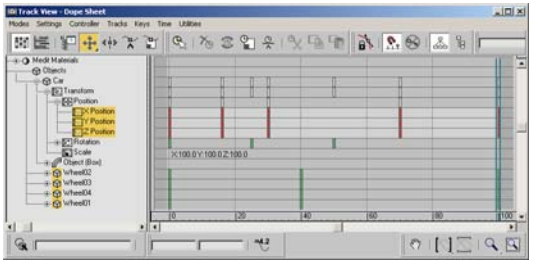
See also

Track View Menu Bar (page 2–521)

Curve Editor Toolbars (page 2–535)

Dope Sheet

Menu bar > Graph Editors > Track View - Dope Sheet
Menu bar > Graph Editors > New Track View > Modes > Dope Sheet
Main toolbar > Curve Editor (Open) > Modes menu > Dope Sheet
Right-click the selected object in the viewport > Curve Editor > Modes > Dope Sheet



The Track View - Dope Sheet editor displays keyframes over time on a horizontal graph. This graphical display simplifies the process of adjusting animation timing because you can see all keys at once in a spreadsheet-like format.

Classical animation technique included the use of an exposure sheet, called an “X” sheet or a Dope

Sheet. The Dope Sheet was a vertical chart that served as instructions to the camera operator. Dialogue and camera actions were indicated over a numbered list that represented each shot, which became a single photographed frame of the animated movie. The classical exposure sheet also included instructions for compositing the cel drawings of animated characters over backgrounds. This device serves as inspiration for the Dope Sheet tool in 3ds Max.

The 3ds Max Dope Sheet editor is similar to the classic X sheet. It displays keyframes over time, only using a horizontal graph (rather than vertical). This provides tools for adjusting the timing of your animation. Here, you can see all the keys in a spreadsheet-type interface. You can select any or all of the keys in a scene, scale them, move them, copy and paste them, or otherwise work directly here, rather than in the viewport. You can choose to select the keys for children, or subtree, or both, so you can make simple changes that affect many objects and their keys at once.

A common use of Dope Sheet is to stagger the movement of a character's limbs so they don't all move simultaneously. If you have a crowd of characters, you could use Dope Sheet to shift movements so they don't all move in unison.

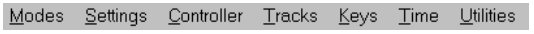
Dope Sheet Menus and Tools

In the Dope Sheet, you can select any or all of the keys in a scene, scale them, move them, copy and paste them, or otherwise work directly, rather than working with objects in viewports. You can choose to select the keys for children, or subtree or both, so you can make simple changes that affect many objects and their keys at once.

Dope Sheet allows for soft-selection of keys which is very useful when working with motion capture data that has keys on every frame.

Dope Sheet provides tools for working directly with time. You can select, cut, copy, paste, and insert and reverse time using the tools on the Time menu.

Like the *Curve Editor (page 2–507)*, Dope Sheet has both a menu bar and toolbars to provide you quick access to tools.

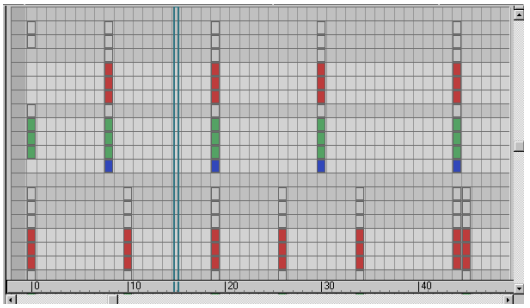


Dope Sheet menu bar

Unlike Curve Editor, Dope Sheet has two modes: Edit Keys and Edit Ranges. These modes change the display in the Key window.

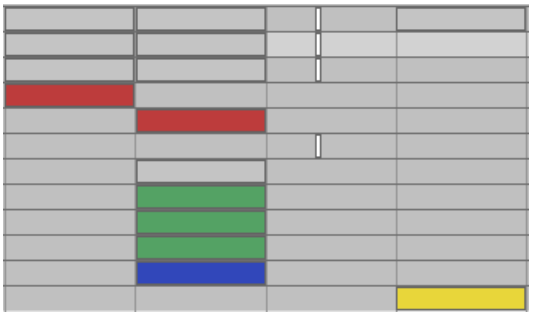
Edit Keys Mode 

When Edit Keys is turned on, the keyframes are displayed as boxes within rectangles on a grid. The keys are color-coded to show what has been keyframed (position is red, scale is yellow, rotation is green, and so on.)



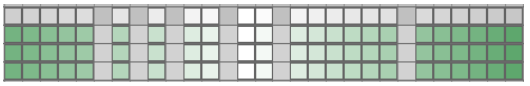
Colored keys

Dope Sheet keys are now displayed as rectangles within boxes so you can easily spot *sub-frame keys*, keys that fall in-between frames. Keys that fill the boxes are on the frame, keys that are small rectangles are sub-frame.



Sub-frame keys

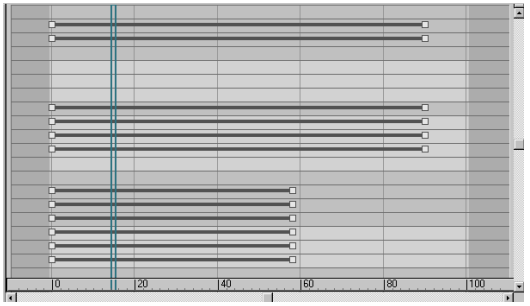
Dope Sheet, just like the Curve Editor, allows you to use soft selection on keys. This is extremely useful when you are dealing with massive quantities of keys, such as in motion-capture data files. Combine this with scaling keys for a means to manipulate motion data.



Soft selection of Dope Sheet keys

Edit Ranges Mode 

When Edit Ranges is turned on, the animation tracks are displayed as range bars; no individual keys are visible. Use Ranges when you only want to change how long something happens, or when it starts and ends, rather than when you need to manipulate particular keys within an action.



Edit Ranges mode

Modify Subtree and Modify Child Keys

When working in Dope Sheet, you can turn on or off Modify Subtree and Modify Child Keys. These let you automatically move the keys for the children, and/or the tracks for the subtree. If you experience a slowdown while working with Dope Sheet, try turning these off, and moving the keys manually instead. Modify Subtree is on by default in Dope Sheet, but off in Curve Editor.

Time Editing

Dope Sheet offers you a variety of tools for working directly with time. You can select a period of time, which includes all the keys within that period, and then perform different operations on that time segment. You can copy and paste time to loop animations, or reverse time so the animation plays backward. You can insert time to add a space to an animation, or delete time to shorten a motion.

Note: For Dope Sheet procedures, see the individual tools and menu choices within the links below.

Dope Sheet Speed Improvements

In 3ds Max, only tracks displayed in the current field of view are computed, optimizing the responsiveness of the Dope Sheet editor.

Dope Sheet now temporarily retains the key caches so tracks that have already been computed as a result of the parent being computed needn't be recomputed. The cache is used instead.

The default auto-navigation settings for the Dope Sheet editor now only auto-expand to the node track for the currently selected object. This reduces the number of tracks whose keys need to be displayed and also helps enforce the top-down workflow the Dope Sheet editor is designed for.

See also

Track View Menu Bar (page 2–521)

Dope Sheet Toolbars (page 2–538)

Select Time (page 2–566)

Edit Ranges (page 2–573)

Edit Keys (page 2–554)

Time Ruler

Main toolbar > Curve Editor (Open) > Time Ruler (below right-hand pane)

Graph Editors menu > Track View - Curve Editor/Track View - Dope Sheet > Time Ruler (below right-hand pane)



The time ruler at the bottom of the Track View Key window measures time. Markings on the time ruler reflect the settings in the *Time Configuration dialog (page 3–725)*. You can drag the time ruler vertically in the Key window to align it with any track.

Track View Shortcuts

Track View functions for which keyboard shortcuts can be set are listed below. Where there is no default keyboard listed and no button shown in the action list, a short description appears in the right column.

To use Track View keyboard shortcuts, the *Keyboard Shortcut Override toggle (page 3–872)* must be on.

See also

Track View (page 2–501)

Keyboard Shortcuts (page 3–871)

Keyboard Panel (page 3–793)

Customize User Interface Dialog (page 3–792)

Track View Function	Default Keyboard Shortcut	Description
Access Hierarchy Select Name Field		Accesses the track selection field at the lower left of Track View window
Access Time Field		Accesses the time field at the lower left of Track View window
Access Track Name Field		Accesses the name of the Track View window at upper right
Access Value Field		Accesses the value field at the lower left of Track View window
Add Keys	A	
Add Note Track		
Add Visibility Track		
Align Keys		
Apply Ease Curve	Ctrl+E	
Apply Multiplier Curve	Ctrl+M	
Assign Controller	C	
Auto Expand Animated		
Auto Expand Keyable		
Auto Expand Limits		
Collapse All		
Collapse Objects		
Collapse Tracks		
Copy Controller	Ctrl+C	
Copy Time		
Cut Time		
Delete Controller		
Delete Ease/ Multiplier Curve		
Delete Note Track		

Track View Function	Default Keyboard Shortcut	Description
Delete Visibility Track		
Draw Curves		
Ease Curve Out-of-Range Types		
Ease/Multiplier Curve Enable Toggle		
Edit Keys Mode		
Edit Ranges Mode		
Edit Time Mode		
Edit Track Set		
Exclude Left End Point Toggle		
Exclude Right End PointToggle		
Expand Object Toggle	O (letter "o")	
Expand Objects		
Expand Track Toggle	ENTER T	
Expand Tracks		
Filters	Q	
Freeze Nonselected Curves		
Function Curves Mode		
Ignore Animation Range		
Insert Time		
Keyable Properties Toggle		
Lock Selection	Spacebar	
Lock Tangents Toggle	L	
Make Controller Unique	U	

Track View Function	Default Keyboard Shortcut	Description
Modify Subtree Toggle		
Move Highlight Down	Down Arrow	
Move Highlight Up	Up Arrow	
Move Keys	M	
Move Keys Horizontal		
Move Keys Vertical		
Move Object Down		Moves an object down in the hierarchy display
Move Object Up		Moves an object up in the hierarchy display
Multiplier Curve Out-of-Range Types		
Nudge Keys Left	Left Arrow	
Nudge Keys Right	Right Arrow	
Pan	P	
Parameter Curve Out-of-Range Types		
Paste Controller	Ctrl+V	
Paste Limit Only		
Paste Time		
Position Ranges Mode		
Properties		
Reduce Keys		
Remove Limit		
Respect Animation Range		
Reverse Time		
Scale Keys		
Scale Time		

Track View Function	Default Keyboard Shortcut	Description
Scale Values		
Scroll Down	Ctrl+Down Arrow	
Scroll Up	Ctrl+Up Arrow	
Select All		
Select Children		
Select Invert		
Select None		
Select Time		
Select Lower Limit		
Set Tangents to Auto		
Set Tangents to Custom		
Set Tangents to Fast		
Set Tangents to Linear		
Set Tangents to Slow		
Set Tangents to Smooth		
Set Tangents to Step		
Set Upper Limit		
Show All Tangents		
Show Keys on Frozen Curves		
Show Selected Key Stats		
Show Tangents Toggle		
Slide Keys		
Snap Frames	S	
Toggle Limit		
Track View Utilities		
Zoom	Z	

Track View Function	Default Keyboard Shortcut	Description
Zoom Horizontal Extents		
Zoom Horizontal Extents	Alt+X	
Zoom Region		
Zoom Selected Object		
Zoom Time		
Zoom Value Extents		
Zoom Values		

Track View Hierarchy

Main toolbar > Curve Editor (Open) > Track View Controller Window

Graph Editors menu > Track View - Curve Editor/Track View - Dope Sheet

The Track View Controller window displays all objects in your scene in a hierarchical fashion. You select objects in the Hierarchy list to:

- Select object and track labels for Track View operations.
- Select object icons to select objects in the scene.
- Select material or map icons when working in the Material Editor.
- Select modifiers in Track View to navigate the modifier stack in the Modify panel.

See also

Hierarchy Right-Click Menu (page 2–516)

Properties (Track View Hierarchy) (page 2–519)

Modify Subtree (page 2–528)

Sound Options Dialog (page 2–520)

Procedures

To select objects in the scene using the Track View Controller window:

1. Do one of the following to select the first object:
   - Click an object's icon to select a single object.
   - Double-click a parent object's icon to select an object and its descendants.
2. Do one of the following to select additional objects:
   - Press Shift and click an object icon to select all objects between the object and the previously selected object.
   - Press Ctrl and click an object icon to add (or remove) the object to the selection.

   When objects are selected in the scene, their icons are also highlighted in Track View.

   All objects in your scene are always displayed in Track View. However, selecting object icons in Track View follows the same rules as selecting objects in the scene:
   - You can't select the icon of a hidden object.
   - You can't select the icon of a frozen object.
   - If you select the icon of an object that is part of a closed group, the group icon and all the objects in the group are selected.

To select and deselect item labels by clicking:

1. Select the first item by clicking its label.
2. Do one of the following to add or subtract from the selection:
   - Press Shift and click an item label to select all items between the item and the previously selected item.
   - Press Ctrl and click an item label to add (or remove) that item to the selection.
   - Press Alt and click an item label to select all item labels in the column.

   Right-clicking an item displays a right-click menu for globally selecting items in the Hierarchy list.

To select items in the Controller window with the right-click menu:

- Right-click an item in the Controller window, then do any of the following:
  - Choose Select All to select all items visible in the Hierarchy list. Collapsed items are not selected. Press Ctrl and choose Select All to select only visible transform items.
  - Choose Select Invert to reverse the selection pattern.
  - Choose Select None to clear all items from the selection.
  - Choose Selected Children to select the linked descendants of a parent object.

  When working with function curves, you can also select a controller icon (instead of the label) to select that function curve for vertex editing.

To select a material or map as the active material in the Material Editor:

1. Open the Material Editor.
2. Open the Medit Materials branch in the Hierarchy list of the Controller window and select a material.

   The selected material becomes the active item in the Material Editor. This is one way to navigate and select a material for editing.

Example: To navigate the modifier stack in the Modify panel:

1. Create a cylinder.
2. In the Modify panel, add a Bend and a Twist modifier to the cylinder.
3. In the Track View controller window, expand the Cylinder branch.
4. Expand the Modified Object branch.

5. Click the icon of the Twist modifier.

   The Twist modifier is displayed on the Modify panel.

6. Click the icon of the Bend modifier in the Hierarchy list.

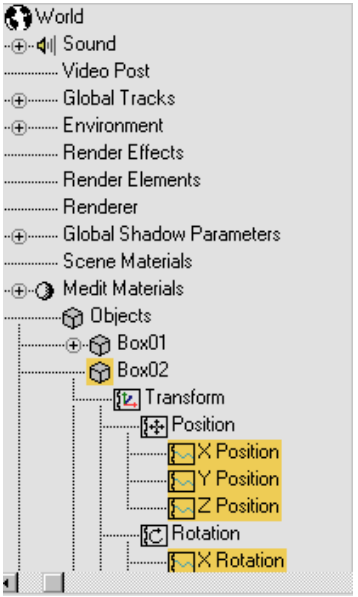
   The Bend modifier is displayed on the Modify panel.

To highlight all tracks containing animation:

1. Create a box and a sphere.
2. Auto Key Click Auto Key, go to frame 10 and move the sphere in X.
3. Go to frame 20 and move the sphere in Y.
4. Go to frame 30 and move the sphere in Z.
5. Deselect the sphere by clicking in a blank area of the viewport.
6. Open the track View — Curve Editor and press Ctrl+A.

   Locate in the Track View Controller window the box and sphere objects. Notice that both their icons are highlighted as well as the sphere's XYZ position tracks.

Interface



Scene Hierarchy

World—World is the root of the scene hierarchy. This track collects all keys in your scene as a single range for quick global operations.

By default, the World track shows the range of the Sound, Environment, Medit Materials, and Scene Materials branches. Modify Subtree in Edit Ranges mode causes the World track to include the range of all tracks in the Objects branch as well.

Note: The default location of the time ruler covers the World animation track. Move the time ruler to see the World track.

Sound—Allows you to synchronize your animation to a sound file or to a metronome. If you use a sound file, a Wave Form track is displayed in the Hierarchy list, and a waveform is displayed in the Edit window. A sound file plays on your installed sound card. You set the Metronome and the Wave Form parameters in the *Sound Options dialog (page 2–520)*.

Global Tracks—Allows you to store controllers for global use. Using expression controllers, for example, you could point to a controller in the Global Tracks from several other tracks. By altering the expression in the Global Tracks, all of the other tracks are changed.

By pasting an instance of a controller in Global Tracks to a number of other tracks, you can change many tracks by altering the controller in Global Tracks.

By default, Global Tracks contains List tracks of the different controller types. To assign a controller, open one of the List tracks, highlight the available track, and then click Assign Controller. Once the controller is assigned, you can point to it from an expression controller assigned to any other track, or you can copy and paste it to any number of matching controller types.

Video Post—Allows you to manage animated parameters for Video Post plug-ins.

Environment—Contains items that control the background and scene environment effects. Examples include Ambient Light, Background definition, Fogs, and Volumetric Lighting.

Render Effects—Contains tracks for the effects added in Rendering menu > Effects. After adding rendering effects, use the tracks here to animate effects parameters such as glow size and color.

Render Elements—Shows you what elements you've chosen to render separately via the Render Scene dialog > Render Elements rollout.

Renderer—Allows you to animate parameters in the renderer. After selecting a type of antialiasing in the Render dialog, you can animate various antialiasing parameters using these tracks.

Global Shadow Parameters—Use these tracks to change or animate shadow parameters for any light that has the Use Global Settings parameter turned on in the Shadow Parameters rollout. (Select a light, then open the Modify panel > Shadow Parameters rollout to toggle the Use Global Settings parameter).

You can change shadow parameters for multiple lights simultaneously. Tracks include *Map Size, Map Range, Map Bias, and Absolute Bias (page 2–1337)*.

Scene Materials—Contains the definitions for all materials in the scene. It is empty until you begin assigning materials to objects. When you select materials in this branch, you are working with instances of materials assigned to objects in your scene. These materials might not be in any of the Material Editor samples.

 **Medit Materials—**Contains global material definitions. The Medit Materials branch contains the 24 material definitions in the *Material Editor (page 2–1409)*. When you select materials from this branch, you are working with global material definitions that might not yet be assigned to objects in your scene.

 **Objects—**Contains a hierarchy for all the objects in your scene.

When *Manual Navigation (page 2–529)* is on, many of these categories display a small plus or minus button to expand/collapse the contents.

Item Categories

Each type of item in the Hierarchy list is represented by an icon. You can use these icons to quickly identify what each item represents.

 **Objects—**Branches contain linked descendants of the object. Branches below contain transforms, materials, and modifiers applied to the object, as well as the object's creation parameters.

Controllers—They contain the animation values for parameters and are the only item in the Hierarchy list that can have a track containing keys. Every controller has an individual icon



Examples of controller icons

Certain types of controllers can contain other controllers. Examples of these are Transform Controllers and List Controllers.

Note: Controllers that have been applied to objects via the Animation menu > Controllers or Constraints submenus automatically have list controllers assigned. Controllers assigned using the Track View Controller menu or right-click Assign Controller choice on the quad menu do not apply list controllers automatically.

Maps—Indicates map definitions. All branches below a map definition are part of that map. This includes values used by parametric maps and other map definitions that are part of a map tree.

Modifier—This icon indicate modifiers and space warp bindings. Branches below a modifier contain the modifiers sub-objects and parameters.

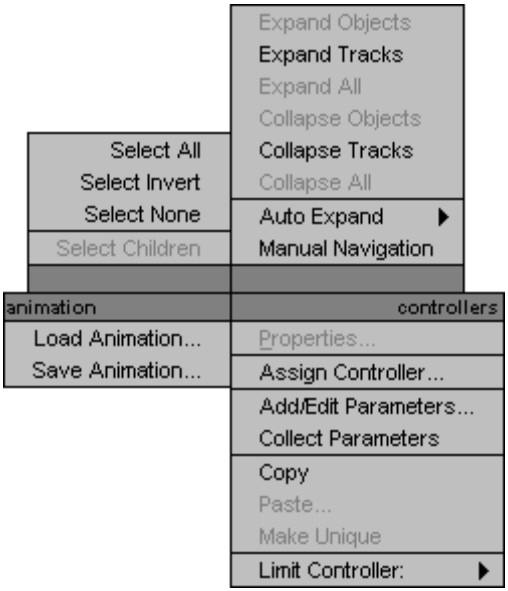
Controller Window Right-Click menu

Right-click any item in the Controller window to display the *right-click menu (page 2–516)*. Use Alt + right-click to display an alternate navigation menu.

Hierarchy Right-Click Menu

Main toolbar > Curve Editor (Open) > Select an item in the Track View controller window and right-click > Hierarchy right-click menu

The Hierarchy right-click menu contains quick-access commands for assigning, copying and pasting controllers, accessing controller properties, as well as tools for navigating, expanding and collapsing the Hierarchy list of the Controller window.

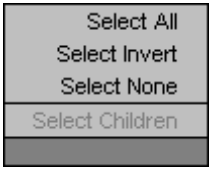


The Hierarchy right-click menu contains quick-access commands for expanding and collapsing the hierarchy.

Tip: Use Alt +right-click with quad menu tools to expand and collapse selected tracks individually.

Tip: The Hierarchy right-click menu displays commands related to the highlighted item while the Alt +right-click menu displays commands related to the item currently located under your mouse cursor.

Interface



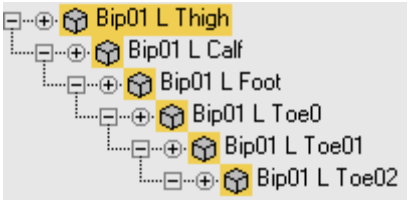
Select All—Selects all tracks that are visible in the Hierarchy list. Collapsed items are not selected.

Select Invert—Inverts the current Hierarchy list selection.

Tip: Inverting an empty selection is the same as selecting all.

Select None—Deselects all visible object tracks in the Hierarchy list. Does not apply to selected objects in the scene (object icons remain highlighted).

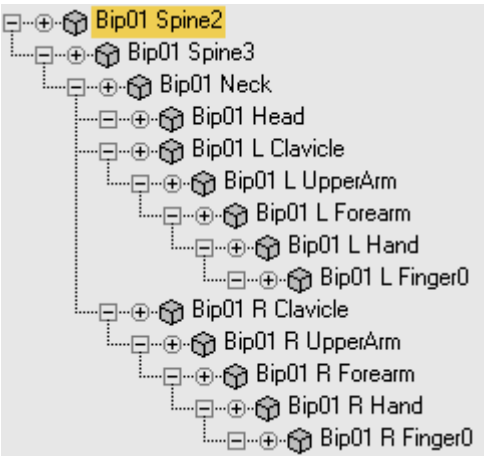
Select Children—Selects all objects descending from the selection by highlighting their icon in the Hierarchy list. Collapsed children are also selected.



Tip: You can also select an object's children by double-clicking its icon.

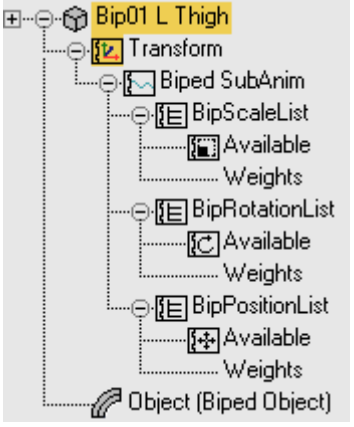


Expand Objects—Expands only the object branch for all descendants of the selected object. The expanded branches are not selected.



Note: If you are using *Auto Expand (page 2–527)* with the Children option turned on, all objects branches are automatically expanded.

Expand Tracks—Expands all branches of the selected item.



Expand All—Expands all branches for all descendants of the selected object.

Collapse Objects—Collapses only the object branch for all descendants of the selected object.

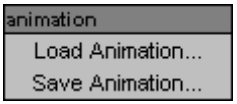
Collapse Tracks—Collapses all branches of the selected item.

Collapse All—Collapses all branches for all descendants of the selected object.

Auto Expand—Expands the Hierarchy list automatically based on submenu selection choices. Submenu choices are: Selected Objects Only, Transforms, XYZ Components, Limits, Keyable, Animated, Base Objects, Modifiers, Materials, and Children.

Manual Navigation—Turns off Auto Expand. Allows you to manually decide when you will collapse and what you will expand. A small minus button in a circle to the left of an item lets you collapse it. This button disappears when Auto Expand is on and Manual Navigation is off.

Tip: To expand or collapse individual tracks, use Alt+right-click.

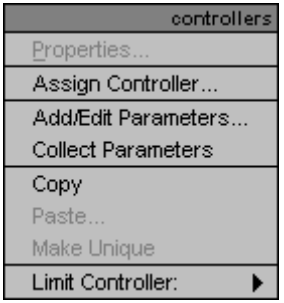


Load Animation—Loads XML animation data from files to the highlighted track or tracks. For more information, see *Saving and Loading Animation (page 3–472)*.

If any of the highlighted tracks has child tracks, animation is loaded to those tracks as well.

Save Animation—Saves XML animation data from the highlighted track or tracks to disk files. For more information, see *Saving and Loading Animation (page 3–472)*.

If any of the highlighted tracks has child tracks, animation is saved from those tracks as well.



Properties—Displays the property dialog of a controller, if available. Not all controller use this dialog, and it is unavailable in these cases.

Assign Controller—Displays the Assign Controller dialog, which offers a list of the available controller for that selection.

Copy—Makes a copy of the controller held in the Track View clipboard.

Paste—Pastes the copied controller to another object or track. Copies can be pasted as instances or as unique copies.

Make Unique—Changes an instanced controller into a unique one. Changes made to instanced controller are reflected in all versions of the controller, unique controllers can be individually edited without affecting anything else.

Alt + Right-click Menu

These commands can also be found on the right-click quad menu described above.

Expand Objects—Expands only the object branch for all descendants of the selected object. The expanded branches are not selected.

Note: If you are using *Auto Expand (page 2–527)* with the Children option turned on, all objects branches are automatically expanded.

Expand Tracks—Expands all branches of the selected item.

Expand All—Expands all branches for all descendants of the selected object.

Collapse Objects—Collapses only the object branch for all descendants of the selected object.

Collapse Tracks—Collapses all branches of the selected item.

Collapse All—Collapses all branches for all descendants of the selected object.

Select All—Selects all tracks that are visible in the Hierarchy list. Collapsed items are not selected.

Select Invert—Inverts the current Hierarchy list selection.

Tip: Inverting an empty selection is the same as selecting all.

Select Children—Selects all objects descending from the selection by highlighting their icon in the Hierarchy list. Collapsed children are also selected.

Tip: You can also select an object's children by double-clicking its icon.

Select None—Deselects all visible object tracks in the Hierarchy list. Does not apply to selected objects in the scene (object icons remain highlighted).

Properties—See *Properties (page 2–519)*.

Properties (Track View Hierarchy)

Main toolbar > Curve Editor (Open) > Select an item in the Track View controller window > Hierarchy Right-click menu > Properties
Graph Editors menu > Track View - Curve Editor > Track View toolbar > Controller > Properties
Right-click an animated object in the viewport. > Curve Editor > Right-click the Key window. > Properties
Motion panel > Assign Controller rollout > Right-click a controller track. > Properties

Properties displays the Properties dialog for the selected item in the Track View. For example, this option displays the Key Info dialog for selected keys, or the *Sound Options dialog (page 2–520)* if you select an item in the Sound Track.

Some controllers, such as Noise, use the Properties dialog as their primary interface for animating parameters. Other controllers do not use the Properties dialog at all, and in those cases it is unavailable.

Procedures

To display properties for animation controllers:

1. Select the controller track in the Controller Window.
2. On the Controller menu choose Properties.

   Or you can right-click and select Properties from the Track View quad menu.

To display the properties dialog for a Bezier key:

The Bezier Controller is the default position controller when an object is animated.

1. In the Track View Key window, select an object's position key.
2. Right-click the key.

   The controller dialog is displayed. The Key Tangency choices are available from the dialog.

   Tip: You can also change tangency handles using the tools in the Key Tangents: Track View toolbar.

To display the Sound Options dialog for adding a sound track:

1. In the Track View Hierarchy, select the Sound item.
2. Right-click and choose Properties.

   The Sound Options dialog is displayed.

   You can use this dialog to assign a sound file to the animation, which can be displayed in Track View or the track bar.

To display the Sound Options dialog for adding a sound track:

1. In the Track View Hierarchy, select the Sound item.
2.  On the Tools toolbar, click Properties, or right-click the Sound label and choose Properties.

   The Sound Options dialog is displayed.

Sound Options Dialog

Main toolbar > Curve Editor - Open > Select an item in the Track View Hierarchy and right-click. > Hierarchy right-click menu > Properties

Graph Editors menu > Track View > Open Track View > Right-click a Sound track in the Track View Hierarchy. > Properties > Sound Options dialog

The Sound track in the Track View Controller window contains two sound options. Use the Metronome to generate a series of beeps, or load a sound file to play back.

Sound is one way to provide the timing for your animation. You can scrub the time slider to locate a particular sound, and coordinate your visuals around it.

You can display this dialog by right-clicking the Sound item in the Track View controller window, or by right-clicking a sound track in the Key window.

Procedure

To add a sound track:

1. In the Track View Controller window, right-click a sound item and choose Properties.

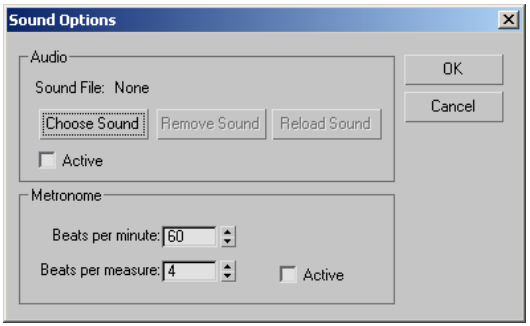
   The Sound Options dialog is displayed.

2. Click Choose Sound, select a sound file, and then click OK.

   A Wave Form branch appears in the Track View Hierarchy. A waveform appears in the Track View Key window.

3. Turn on Real Time in the Time Configuration dialog to hear the sound track when you play your animation.

Interface



Audio group

Choose Sound—Displays a file selector where you choose a sound file.

Remove Sound—Deletes the waveform from Track View.

Reload Sound—Reloads the last loaded file. This allows the sound file to be reloaded when it has been changed by another program.

Active—Plays the waveform during animation play back.

Metronome group

The Metronome track produces a regular beat using two tones. You specify three settings in the Sound Options dialog to control the Metronome:

Beats per minute—Sets the frequency of beats. The beats display in the Metronome tracks as black vertical bars. Default=60.

Beats per measure—Sets which beat is emphasized with a higher pitch tone. Emphasized beats display in the Metronome track as black pluses. For example, a setting of 4 generates a higher pitched tone every fourth beat. Default=4.

Active—Plays the beats during animation playback.

Track View Menus

Track View > Menu bar

A menu bar appears at the top of Track View in both Curve Editor and Dope Sheet modes and the expanded track bar layout. The Track View menu bar is contextual; it changes slightly between Curve Editor and Dope Sheet modes.

The commands available on the Track View menus can also be accessed on the Curve Editor and Dope Sheet toolbars. Certain tools, however, appear only on the toolbars, and do not appear in the menus.

Interface

Modes—Lets you choose between Curve Editor and Dope Sheet. See *Modes Menu (page 2–521)*.

Controller—Assigns, copies, and pastes controllers, and makes them unique. Here you also add looping. See *Controller Menu (page 2–521)*.

Tracks—Adds Note and Visibility tracks. See *Tracks Menu (page 2–524)*.

Keys—Adds, removes, slides, and scales keys. Also includes soft selection, align to cursor, and snap frame. See *Keys Menu (page 2–524)*.

Curves—(Curve Editor only) Applies or removes Ease and Multiplier curves. See *Curves Menu (page 2–525)*.

Time—(Dope Sheet only) Adjusts time. See *Time Menu (page 2–526)*.

Options—Controls the behavior (auto-expansion, etc.) of the Hierarchy list window. Also contains controls that improve performance. See *Options Menu (page 2–526)*.

Display—Affects curve, icon, and tangent display. See *Display Menu (page 2–530)*.

View—Contains Pan and Zoom commands. See *View Menu (page 2–532)*.

Utilities—Randomizes or creates out-of-range keys. Also selects keys by time and current value editor. See *Utilities Menu (page 2–533)*.

Modes Menu

Track View > Modes menu

Graph Editors menu > Track View - Dope Sheet > Modes menu

Main toolbar > Curve Editor (Open) > Modes

The Modes menu lets you switch between the Curve Editor and Dope Sheet when working in Track View.

Curve Editor—Displays and allows for editing of animation function curves.

Dope Sheet—Displays animation as a spreadsheet of keys available for editing.

Controller Menu

Graph Editors menu > Track View – Curve Editor > Controller menu

Graph Editors menu > Track View – Dope Sheet > Controller menu

Main toolbar > Curve Editor (Open) > Track View > Controller menu

The Controller menu gives you tools for working with controllers that are assigned to objects or tracks when in the Curve Editor or Dope Sheet.

Assign—Lets you select tracks and then assign a controller to that selection. See *Assign Controller (page 2–546)*.

Delete Controller—Lets you delete certain controllers, that can't otherwise be replaced (Visibility tracks, Image Motion Blur Multiplier, Object Motion Blur, On/Off). See *Delete Controller (page 2–549)*.

Keyable—Toggles the ability to be keyed of selected controller tracks. Use this with the Show Keyable button on the toolbar to see whether a track is keyable.

Copy—Puts a copy of the selected controller track into the Track View buffer. See *Copy Controller (page 2–544)*.

Paste—Copies the controller track in the Track View buffer onto the selected track of another object or objects. You have the option of pasting as a copy or as an instance. See *Paste Controller (page 2–545)*.

Collapse Controller— Converts procedural animation tracks to Bezier, Euler, Linear, or TCB keyframe controller tracks. Can also be used to convert any controller to these type of controllers. Allows for key reduction by using a Samples parameter. See *Collapse Controller (page 2–522)*

Enable Anim Layer—Assigns a Layer controller to each highlighted track in the hierarchy window. See *Animation Layers (page 2–326)*.

Note: You must first set keys on the desired tracks before they can be layer-enabled.

Ignore Animation Range—Ignores the animation range for the selected controller track. When set, the track plays independently of its range and its background changes color. See *Ignore Animation Range (page 2–549)*.

Respect Animation Range—Respects the animation range for the selected controller track. When set, the track plays only within its range. See *Respect Animation Range (page 2–550)*.

Make Unique—Lets you turn an instanced controller into a unique controller. If a controller is instanced, making changes to it will affect where ever it is copied. If the controller is unique, then changes to it will not affect anything else. See *Make Controller Unique (page 2–550)*.

Out of Range Types—Lets you extend the animation beyond the existing keyframes. Used primarily to loop and otherwise cycle your animation without having to copy keys. See *Parameter Curve Out-of-Range Types (page 2–551)*.

Properties—Displays the Properties dialog, which gives access to the key interpolation types. Different controller types will offer individual options here. For example a Position XYZ controller will offer Fast, Slow, Linear, Smooth, Step, Bezier and Auto Tangent as key options, while a TCB controller will not show any of those controls. For some controllers this is the primary gateway to the animation parameters. See *Properties (Track View Key Window) (page 2–560)*.

Collapse Controller

Track View > Track View menu > Controllers > Collapse Controllers

The Collapse Controller tool allows you to convert procedural animation tracks into tracks with keyframes. A procedural track is an expression track, wired parameter, noise controller, or even something like a list controller.

You can also use this tool to convert any controller track to a simple keyframed track. You can convert rotation tracks to either Euler or TCB, and collapse Position and Scale tracks to either Bezier or Linear controllers.

You can collapse an entire track, or collapse a range of frames within a track by setting the Start and End frames. A Samples parameter lets you define how many frames will be placed between keys.

Collapse Controller can be used to collapse any number of selected tracks at once.

This functionality is similar to the Collapse button found in the Collapse Transform group of the Motion panel > *Trajectories rollout (page 2–301)*.

Warning: If you are using a Path constraint with the Follow option, using Collapse Controller to collapse the position and rotation tracks will not preserve rotation properly, because additional rotation is applied by the Path constraint. If you want to collapse this kind of track, use Collapse in the Collapse Transform group of the Motion panel > Trajectories rollout instead.

Procedure

To collapse a procedural controller track:

1. In a viewport, select an object with a procedural controller, then right-click and choose Curve Editor from the quad menu.

   The Track View – Curve Editor is displayed.

2. In the Curve Editor's Controllers window, select the track you want to collapse. Then choose Controller > Collapse Controller from the Track View menu.

   The Collapse Controller dialog appears.

3. Change the Samples to specify the number of frames between keys.

   The default value of 1.0 creates a key for every other frame, which might be more than you want to work with.

4. Select the type of controller that you want the keys to use after collapsing.

   Position and Scale tracks collapse to either Bezier or Linear. Rotation tracks collapse to either Euler or TCB.

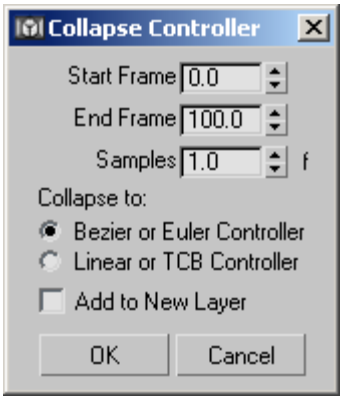
   Note: TCB controllers will not display function curves.

5. Click OK.

   The Key window displays the curves for the collapsed animation. The track bar displays the keys, as does the Track View – Dope Sheet.

   Tip: If you want to collapse the controller into a weighted list controller, turn on Add To New Layer. The original controller will be stored on a layer with a weight of 0.0. It will have no effect, but will be available for future use, if need be.

Interface



Start Frame—Sets the first frame of the range of animation to be collapsed. Default=The first frame of the original controller's active range.

End Frame—Sets the last frame of the range of animation to be collapsed. Default=The last frame of the original controller's active range.

Warning: If you choose a range of frames smaller than the original range, the animation within the range will be collapsed, and the remaining motion *will be discarded* unless Add to New Layer is on.

Samples—Sets the granularity of the collapse operation. This value defines the number of frames between keys. Default=1.0.

Collapse to—Sets the type of controller the procedural track will become:

- **Bezier or Euler Controller—** When this is chosen, Position, Scale, and scalar parameter tracks will

be assigned Bezier controllers. Rotation tracks will be assigned Euler controllers.

- **Linear or TCB Controller—**When this is chosen, Position, Scale, and scalar parameter tracks will be assigned Linear controllers. Rotation tracks will be assigned TCB controllers.

  Note: TCB controllers will not display function curves.

Add to New Layer—When this is turned on, the collapsed controller is added as a new layer in a weighted *list controller (page 2–342)*. Default=off.

This option is useful because it keeps the old controller on another layer with a weight of 0.0. This lets you restore the original animation easily if you need to.

Tracks Menu

Graph Editors menu > Track View - Curve Editor > Tracks menu

Graph Editors menu > Track View - Dope Sheet > Tracks menu

Main toolbar > Curve Editor (Open) > Tracks menu

There are two types of special tracks that can be accessed from the Tracks menu.

Note Track—Lets you add or remove note tracks to your scene. Note tracks can be used for variety of purposes to add information to any track in Track View. See *Add Note Track (page 2–552)*.

Visibility Track—Lets you add or remove visibility tracks to objects your scene. You can also keyframe visibility by changing the visibility parameter in the Object Properties dialog when Auto Key is on. See *Add Visibility Track (page 2–556)*.

Keys Menu

Track View > Keys menu

The commands on the Keys menu give you control over keys when working in either the Curve Editor or the Dope Sheet.

Add Keys—Adds keys in the Curve Editor or Dope Sheet. See *Add Keys (Dope Sheet) (page 2–560)*.

Reduce Keys—Reduces the amount of keys on curves, or in edit keys mode in Dope Sheet. See *Reduce Keys (page 2–572)*.

Move—Moves keys either vertically (in value) or horizontally (in time). See *Move Keys (Dope Sheet) (page 2–558)* or *Move Keys (Curve Editor) (page 2–579)*.

Slide—Moves keys, and adjacent keys slide away to make room for the move. See *Slide Keys (page 2–558)*.

Scale Values—Proportionally increase or decrease the key values (in space, not time). Use in combination with the Scale Value Origin Slider. See *Scale Values (page 2–581)*.

Scale Keys - Time—Proportionally increase or decrease the key values (in space, not time). Use in combination with the Scale Value Origin Slider. See *Scale Keys - Time (page 2–559)*.

Use Soft Select—When turned on, affects keys adjacent to a selection set of keys according to a falloff threshold. Works on Curves and Dope Sheet edit keys mode. See *Use Soft Select (page 2–525)*.

Soft Select Settings—Use Soft Select Settings to display the dialog that lets you control the falloff range of the soft selection. Select some keys, then watch the Track View Key window while you adjust the range value. See *Soft Selection Settings (page 2–525)*.

Note: By default, the Soft Selection dialog appears as a toolbar on the bottom of the Track View dialog.

Align to Cursor—Proportionally increases or decreases the key values (in space, not time). Use in combination with the Scale Value Origin Slider. See *Align to Cursor (page 2–556)*.

Snap Frames—When this is on, keys always snap to frames. When off, you can move keys to sub-frame positions. See *Snap Frames (page 2–554)*.

Use Soft Select

Track View > Keys menu > Use Soft Select

When on, Use Soft Select affects keys adjacent to a selection set of keys according to a falloff threshold. This applies to both Curve Editor keys and Dope Sheet Edit Keys modes.

To adjust the Soft Selection range and falloff, choose Soft Select Settings on the Keys menu.

Tip: Keys are "soft-selected" across time only (horizontally).

Procedure

To soft select keys in Track View:

1. In either Dope Sheet > Edit Keys mode or in the Curve Editor, select a single key in the middle of the animation curve or graph.
2. From the Keys menu, choose Use Soft Select.
3. From the Keys menu, choose Soft Select Settings

   The Soft Selection Settings dialog appears.
4. Change the Range and observe the soft selection display in the keys window.

See also

Soft Selection Settings (page 2–525)

Soft Selection Settings

Track View menu bar > Keys > Use Soft Select

Track View menu bar > Utilities > Track View Utilities > Soft Selection Settings Manager

Displays a dialog with controls to adjust the range and falloff of the soft selection of keys in the Dope Sheet and Curve Editor.

Interface



Soft—Toggles usage of the soft selection. When this is off, only the selected keys are used, the soft selection is not.

Range—Determines the range of keys affected by the soft selection.

Falloff—Determines the distribution of the strength of the soft selection over the range of keys.

Curves Menu

Graphic Editors menu > Track View - Curve Editor > Curves menu

Main toolbar > Curve Editor (Open) > Curves menu

The Curves menu is only available when working in Track View's Curve Editor mode. The tools on this menu facilitate curve adjustments.

Apply – Ease Curve—Applies curves to selected tracks that allow you to effect the timing of the animation. See *Apply Ease Curve / Apply Multiplier Curve (page 2–584)*.

Apply – Multiplier Curve—Applies curves to selected tracks that allow you to effect the strength of the animation. See *Apply Ease Curve / Apply Multiplier Curve (page 2–584)*.

Remove—Removes Ease and Multiplier Curves. See *Remove Ease/Multiplier Curve (page 2–585)*.

On/Off—Turns Ease and Multiplier Curves on or off. See *On/Off (Curves) (page 2–585)*.

Ease Curve Out-of-Range Types—Applies Eases to Parameter Out of Range keys. See *Ease Curve Out-of-Range Types (page 2–585)*.

Multiplier Curve Out-of-Range Types—Applies multiplier curves to Parameter Out of Range keys. See *Multiplier Curve Out-of-Range Types (page 2–586)*.

Time Menu

Graph Editors menu > Track View - Dope Sheet > Time menu

Use tools on the Time menu to edit, adjust, or reverse time. The Time menu is available only when Track View is in Dope Sheet mode. The commands on the Time menu are available only when Edit Keys is on. If you are using Edit Ranges, these commands will not be accessible. If you are using Edit Keys, once you select a time range, the remainder of the time menu commands will become available for use.

Select—Select a time range.

Insert—Add blank periods of time into a selected range.

Cut—Remove a time selection.

Copy—Copy a time selection. Includes any keys within the time selection

Paste—Duplicates a copied or cut selection.

Reverse—Rearranges the order of keys within a time range, flips time from back to start.

Options Menu

Track View > Options menu

The Options menu contains a series of toggles and switches that control how items are handled in the Track View window.

Interactive Update (page 2–526)

Sync Cursor Time (page 2–527)

Auto Expand (page 2–527)

Auto Select (page 2–528)

Auto Scroll (page 2–528)

Modify Subtree (page 2–528) (Dope Sheet only)

Modify Child Keys (page 2–529) (Dope Sheet only)

Manual Navigation (page 2–529)

Interactive Update

Track View > Options menu > Interactive Update

Interactive Update controls whether editing keys in Track View updates the viewports in real time. Also, turning off Interactive Update allows faster animation playback in some cases.

Interactive Update—When on, editing keys in Track View updates the viewports with any changes that affect the current frame while your mouse button is down. When off, the viewports are updated only after you release the mouse button. Default=Off.

For example, if you're at frame 10 and there are position keys at frame 0, 20, 30, and 50, then editing either a key at frame 0 or 20 with Interactive Update on potentially lets you see the changes in the viewport as you drag the key. However, editing a key at frame 50 is unlikely to cause a position

change at the current frame, so you won't see any difference with Interactive Update on or off.

To view the results of editing keys anywhere in Track View in real time, use *Sync Cursor Time (page 2–527)* instead. Turning on Sync Cursor Time automatically turns on Interactive Update as well, so you can see all changes interactively.

Keep this turned off when you are working in big files and moving large numbers of keys. Default=Off.

In some circumstances, simply playing back an animation or dragging the time slider with Track View open, without editing keys, can cause performance issues. If you experience this, and Interactive Update is on, turn it off to improve playback performance.

Sync Cursor Time

Track View > Options menu > Sync Cursor Time

Sync Cursor Time—Snaps the time slider to the cursor position. When on, clicking the cursor in the Track View window moves the time slider to that spot in time. When you release the mouse button, the time slider jumps back to its previous position.

When off, the time slider does not jump to the cursor position; you need to click the time slider and move it by hand. Default=Off.

Turning on Sync Cursor Time also turns on *Interactive Update (page 2–526)* so you can edit keys anywhere in Track View by dragging them and see your results right away. The time slider jumps to the frame at which the key resides and the viewport displays your changes in real time.

Auto Expand

Track View > Options menu > Auto Expand submenu

Auto Expand determines the behavior of the controller window display based on choices made from a submenu. To turn Auto Expand off with a single click, choose Manual Navigation from the Options menu: The Auto Expand settings are then disregarded.

When you are working on a specific animation task, turn off the unnecessary options to focus the controller window on what you need to see.

Note: The default auto-navigation setting for the Dope Sheet editor auto-expands only to the node track for the currently selected object. This reduces the number of tracks whose keys need to be displayed and also helps enforce the top-down workflow for which the Dope Sheet editor is designed.

Interface

Selected Objects Only—When this is on, the controller window displays the tracks for highlighted objects only. Default=On.

Transforms—Expands the Hierarchy list to display the highlighted object's Transform track. Default=On.

XYZ Components—Expands the highlighted object's Transform track to display individual XYZ components contained in each Transform controller (such as Position and Scale).

Limits—Expands the highlighted object's Limits track to display its parameters (such as Upper Limit and Lower Smoothing). To be used in conjunction with the *Limit Controller (page 2–335)*. Default=On.

Keyable—Expands the highlighted object's Hierarchy list to display keyable tracks. You must

also enable Transform or XYZ Components to see results. Also expands the Environment and Global Shadow Parameters Hierarchy lists.

Animated—Expands the highlighted object's Hierarchy list to display animated tracks.

Base Objects—Expands the highlighted object's base object track to display its parameters (such as Height/Width/Length).

Modifiers—Expands the highlighted object's modifier track to display modifiers applied to that object.

Materials—Expands the highlighted object's material track to display materials parameters.

Children—Expands the Hierarchy list to display all children starting from the highlighted object.

Auto Select

Track View > Options menu > Auto Select

Provides options to determine which types of tracks are selected when a Track View window is opened, or node selection changes. Options include Animated, with submenu choices of Position, Rotation and Scale.

When this is turned on, animated curves are automatically selected when opening the controller window, using the submenu choices as well.

Additional filtering of the controller window can be accomplished by using Filters, such as combining only selected tracks with only animated tracks.

Auto Scroll

Track View > Options menu > Auto Scroll submenu

Provides options to control the automatic scrolling of the controller window in Dope Sheet and Curve Editor. When these are chosen, the choice is displayed at the top of the controller window.

Options include Selected and Objects.

Interface

Selected—When this is on, the controller window automatically scrolls to move the viewport selection to the top of the controller window.

Objects—When this is turned on, the controller window automatically scrolls to show all the objects in the scene in the controller window.

Modify Subtree 

Track View - Dope Sheet > Dope Sheet toolbar > Modify Subtree

Track View - Dope Sheet > Options menu > Modify Subtree Keys

Modify Subtree lets you move and scale on all sub-tracks for any/all tracks on a node. It has two modes: Edit Keys and Edit Ranges.

Note: *Modify Child Keys (page 2–529)* works similarly, but for child nodes of the current object. Modify Subtree lets you edit timing for a subset of tracks for an object, while Modify Child Keys lets you edit the timing of an entire linked structure, group, or character.

Procedure

To drag the ranges and linked descendants of an object:

1.  In the Dope Sheet Editor, click Edit Ranges.

    Modify Subtree is on by default.

2. Drag the World range bar or the Objects range bar.

   With Modify Subtree on, a range bar is displayed in the Objects branch. The Objects branch is the default parent of all named objects in the scene.

   Dragging a parent Objects range bar with Modify Subtree on affects all tracks subordinate to the object and all tracks of its linked descendants.

   Dragging the World range bar with Modify Subtree on affects all tracks in the scene, including Sounds, Materials, and all tracks of all objects.

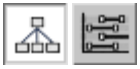
3.  Turn Modify Subtree off.

   The World range bar and the Objects range bar are no longer available. In this mode you can affect the range of an individual object or track without affecting its descendants.

Interface

 **Modify Subtree (Edit Keys)—**Edits (cuts, pastes, or moves) keys in the parent track. Anything you do to the keys in the parental track affects the child keys as well.

Note: Adding keys only affects the current track.

 **Modify Subtree (Edit Range)—**Affects the tracks of an object and all of its descendent objects. When you edit the range of a parent object, the child objects are also affected.

Modify Child Keys 

Track View - Dope Sheet > Options menu > Modify Child Keys

Track View - Dope Sheet > Track View toolbar > Modify Child Keys button

Provides the ability to turn changes on and off down the hierarchy when working in Dope Sheet mode. Changes that you have made to a parent object can be added to the children by clicking this toolbar button. Similarly, if you have made changes with Modify Subtree on, and you want to remove the changes from the children, clicking Modify Child Keys will remove those changes from the children that you previously made to the parent.

This tool is primarily designed for use in Dope Sheet Edit Ranges mode.

Modify Subtree allows you to move, scale, edit time, and so on, on all sub-tracks for any or all tracks on a node. Modify Child Keys does the same thing, but also extends to child nodes as well. Modify Subtree lets you edit timing for an object of subset of tracks on an object, while modify child keys lets you edit the timing of an entire linked structure, group, or character.

Manual Navigation

Main toolbar > Curve Editor (Open) > Options menu > Manual Navigation

Graph Editor > Track View - Curve Editor > Right-click the controller window. > Manual Navigation.

Manual Navigation turns off the Auto Scroll features of the controller window and allows you

to select which tracks will display on an individual basis.

The default behavior of the Controller window is to automatically display the selected object animated tracks, and to hide them when the object is not selected in the viewport. Turning on Manual Navigation changes this behavior so that deselecting the object will not make the curves disappear from view.

When Manual Navigation is engaged, buttons appear in the Controller window next to the tracks, that allow you to expand or collapse individual containers such as objects or materials.

Tip: Use Alt+right-click to quickly access tools for expanding and collapsing selected tracks.

To collapse individual components in the Controller window:

1. On the Graph Editors menu, choose Track View - Curve Editor.
2. Right-click the controller window and turn on Manual Navigation.
3. Click any item label in the controller window.

   An small minus sign within a circle appears to the left of the entry.
4. Click the minus sign in the circle.

   The tracks collapse. A plus sign within a circle appears.

   When Manual Navigation is turned off, tracks expand based on Auto-Expand choices made on the Settings > Auto Expand menu.

Display Menu

Track View > Curve Editor > Display menu

Display is only available when Track View is in Curve Editor mode. The toggles in the Display menu let you adjust and customize how the items in the Curve Editor appear.

Selected Key Stats—Displays the statistics for the selected keys in the function curve window. This is useful because you only see statistics for the keys you are working on. See *Show Selected Key Statistics (page 2–595)*

All Tangents—Displays all the tangent handles for all keys in the Curve Editor.

Custom Icons—Changes the display of the icons in the Hierarchy list from 2D to 3D shaded.

Keyable Icons—Displays a keyable icon that lets you define if a track is keyable or not. Red icons are keyable tracks, Black icons are not. Click the icon to toggle between these or use Keyable on the Controllers menu. See *Keyable Icons (page 2–531)*.

Hide Non-Selected Curves— When this is on, if you deselect the object in the viewport, the function curve will also disappear from view in the Curve Editor. Default is on. See *Hide/Show Non-Selected Curves (page 2–532)*.

Show Non-Selected Curves—Turn this on so you can deselect an object in the viewport and still see its curves. Default is off. See *Hide/Show Non-Selected Curves (page 2–532)*

Freeze Non-Selected Curves—Displays nonselected curves, but doesn't allow you to edit them. Only available when Show Non-Selected Curves is on. Default is on. See *Freeze Non-Selected Curves (page 2–587)*.

Filters—provides controls to filter the display in Curve Editor. There are a wide range of options to show, hide and display. See *Filters (page 2–541)*.

Show All Tangents 

Track View > Curve Editor toolbar > Show All Tangents

Curve Editor menu bar > Display > All Tangents

This feature shows all tangent handles on any displayed curves.

You can make changes to multiple tangent handles simultaneously by using Show All Tangents in combination with Lock Tangents.

See also

Show Tangents (page 2–582)

Keyable Icons 

Main toolbar > Curve Editor (Open) > Display menu > Keyable Icons

Menu bar > Graph Editors > Track View - Curve Editor > Track View toolbar > Show Keyable Icons

Keyable icons provide a method to tag a track so that it can receive keyframes or be exempted from keyframes. An icon is displayed in the Controller window next to the name of track to indicate whether or not the track is keyable. You can toggle the state of the icon to define the keyable property.

You can also use the Track View Controller menu > Keyable command to make tracks keyable in a single operation.

You can define keyboard shortcuts for making tracks keyable. By combining the use of keyable icons with key filters, you can use Set Key animation mode to add keyframes to just the tracks you want to work with, and avoid keyframing other tracks.



Red key icon means a track is keyable.

When the keyable icons are visible, click the red icon to turn off the track.



Black key in a circle indicates a track is not keyable.

Warning: When you are working with Set Key animation and you have used key filters to select object parameters or materials, *all* of the parameters will be keyed unless you turn off the track's keyable property.

Tip: You can assign a keyboard shortcut for making tracks keyable, by choosing Keyable Property Toggle in the Track View group, on the Keyboard panel of the Customize User Interface dialog.

Procedures

To make an individual track keyable:

1. Select the object in the viewport.
2. Right-click and choose Curve Editor.

   The Function Curve Editor opens, with the selected object tracks visible.
3.  On the Track View toolbar, click Show Keyable Icons.
4. Click the red keyable icons for the tracks for which you want to prevent animation.

   The tracks are changed, now marked with a black key in a circle. These tracks will not receive keyframes. Only the tracks marked with red keyable icons will.

   Note: Keyable tracks work with both Auto Key and Set Key animation modes.

To make multiple tracks keyable:

1. In Track View, with the object selected, hold down the Ctrl key and click each track to

create a selection set of tracks for which you want to prevent animation. This can be for one or multiple objects

Tip: You can hold down the Shift key to select a group of sequential tracks at once. Alternately you can hold down the Alt key to select all tracks at the same level as a given track at once.

Note: If you select just a parent track such as Position, Controller > Keyable will toggle all of its sub-tracks, even if they're not selected.

2. From the Controller menu, choose Keyable.

   The selected tracks are now defined as keyable.

3. From the Display menu, choose Keyable Icons.

   The keyable icon appears next to the tracks. The icon appears red for the tracks that are keyable, and black for the ones that are not.

   Tip: Use the same procedure to make multiple tracks not keyable. The Controller > Keyable command toggles the keyability of the tracks.

   Note: If you have a mixed keyable situation it will toggle the state of the first track and then set all remaining tracks to match the value of the first track

Hide/Show Non-Selected Curves

Main toolbar > Curve Editor (Open) > Display > Hide Non-Selected Curves or Show Non-Selected Curves

These commands either hide or show function curves that are not selected in the controller window.

When Hide Non-Selected Curves is on (the default), the curve will disappear when another track is chosen.

When Show Non-Selected Curves is on, the curve will still be visible in the Key window when another track is chosen.

If you turn on Show Non-Selected Curves, then you can also use Freeze Non-Selected Curves. This allows you to see the other curve but not edit it inadvertently.

View Menu

Track View > View menu

These commands replicate those found on the Navigation toolbar in Track View.

Pan—Lets you move the window. See *Pan (Track View) (page 2–595)*.

Zoom—Lets you zoom in and out. See *Zoom Track View Key Window (page 2–596)*.

Zoom Region—Lets you zoom in to a rectangular area. See *Zoom Region (Track View) (page 2–597)*.

Zoom Horizontal Extents—Zooms to the active time segment. See *Zoom Horizontal Extents (page 2–595)*.

Zoom Horizontal Extents Keys—Zooms to show all keys. See *Zoom Horizontal Extents (page 2–595)*.

Zoom Value Extents—(Curve Editor only) Zooms vertically so that you can see the full height of the curve.

Zoom Values—(Curve Editor only) Zooms the contents of the keys window vertically. Drag upward to increase magnification, or downward to decrease magnification.

Utilities Menu

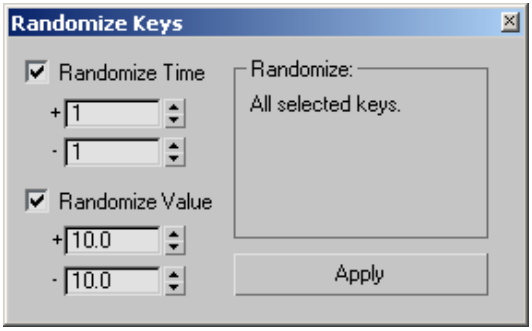
Graph Editors menu > Track View - Curve Editor > Utilities menu

Graph Editors menu > Track View - Dope Sheet > Utilities menu

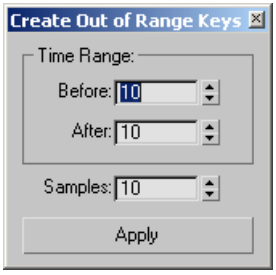
Main toolbar > Curve Editor (Open) > Utilities menu

The *Utilities menu (page 2–561)* gives you access to the Track View Utilities dialog. The dialog shows a list of tools you can use when working with keys. This menu is available in both Curve Editor and Dope Sheet modes.

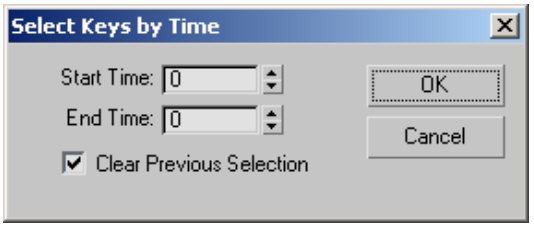
Randomize Keys—Changes the values of selected keys randomly based on range thresholds. Use this on either values or time. For details, see *Randomize Keys Utility (page 2–562)*.



Create Out of Range Keys—Creates new keys for selected tracks based on Out-of-Range Parameters. Adjust the samples value to change the granularity of the key creation. For details, see *Create Out of Range Keys Utility (page 2–562)*.



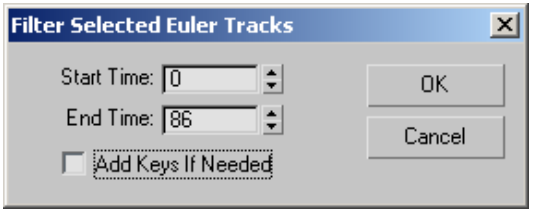
Select Keys by Time—Allows you to select the keys within a time range. Turn off Clear Previous Selection to create discontinuous selection sets. For details, see *Select Keys By Time Utility (page 2–563)*.



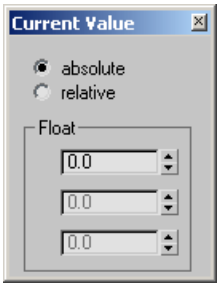
Soft Selection Settings Manager—Displays a soft selection dialog at the bottom of the Track View window. For details, see *Soft Selection Settings (page 2–525)*.



Euler Filter—Corrects rotation anomalies by filtering Euler rotation tracks. For details, see *Euler Filter (page 2–564)*.



Current Value Editor—Provides transform type-in capability from within the Track View modes. Allows you to choose between absolute and relative value editing. The name of the controller appears above the axis choices. For details, see *Current Value Editor (page 2–565)*.

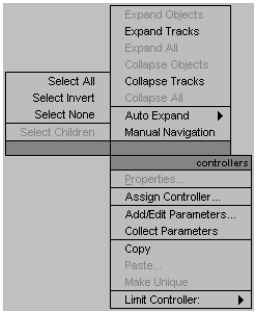


This utility starts a floating Current Value window that works for either Dope Sheet – Edit keys mode or Curve Editor. It doesn't work for Edit Ranges. Not for use with object parameters, the current value editor is intended for use primarily with Transform controllers.

Track View Quad Menu

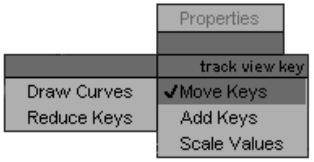
Right-clicking the Key window or Hierarchy list brings up a quad menu that offers instant access to tools. The menu is context-sensitive, so its contents vary, depending on what is highlighted. For example, when the Hierarchy list is active, you can assign copy and paste controllers, and open properties which you can animate. You can set the Auto Expand and Manual Navigation controls to change the behavior of the Controller window display.

The commands from the alternative quad menu (Alt+right-click) are integrated in the *right-click quad menu (page 2–516)*. The Select, Expand, and Collapse commands are therefore accessible along with the existing quad menu commands, improving workflow efficiency.



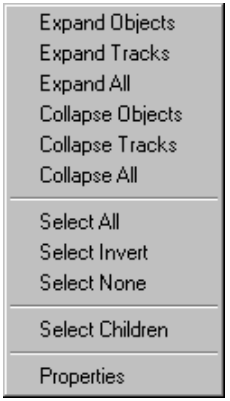
Controller window quad menu

When the keys window is active you can draw curves, add keys, move keys and scale values. You can also reduce keys.



Key window quad menu

In the Controller window, if you hold down Alt+right-click you can display an alternate quad menu, which is also integrated with the main quad menu. This menu was available in earlier versions of the program as a default right-click menu in Track View.



Alt+right-click for alternate quad menu for Controller window

You can add additional commands to these quad menus. You can customize the Track View quad menu the same as any other quad menu.

Procedure

To customize the Track view quad menu:

1. On the Customize menu choose Customize User Interface.
2. On the Customize User Interface dialog, click the Quad tab.
3. From the Quad pull-down menu, choose Track View Quad, or Track View Key Quad.
4. Change the Group to Track View, and then drag items from the table of actions on the left to the window in the lower right.

Track View Toolbars

Curve Editor Toolbars

Main toolbar > Curve Editor (Open) > Curve Editor toolbar

Track View - Dope Sheet > Modes menu > Curve Editor > Curve Editor toolbar

The Curve Editor toolbars contain tools that are also available from the Track View menu bar. The toolbars can be floated, docked, and rearranged as you like. These include the Keys toolbar, The Key Tangents toolbar, and the Curves toolbar, all of which are displayed by default. You can also display other toolbars for quick access to additional tools. To show the hidden toolbars, right-click an empty area between toolbars, choose Show Toolbars, and then pick the toolbar you want to display.

Key Tools Toolbar

 **Filter—**Use this to determine what is displayed in the Controller window and the Key window. See *Filters (page 2–541)*.

 **Move Keys—**Moves keys freely both horizontally and vertically on the function curve graph. See *Move Keys (page 2–558)*.

 **Move Keys—Horizontal—**Moves keys only horizontally on the function curve graph. See *Move Keys (page 2–558)*.

 **Move Keys — Vertical—**Moves keys constrained vertically on the function curve graph. See *Move Keys (page 2–558)*.

 **Slide Keys—**Use Slide Keys in the Curve Editor to move a group of keys and slide adjacent keys away as you move. See *Slide Keys (page 2–558)*.

 **Scale Keys—**Use Scale Keys to compress or expand the amount of time between keyframes. Works both in Curve Editor and Dope Sheet modes. See *Scale Keys - Time (page 2–559)*.

 **Scale Values—**Proportionally increases or decreases the values of the keys, rather than moving the keys in time. See *Scale Values (page 2–581)*.

 **Add Keys—**Creates keys on existing curves on the function curve graph or Dope Sheet. See *Add Keys (Dope Sheet) (page 2–560)*.

 **Draw Curves—**Use this to draw new curves, or revise existing ones by sketching directly on the function curve graph. See *Draw Curves (page 2–582)*.

 **Reduce Keys—**Use this to reduce the amount of keys in a track. See *Reduce Keys (page 2–572)*.

Key Tangency Toolbar

 **Set Tangents to Auto—**Located on the Key Tangents Track View toolbar, select keys then click this button to set tangents to Auto Tangent automatically. A flyout also lets you set in tangents and out tangents individually to Auto. Selecting the handles of an Auto tangent changes them to custom, and makes them available for editing.

 **Set Tangents to Custom—**Sets keys to custom tangents. Select the key then click this button to make the key handles available for editing. Set in and out tangents individually using the flyout. Use the Shift key to break continuity when using handles.

 **Set Tangents to Fast—**Sets key tangency to fast in, fast out, or both in and out, depending on your choice from the flyout.

 **Set Tangents to Slow—**Sets key tangency to slow in, slow out, or both in and out, depending on your choice from the flyout.

 **Set Tangents to Step—**Sets key tangency to step in, step out, or both in and out, depending on your choice from the flyout. Use step to freeze motion from one key to the next.

 **Set Tangents to Linear—**Sets key tangency to linear in, linear out, or both in and out, depending on your choice from the flyout.

 **Set Tangents to Smooth—**Sets key tangency to smooth. Use this to even up discontinuous motion.

Curves Toolbar

 **Lock Selection—**Locks the key selection. Once you have created a selection, turn this on and you can't inadvertently select something else. See *Lock Selection (page 2–555)*.

 **Snap Frames—**Restricts key movement to frames. Keys that are moved will always snap to frames when this is on. When this is off, you can move a key so it falls between frames and becomes a sub-frame key. Default=on. See *Snap Frames (page 2–554)*.

 **Parameter Out-of-Range Curves—**Use this to repeat keyframed motion beyond the range of the keys. Includes options for Loop, Ping Pong, Cycle, or Repeat relatively, as well as constant and linear. If you use Parameter Out-of-Range types, you can later create keys using Track View > Utilities > Create Out-of-Range Keys. See *Parameter Curve Out-of-Range Types (page 2–551)* and *Create Out of Range Keys Utility (page 2–562)*.

 **Show Keyable Icons—**Displays an icon that defines a track as keyable or not. Use this to set keys only on the tracks you want to keyframe. Turning off a track in Track View also restricts the movement in the viewport. Red keys indicate keyable tracks, black keys are not keyable. See *Keyable Icons (page 2–531)*.


Show Tangents—Hides or displays tangent handles on the curves. Use this to hide the handles on individual curves. See *Show Tangents (page 2–582)*.


Show All Tangents—Hides or displays all tangent handles on the curves. Use this to hide the handles quickly when many keys are selected. See *Show All Tangents (page 2–531)*.


Lock Tangents—Locks the selection of multiple tangent handles, so you can then manipulate several handles at once. When Lock Tangents is off, you can only manipulate one key tangency at a time. See *Lock Tangents (page 2–583)*.

Biped Toolbar

These tools let you choose which animation curves to display in the Curve Editor. You can toggle between the position and rotation curves, as well as toggle the separate curves representing the X, Y, and Z axes of the current biped selection.


Show Biped Position Curves—Displays the position curves for the animated biped selection. Also chooses *Pos Curve* from the Curve Type drop-down list of the *Animation Workbench Toolbar (page 2–1014)*.

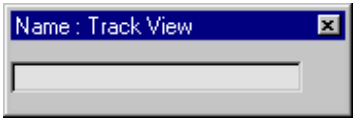

Show Biped Rotation Curves—Displays the rotation curves for the animated biped selection. Also chooses *Rot Curve* from the Curve Type drop-down list of the *Animation Workbench Toolbar (page 2–1014)*. Default=on.


Show Biped X Curves—Toggles the X axis of the current animation or position curves. Also toggles the X button of the *Animation Workbench Toolbar (page 2–1014)*. Default=on.


Show Biped Y Curves—Toggles the Y axis of the current animation or position curves. Also toggles the Y button of the *Animation Workbench Toolbar (page 2–1014)*. Default=on.


Show Biped Z Curves—Toggles the Z axis of the current animation or position curves. Also toggles the Z button of the *Animation Workbench Toolbar (page 2–1014)*. Default=on.

Name: Track View Toolbar



Name—Track View lets you name the Track View. By entering a name in this field, you create a named Track View that you can recall using the Graph Editors menu > Saved Track View submenu. Also available in Dope Sheet mode.

Track Selection Toolbar


Zoom Selected Object—Puts the currently selected object at the top of the Hierarchy list. Also available in Dope Sheet mode. See *Zoom Selected Object (page 2–588)* for details.

Note: If multiple objects are selected, the Hierarchy list zoom is based on an alphabetical sorting.



Select By Name—Highlights filtered tracks in the Controller window based on the current field value. Also available in Dope Sheet mode. See *Select By Name (page 2–589)* for details.

Tip: You can use commas (,) to include multiple names in your selection.



Track Set List—Allows you to assign a name to the current track selection (referred to as a track set), and then later reselect those tracks by choosing the respective track set from the list. Also available in Dope Sheet mode. See *Track Set List (page 2–590)* for details.

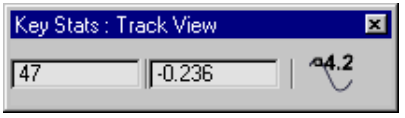
You can edit track set lists via the Edit Track Set button, which opens the *Track Sets Editor dialog (page 2–591)*.

Navigation Toolbar



Navigation—Track View provides tools to pan and zoom the Key window. You can zoom to the extents of the time or the values. Also available in Dope Sheet mode. See *Status Bar/View Controls (page 2–588)*.

Key Stats Toolbar



Key Stats—Track View provides tools for the display and type-in transform of key values. Also available in Dope Sheet mode.

Dope Sheet Toolbars

Main toolbar > Curve Editor (Open) > Modes menu > Dope Sheet > Dope Sheet toolbar

These are the default toolbars that are visible when you open the Dope Sheet.

Keys: Dope Sheet Toolbar


Edit Keys—Displays a Dope Sheet Editor mode that shows the keys as boxes on a graph. Use this to mode to insert, cut, and paste time. See *Edit Keys (page 2–554)*.


Edit Ranges—Displays a Dope Sheet Editor mode that shows the keyed tracks as range bars. See *Edit Ranges (page 2–573)*.


Filter—Use this to determine what is displayed in the Controller window and the Dope Sheet - Key window. See *Filters (page 2–541)*.


Slide Keys—Use Slide Keys in the Dope Sheet to move a group of keys and slide adjacent keys away as you move. Only slides keys on the same controller track. See *Slide Keys (page 2–558)*.


Add Keys—Creates keys on existing tracks on the Dope Sheet grid. Combine this tool with Current Value editor to adjust the key values numerically. See *Add Keys (Dope Sheet) (page 2–560)*.


Scale Keys—Use Scale Keys to compress or expand the amount of time between keyframes. Works both in Curve Editor and Dope Sheet modes. Uses the time slider as an origin point to or from which to scale. See *Scale Keys - Time (page 2–559)*.

Time: Dope Sheet Toolbar


Select Time—Lets you select a time range. Time selections include any keys that might be included within the time range. Use Insert Time, then Select Time to choose the time range. See *Select Time (page 2–566)*.


Delete Time—Removes selected time from the selected tracks. Cannot be applied globally to shorten the time segment. This removes keys but leaves "blank" frames behind. See *Delete Time (page 2–567)*.


Reverse Time—Reverses the keys on selected tracks within a selected time segment. See *Reverse Time (page 2–569)*.


Scale Time—Scales the keys within a selected time segment for selected tracks. See *Scale Time (page 2–570)*


Insert Time—Allows you to insert a range of frames as a time insertion. Existing keys slide out of the way to make room for inserted time. Once you have made a time selection with Insert Time, you can then use all the other time tools. See *Insert Time (page 2–570)*.


Cut Time—Deletes time selections from the selected tracks. See *Cut Time (page 2–567)*.


Copy Time—Duplicates the selected time selection so it will be available for pasting. See *Copy Time (page 2–568)*.


Paste Time—Adds cut or copied time selections into the selected tracks. See *Paste Time (page 2–568)*.

Display: Dope Sheet Toolbar


Lock Selection—Locks the key selection. Once you have created a selection, turn this on so that you cannot inadvertently select something else. See *Lock Selection (page 2–555)*.

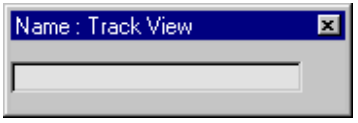

Snap Frames—Restricts key movement to frames. Keys that are moved will always snap to frames when this is on. When this is off you can move a key so it falls between frames and become a sub-frame key. Default=on. See *Snap Frames (page 2–554)*.


Show Keyable Icons—Shows an icon that defines a track as keyable or not. Use this to set keys only on the tracks you desire to keyframe. Turning off a track in Track View will also restrict the movement in the viewport. Red keys show keyable tracks, black keys are not keyable. See *Keyable Icons (page 2–531)*.


Modify Subtree—When this is on, allows key manipulations to parent tracks to affect the tracks down the hierarchy. On by default in Dope Sheet mode. See *Modify Subtree (page 2–528)*.


Modify Child Keys—If you modify the parent without Modify Subtree on, click Modify Child Keys to apply the change to the child keys. Similarly, if you modify the parent with Modify Subtree on, Modify Child Keys toggles those changes off. See *Modify Child Keys (page 2–529)*.

Name: Track View Toolbar



Name—Lets you name the Track View. By entering a name in this field, you create a named Track View window that you can later recall using the Graph Editors menu > Saved Track View submenu. Also available in Curve Editor mode.

Track Selection Toolbar



Zoom Selected Object—Puts the currently selected object at the top of the Hierarchy list. Also available in Curve Editor mode. See *Zoom Selected Object (page 2–588)* for details.

Note: If multiple objects are selected, the Hierarchy list zoom is based on an alphabetical sorting.



Select By Name—Highlights filtered tracks in the Controller window based on the current field value. Also available in Curve Editor mode. See *Select By Name (page 2–589)* for details.

Tip: You can use commas (,) to include multiple names in your selection.



Track Set List—Allows you to assign a name to the set of currently highlighted tracks (referred to as a track set), and then later re-select those tracks by choosing the respective track set from the list. Also available in Curve Editor mode. See *Track Set List (page 2–590)* for details.

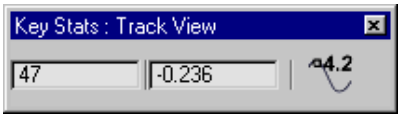
You can edit track set lists via the Edit Track Set button, which opens the *Track Sets Editor dialog (page 2–591)*.

Navigation Toolbar



Navigation—Track View provides tools to pan and zoom the Key window. You can zoom to the extents of the time or the values. Also available in Curve Editor mode. See *Status Bar/View Controls (page 2–588)*.

Key Stats Toolbar



Key Stats—Track View provides tools for the display and type-in transform of key values. Also available in Curve Editor mode.

Controller Toolbar

Main toolbar > Curve Editor (Open) > Right-click the starting area of a toolbar (containing two vertical bars). > Show Toolbars > Controllers : Animation Editor

This toolbar contains the basic tools you need to work with controllers in Track View. This toolbar is hidden as a default. Right-click the Track View toolbar and choose Show Toolbars > Controllers: Track View to display these tools.

All of these commands are always available via the Track View Controller menu.



Filters—Click to display the *Filters dialog (page 2–542)*, which you can use to control what Track View displays in the Controller and Key windows.



Copy Controller—Use this to copy a controller and its animated tracks from a selected object.



Paste Controller—Pasts the copied or cut controller onto a new object or selection of object tracks.

 **Assign Controller—**Use this to assign a new controller to an object. All objects have a default controller assigned; use this to change the default controller to a different one. Select the controller track in the Controller window, then use Assign Controller to select a new one.

 **Delete Controller—** Delete a controller from an object. The controller will be replaced with a default controller.

 **Make Controller Unique—**Changes an instanced controller to a unique controller. This lets you make changes to the controller without affecting any other object tracks.

Procedure

To display the Controller toolbar:

1. Right-click the blank area on the Track View toolbar, to the right of the Modify Child Keys button.
2. From the right-click menu, choose Show Toolbars.
3. Choose Controllers: Track View from the list.

Ranges: Dope Sheet Toolbar

Main toolbar > Curve Editor (Open) > Right-click the starting area of a toolbar (containing two vertical bars). > Choose Show Toolbars > Ranges : Dope Sheet.

This toolbar, hidden by default, provides tools you use when working with ranges.

 **Edit Ranges—**Changes Dope Sheet to display range bars rather than tracks of keys. See *Edit Ranges (page 2–573)*.

 **Position Ranges—**Adjusts the relationship between a range bar and its keys. Turning this on will allow you to see the keys while moving the range bar.

 **Recouple Ranges—**Resizes the range bar to fit the first and last key of the selected track.

Extras: Dope Sheet Toolbar

Main toolbar > Curve Editor (Open) > Mode > Dope Sheet > Right-click the blank area of the toolbar to the right of Modify Child Keys. > Show toolbar > Extras: Dope Sheet toolbar

Provides extra tools for use in Dope Sheet mode.

 **Exclude Left End Point—**Prevents a key at the first frame of a selected block of time from being copied to the clipboard.

 **Exclude Right End Point—**Prevents a key at the last frame of a selected block of time from being copied to the clipboard.

Filters 

Main toolbar > Curve Editor (Open) > Track View toolbar > Filters

Graph Editors menu > Track View - Curve Editor/Track View - Dope Sheet > Track View toolbar > Filters

Graph Editors menu > Track View - Curve Editor/Track View - Dope Sheet > Track View menu > Display > Filters

Filters lets you determine which categories of items appear in Track View. When you click the Filters button, Track View's *Filters dialog (page 2–542)* is displayed.

You can also right-click the Filters button for quick selection of items.

Tip: You can set up a default filter configuration. Open a single Track View, set the filters the way you want them to come up, and close the Track View. Save the (empty) scene as *maxstart.max*. This scene file is automatically loaded when you start the program.

Procedures

To position a selected object at the top of the Track View Hierarchy:

1. In a viewport, right-click a selected object.
2. From the right-click menu, choose Curve Editor or Dope Sheet.

   Track View opens with the selected object at the top of the Track View Hierarchy.

To display the animated transform tracks for an object in Track View:

1. Select an object, then open a Track View window.
2. Right-click Filters.
3. From the right-click menu, choose Animated Tracks Only.
4. Right-click Filters again.
5. From the right-click menu, choose Selected Objects Only.

Interface

See *Filters Dialog (page 2–542)*.

Filters Dialog (Track View)

Main toolbar > Curve Editor (Open) > Toolbar > Filters

Graph Editors menu > Track View - Curve Editor > Display menu > Filters

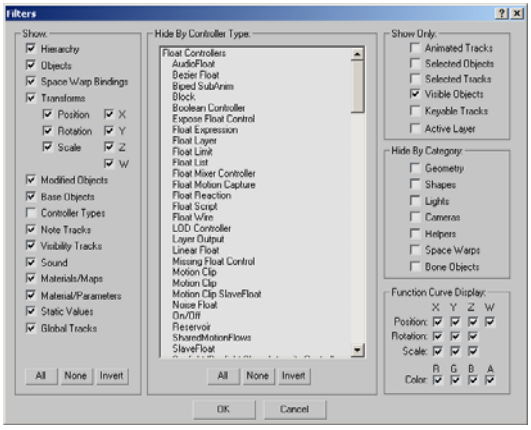
You use the Filters dialog to choose what to display in Track View. For example, you can limit the view to animated tracks only, or tracks for selected objects. This dialog also controls function curve display and transform display for Position, Rotation, Scale, and X, Y, and Z axes individually.

Procedure

To choose filter options:

-  On the Track View toolbar, click Filters.
- Choose any of the filter options from the dialog.

Interface



Show group

The Show group has options to display any of the following in the Hierarchy list window:

- Hierarchy
- Objects
- Space Warp Bindings

- Transforms (Position, Rotation, Scale, X/Y/Z/W axes in any combination)
- Modified Objects
- Base Objects
- Controller Types (off by default)
- Note Tracks
- Visibility Tracks
- Sound
- Materials/Maps
- Material/Parameters
- Static Values
- Global Tracks: non-object tracks

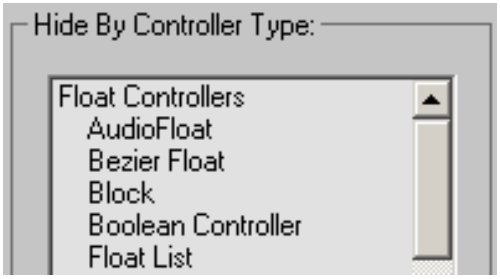


All—Sets all Show check boxes to on.

None—Sets all Show check boxes to off.

Invert—Reverses the state of all Show check boxes.

Hide by Controller Type group

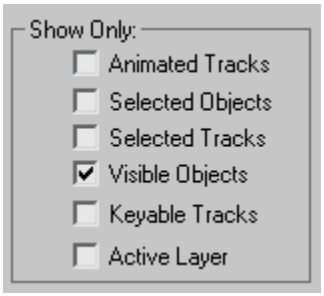


Hide By Controller Type contains a window displaying a list of all controller types in 3ds Max. Choose one or more controller types to prevent them from appearing in the Hierarchy list. You can use the standard multiple-selection methods of Ctrl+click, Shift+click, or drag.

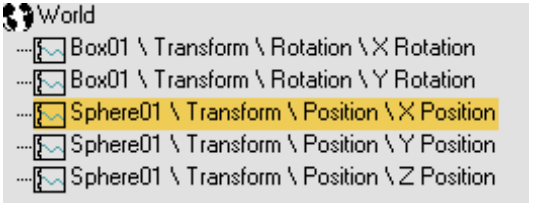
All/None/Invert—Selects either all items in the list, none of the items in the list, or inverts the current selection.

Note: When you hide a controller, its subcontrollers (if any) are hidden as well. For example, if you hide a PRS Transform controller, its Position, Rotation, and Scale controllers are also hidden.

Show Only group



Animated Tracks—Displays only tracks that contain animation. Only the animated tracks are shown, each with its full hierarchy, whether or not the *Hierarchy switch* is on.



Selected Objects—Displays only items for objects selected in the scene. As you select objects in the scene, the display in the Hierarchy list changes to show the current selection. The display of sound and materials branches are not affected by this filter.

Selected Tracks—Displays only items that you select prior to setting this filter. All of the displayed items are left-justified in the Hierarchy list regardless of their level in the hierarchy.

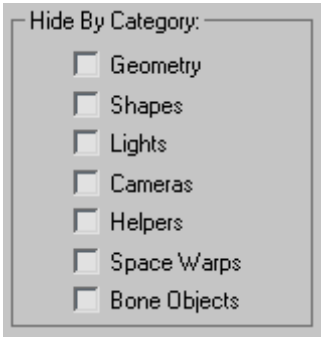
Visible Objects—Determines whether objects that are hidden in the viewports will appear in the Track View Hierarchy. This does not consider objects with tracks, but only objects that have been hidden. Default=on.

Keyable Tracks—Displays only tracks that can receive keys. This property is toggled using the Keyable property available on the Track View Controller menu, or by clicking the keyable icon displayed using Show Keyable Icon. Default=off.

Active Layer— For each object listed in the hierarchy window with *animation layers (page 2–326)* enabled, displays only the active layer, along with all nested controllers.

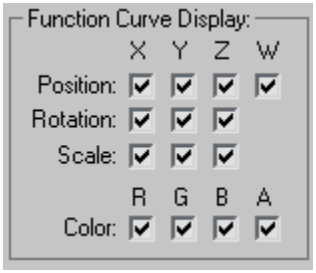
Tip: Newly available in *Customize User Interface (page 3–792)* are actions and icons for all six Show Only options. Find them in the Track View group; their names all start with "Filter".

Hide by Category group



Contains a list of check boxes that let you hide tracks based on categories similar to those found in the Display panel. Turning on one of these hides the entire category type and any subcomponents.

Function Curve Display group



Check boxes in this group, when active, specify which transforms are suppressed, for which axis, and which RGB color values are suppressed. This is only used for controllers such as the Bezier Position controller that displays all three axis with one track selected.

Note: The "W" and "A" Filters are for use with the Point4 controller. The W is simply a fourth filter for position, and the A is for use with the floating-point RGBA controller, which is based on the Point4 controller.

Copy Controller

Main toolbar > Curve Editor (Open) > Highlight a controller track in the controller window. > Controller menu > Copy

Main toolbar > Curve Editor (Open) > Highlight a controller track in the controller window. > Right-click any blank space on the Track View toolbar > Show Toolbars > Controllers: Track View > Copy Controller

Copy Controller copies the selected item to the Track View clipboard.

You can use Copy and Paste to copy object and modifier tracks, controllers, objects, and containers, depending on what you are working on.

The Controller tools are also available as buttons on the Controllers: Track View toolbar. This toolbar is hidden by default; to unhide right-click the toolbar and choose Show Toolbars > Controllers: Track View.

Not all tracks can be copied and pasted. The basic restrictions for using Copy and Paste in Track View are:

- Only a single selected item can be copied.
- A copied item can only be pasted into another item of the same type.

- A copied item can be pasted into a selection of multiple items only if all of the items are of the same type.

Procedures

To copy and paste a modifier:

1. In the Track View Hierarchy, highlight a modifier track.
2. From the Controller menu, choose Copy.
3. Click the object track for an object in the Track View Hierarchy.
4. From the Controllers menu, choose Paste.

   The Paste dialog is displayed.
5. Set the options in the Paste dialog, and click OK.

   The pasted modifier is inserted above the object track you selected. Note that it does not replace any of the tracks in the target object.

To copy and paste a controller:

1. In the Track View Hierarchy, highlight a controller track.
2.  On the Track View Controllers menu, click Copy Controller.
3. Highlight the controller track for an object in the Track View Hierarchy.

   Make sure you choose the same type of transform that you copied.
4. From the Controllers menu, choose Paste.

   The Paste dialog is displayed.
5. Set the options in the Paste dialog, and click OK.

   The controller is pasted onto the selected object.

Paste Controller 

Main toolbar > Curve Editor (Open) > Highlight the track you want to paste to. > Controller menu > Paste.

Main toolbar > Curve Editor (Open) > Highlight the controller track in the controller window. > Right-click any blank space on the Track View toolbar > Show Toolbars > Controllers: Track View toolbar > Paste Controller

Keyboard > Ctrl+V

The Paste controller pastes the contents of the Copy Controller buffer. You can use the Copy and Paste buttons to copy tracks, controllers, objects, and containers.

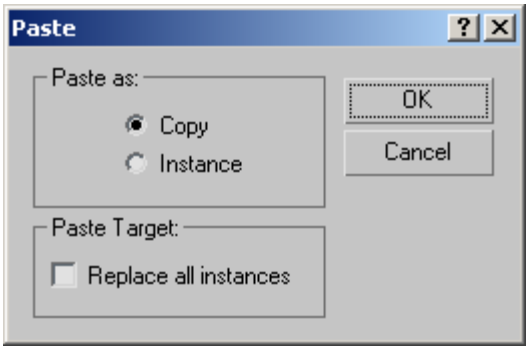
All the Controller tools can also be found on the Controllers: Track View toolbar, which is hidden by default. Right-click the toolbar, then choose Show Toolbars > Controllers: Track View to display the icons for Assign, Copy, Paste. Delete Controller and Make Controller Unique.

Tip: If you need a gizmo to follow a dummy object, use Copy Controller and Paste Controller to copy the animation from the dummy to the gizmo.

Procedures

See *Copy Controller (page 2–544)*.

Interface



Paste Controller—Pastes an item from the Track View clipboard.

The Paste dialog contains the following options:

Copy/Instance—Determines whether the cloned item is a simple copy or an instance of the original.

Replace All Instances—When you paste to any track that contains an instanced object, replaces all instances of the object with the object on the clipboard. If off, only the object in the selected track will be replaced. Any instanced objects will remain as they are.

Assign Controller 

Main toolbar > Curve Editor (Open) > Select the controller track in the controller window. > Controller menu > Assign

Main toolbar > Curve Editor (Open) > Select the controller track in the controller window. > Right-click the controller track and choose Assign Controller.

Keyboard > **C**

Use Assign Controller to assign animation controllers to any animatable parameter or track in Track View.

Animation controllers provide powerful tools for animating all the objects and materials in a scene. For example, rather than keyframing the position of an object in your scene, the object can follow a spline using the Path constraint, react to any animated parameter using a Reaction controller, or move to the beat of music using the Audio controller. You can combine controllers with a List Controller. You can drive a single vertex or control point on a complex object by a variety of controllers.

You can also assign controllers using the right-click menu in the Track View Controller window. Or you can assign controllers in the Motion panel, in Schematic View, and by using the Animation menu.

Constraints and Controllers

Technically, there is no difference between a controller and a constraint. A constraint is simply a controller that requires the use of a second object. For example, a Path constraint is a controller that requires a spline object for a path.

Special-Case Controllers

Special-case controllers are not assigned manually with the Assign Controller command. They are applied automatically during certain procedures.

Barycentric Morph Controller is applied by selecting an object and clicking Command panel > Geometry > Compound Objects > Morph.

A Master Point Controller is assigned when animating vertices, control points, or vectors in the sub-object mode of an Editable Mesh, Editable Spline, Editable Patch, NURBS surface, or FFD modifier.

Slave Controllers can be assigned manually, but is also automatically applied to selected tracks when a Block controller is created in Track View Global Tracks. A Slave controller transfers key data to a Block controller. Slave controllers are described in the Block controller topic.

Available Animation Controllers and Constraints

Attachment Constraint (page 2–393): A Position controller that attaches an object's position to a face on another object. The target object doesn't have to be a mesh, but must be able to convert itself to a mesh.

Audio Controller (page 2–309): Animates a parameter or object using a sound file.

Bezier Controller (page 2–310): Smoothes function curves between transform keys, giving the effect of continuous, natural movement. The Bezier

controller is the default controller for position and scale.

Barycentric Morph Controller (page 2–309): Applied during the creation of a morph object. Each morph key represents a series of weights for all morph targets.

Block Controller (page 2–313): Group tracks from multiple objects into Blocks. Blocks can be copied, pasted, scaled and saved. Use Blocks to quickly re-create a complex motion anywhere in an animation.

Color RGB Controller (page 2–317) (Point 3 XYZ Controller): Creates separate tracks for red, green, and blue.

Euler XYZ Rotation Controller (page 2–318): Combines separate, single-value float controllers to specify an angle of rotation about each of the X, Y, and Z axes. You can specify the order the axes are rotated.

Expression Controller (page 2–320): Mathematical expressions control object parameters such as length, width, and height, or transform and modifier values, such as position coordinates.

IK Controller (page 2–440): Assigned to a hierarchy through the use of an IK Solver. Can be automatically assigned to bones at the time of creation by using Assign to Children.

Linear Controller (page 2–341): Makes the function curves between each of the keyframes into straight lines.

List Controller (page 2–342): Combines controllers.

Link Constraint (page 2–403): Animates the transfer of hierarchical links from one object to another. You can, for example, have the same object passed from object to object.

LookAt Constraint (page 2–406): Forces that object to constantly look at another. Formerly a Transform controller, this is now a rotation controller.

Master Point Controller (page 2–346): Automatically assigned when animating vertices, control points, or vectors in sub-object mode. This controller simplifies managing potentially hundreds of tracks in Track View.

Motion Capture Controller (page 2–347): Controls and records an object's motion or parameter using an external device. Supported devices are mouse, keyboard, MIDI device, and joystick.

Noise Controller (page 2–353): Generates random values, which you can see as peaks and valleys in the function curve.

On/Off Controller (page 2–355): Provides on and off control.

Boolean Controller (page 2–316): An improved version of the On/Off Controller that allows for adding keys without disrupting the existing keyframing.

Path Constraint (page 2–398): Assigns a spline as a trajectory (path) for an object so that the object follows the path.

Position XYZ Controller (page 2–356): Splits the X, Y, and Z components into three separate tracks.

PRS Transform Controller (page 2–357): Creates subcontrollers for position, rotation, and scale.

Reaction Controller (page 2–358): Allows an object or parameter to react to another object or parameter. For example, you can change the scale of one object based on the Z position of another object.

Scale XYZ Controller (page 2–371): Provides separate scale tracks for the X, Y, and Z axes.

Script Controller (page 2–372): Controls objects and parameters with the 3ds Max scripting language.

Spring Controller (page 2–375): Adds secondary dynamics effects to any point or object position.

Smooth Rotation Controller (page 2–374): Creates smooth rotations.

Surface Constraint (page 2–396): Positions an object along the surface of another object.

TCB Controller (page 2–377): Provides Tension, Continuity, and Bias controls for the function curves between keys.

Transform Script Controller (page 2–379): Provides scripted PRS control for objects.

Waveform Controller (page 2–381): A float controller that provides regular, periodic waveforms. For example, blinking lights and rhythmic object motion.

XRef Controller (page 2–383): Lets you externally reference any type of Transform controller from another scene file.

Procedure

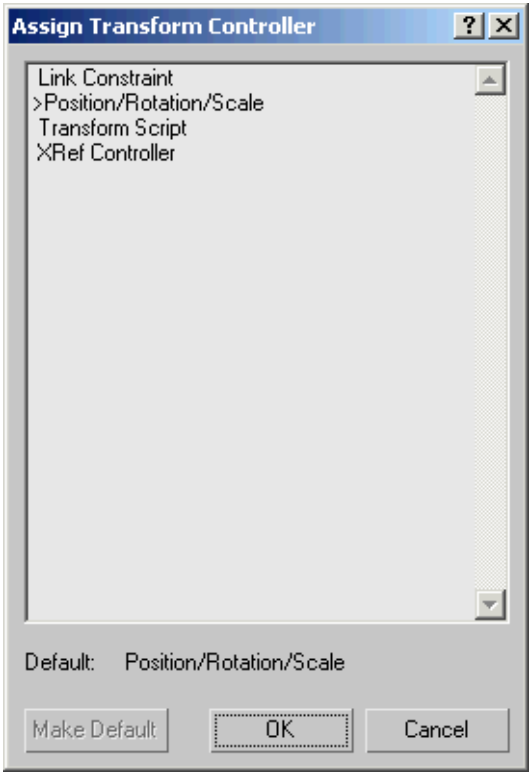
To assign an animation controller in Track View:

1. In the Track View Hierarchy, select one or more parameter items of the same type.
2. From the Controller menu, choose Assign.
3. Choose a controller type from the Assign Controller dialog.

   If a parameter has already been animated, then assigning a new controller has one of the following effects:

   - The existing animation values are recalculated to produce a similar animation with the new controller. For example, replacing Position XYZ with Bezier Position closely preserves the animation.
   - The existing animation values are discarded. For example, replacing Smooth Rotation with Noise Rotation discards the Smooth Rotation animation values.

Interface



Assign Controller—Choose a controller type from the Assign Controllers dialog. Depending on the type of track you have selected, the Choose Controller dialog displays a subset of the different types of controllers.

Delete Controller 

Main toolbar > Curve Editor (Open) > Select a deletable track in the Controller Window. > Right-click the toolbar > Show Toolbars > Controllers: Track View > Delete Controller.

Main toolbar > Open Track View > Select a deletable track in the Track View Hierarchy. > Right-click on a blank area of the Track View toolbars and choose Show Toolbars > Controllers toolbar > Delete Controller.

Some node subcontrollers can be deleted. Visibility tracks, Image Motion Blur Multiplier, Object Motion Blur, and On/Off are examples of controllers that can be deleted. Most controllers, however, are not deletable.

You can also delete a Visibility track by using Tracks > Visibility Tracks > Remove.

Procedure

To delete a deletable controller:

1. In the Controller window, select a deletable controller.
2. On the Track View Controllers menu, choose Delete Controller.

   The controller is deleted.

   Note: Most controllers cannot be deleted, only replaced. The Delete Controller menu item is available only if the controller track selected is an applicable type.

Ignore Animation Range

Track View > Select a controller track > Controller menu > Ignore Animation Range

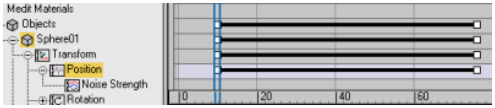
Ignore Animation Range allows a parametric, non-keyable controller track to be active throughout the entire length of the animation, independent of the track's current animation range. Choosing this option changes the track's background color to purple.

This setting is active by default on all new controller tracks (turn off the Override Parametric Controller Range By Default option in the *Animation Preferences (page 3–828)* to change it). However, controller tracks from older scenes are set to *Respect Animation Range (page 2–550)* to maintain their original behavior.

Procedure

Example: To set a controller track to ignore the animation range:

1. Create a sphere.
2. On the Motion Panel > Assign Controller rollout, highlight the sphere's Position track.
3.  Click the Assign Controller button, and then choose Noise Position from the Assign Controller dialog.
4. The Noise Controller dialog opens. Keep the current settings and close the dialog.
5. Select the sphere, and then right-click it. From the quad menu, choose Dope Sheet.
6. Expand the hierarchy items in the Controller window until you find the sphere's Position track.
7. Drag the track's range bar until it starts at frame 10.



8. Play the animation.

   Because Ignore Animation Range is on by default, the sphere moves randomly even before entering the track's animation range.

Respect Animation Range

Track View - Dope sheet > Select a Controller track > Controller menu > Ignore Animation Range

Respect Animation Range constrains a parametric, non-keyable controller track to be active only within the track's current animation range. Choosing this option changes the track's background color to gray.

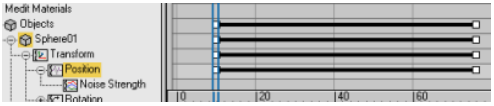
Note: While *Ignore Animation Range (page 2–549)* is the default setting (turn off the Override Parametric Controller Range By Default option in the *Animation Preferences (page 3–828)* to change it), controller tracks from older scenes are set to this setting to maintain their original behavior.

Procedure

Example: To set a controller track to respect the animation range:

1. Create a sphere.
2. On the Motion Panel > Assign Controller rollout, highlight the sphere's Position track.
3.  Click the Assign Controller button, and then choose Noise Position from the Assign Controller dialog.
4. The Noise Controller dialog opens. Keep the current settings and close the dialog.
5. Select the sphere, and then right-click it. From the quad menu, choose Dope Sheet.
6. Expand the hierarchy items in the Controller window until you find the sphere's Position track.
7. On the Controllers menu, choose Respect Animation Range.

   The background color of the track changes to gray.
8. Drag the track's range bar so it starts at frame 10.



9. Play the animation.

   Because Respect Animation Range is on, the sphere starts moving only when it is within the track's animation range.

Make Controller Unique 

Main toolbar > Curve Editor (Open) > Select an instanced Controller track > Controller menu > Make Controller Unique.

Keyboard: **U**

Make Controller Unique converts an instanced clone of a controller to a copy that's unique to the current object. Objects and modifiers can also be made unique.

The result depends on whether or not the data flow branches at the selected container:

- If the data flow does not branch anywhere below the selected container, nothing happens.
- If the data flow does branch at or below the selected container, the data flow above the branch is split from the current data flow as a unique object. The data flow from the selected container to the master object is copied and attached to the new unique object.

All the Controller tools can also be found on the Controllers: Track View toolbar which is hidden by default. Right-click any blank area of the Track View toolbar, then choose Show Toolbars > Controllers: Track View to display the buttons for Assign, Copy, Paste. Delete Controller and Make Controller Unique.

Procedure

To convert an instanced controller to a unique one:

1. Select an instanced controller.
2. On the Controllers menu choose Make Unique, or press U on the keyboard.

   Tip: If you have instanced modifiers you can make them unique by choosing the modifier in the Modifier Stack, right-clicking and choosing Make Unique.

Parameter Curve Out-of-Range Types 

Main toolbar > Curve Editor (Open) > Controller menu > Out-of-Range Types

Graph Editors > Track View - Curve Editor > Curves: Track View Toolbar > Parameter Curve Out-of-Range Types

With Parameter Curve Out-of-Range Types you can specify how an object is to behave outside the range of the keys you've defined.

Use Parameters Curve Out-of-Range Types to create loops and cycles. The idea is to create a short pattern of keys that produce an animated effect, and then loop or cycle through those keys throughout the animation.

Using the Parameter Curve Out-of-Range Types dialog, you select from four ways for repeating your animation and two ways for applying a linear value.

Tip: Use the *Create Out of Range Keys utility (page 2–562)*, found in Track View Utilities, to turn the parameter out-of-range animation into editable keyframes.

Typical Steps for Using Parameter Out-of-Range Types

1. Animate an object to loop or cycle.
2. In the Curve Editor select the track you want to loop.
3.  Click Parameter Out-of-Range Types to select how your animation behaves outside the time covered by the range bar.
4. Choose the type of out of range type you want in the dialog, then close it and play the animation. The Curve editor will display the loop or cycle with a dotted line.

   Tip: You can create keys from the out of range type by using the Track View Utilities > Create Out of Range Keys.

Procedure

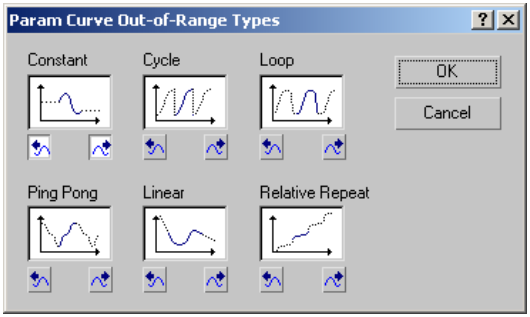
To use Parameter Out-of-Range Types to create a loop:

1. Create a Box primitive on the left side of the Front viewport.
2. Turn on Auto Key and move the time slider to frame 10.
3. Move the box to the right side of the Front viewport.
4. Open Track View, right-click Filters, and then select Animated Tracks Only.
5. Choose the Position track for the box in the Track View Hierarchy, and then click Parameter Curve Out-of-Range types.
6. In the Parameter Curve Out-of-Range Types dialog, select Ping Pong for both the in and out types and then click OK.
7. Click Play in the Viewport Controls.

   The box moves back and forth repeatedly.

   While the animation is playing, experiment by adjusting the keys in the Curve Editor, or in the Dope Sheet.

Interface



Select from the options for repeating your animation and options for applying a linear value.

Constant—Holds the value of the end key of the range for all frames. Use constant when you want no animated effect before the first key of the range or after the last key. Constant is the default out-of-range type.

Cycle—Repeats the same animation as within the range. If the first and last keys in the range have different values, the animation will show an abrupt "jump" from the last key to the first. Use Cycle when you want an animation to repeat, but do not need the ends to match.

Loop—Repeats the same animation as within the range, but interpolates between the last key and first key in the range to create a smooth loop. If the first and last key are both at the extreme ends of the range, Loop will behave exactly like Cycle. If you use Position Ranges to extend the range bar beyond the keys, the added length determines the amount of time used to interpolate between the last key and the first key. Use loop with an extended range bar to produce smoothly repeating animation.

Ping Pong—Alternates between a forward and backward repeat of the animation within the range. Use Ping Pong when you want your animation to alternate back and forth.

Linear—Projects the animation value along a line tangent to the function curve at the end of the range. Use Linear when you want the animation to enter and leave the range at a constant velocity.

Relative Repeat—Repeats the same animation as within the range but offsets each repetition by the value at the end of the range. Use Relative Repeat to create animations that build on each other as they repeat.

Add Note Track 

Main toolbar > Curve Editor (Open) > Modes > Dope Sheet > Edit Keys > Tracks menu > Note Tracks > Add.

With Add Note Track, you can insert a note track below a highlighted track in the Track View Hierarchy. You can make note keys and associated notes at specific time locations to keep track of what keys do. Do this by adding keys in Dope Sheet > Add Keys mode, then right-click a note key to see the Notes dialog. The information you enter will appear in the Key window under the keys.

Any item in the Hierarchy list can have a note track added as a branch below it. You can add a note track to an item in any of the Track View modes.



Note track for animation of door opening

Once you have created a note track for an item, use Add Keys to insert note keys. Use the Notes dialog to edit the notes.

You can have toolbar access to the Note Tracks by unhiding the Tools: Track View toolbar.

Procedures

To add a note track:

1.  On the Track View Dope Sheet toolbar, click Edit Keys.
2. Highlight one or more item labels in the Hierarchy list.
3. From the Tracks menu, choose Note Tracks > Add.

   A note track is inserted as a branch directly below each highlighted item.

To add a note key to a note track:

1. Select the Note track (in Dope Sheet Edit Keys mode).
2.  Click Add Keys.
3. Click a note track in the Keys Window to place a note at that particular frame.

To edit notes:

1. Right-click a note key to display the Notes dialog.
2. Click inside the edit box, and then start typing.

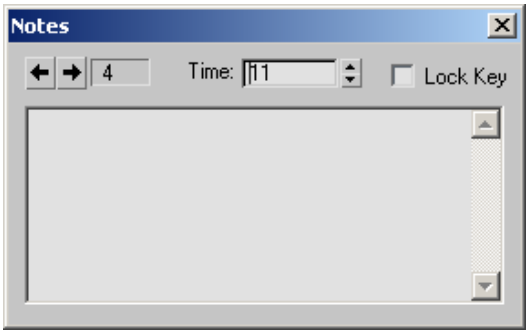
   The first line of your note appears as a label to the right of the note key. Type a one- or two-word description as the first line of your note. Close the dialog using the X button at the top right corner to enter the body of the note.

Interface

Add Note Track—Inserts a note track directly below the highlighted items in the Hierarchy list. This is available from the Tracks menu or from the Tools: Track View toolbar.

Notes dialog

Adds or edits notes on the animation in Track View. The first line of your note appears as a label to the right of the note key. You should type a short one- or two-word description as the first line of your note. Press **Enter** to start the body of the note. Press **Enter** or click within the text field to start the body of the note. Close the Notes dialog by clicking the X button at the top-right corner of the dialog.



Note Number—Indicates which note key you are working with. Click the left arrow to move back to the previous note or the right arrow to move forward to the next note.

Time—Sets the time position for the note key. Change the value to move the note to a new time.

Lock Key—Locks the key to the Time field. When set, the key ignores all operations performed with Move, Slide, and Scale. The only way to change the key's time is to use the Time field in the Notes dialog.

Remove Note Track 

Track View > Tracks menu > Note Track > Remove

Use Remove Note Track to delete a note track and its associated keys.

You can have toolbar access to Remove Note Track by unhiding the Tools: Track View toolbar. This toolbar is available only in the Curve Editor, however. The tooltip is labeled Delete Note Track.

Procedure

To delete a note track from items:

1. Highlight one or more note tracks in the Hierarchy list.
2. From the Tracks menu, choose Note Track > Remove.

Editing Keys

Topics in this section relate to editing animation keys in Track View. The topics are:

Edit Keys (page 2–554)

Snap Frames (page 2–554)

Lock Selection (page 2–555)

Align to Cursor (page 2–556)

Add Visibility Track (page 2–556)

Move Keys (Dope Sheet) (page 2–558)

Slide Keys (page 2–558)

Scale Keys - Time (page 2–559)

Add Keys (Dope Sheet) (page 2–560)

Properties (Track View Key Window) (page 2–560)

Edit Keys 

Graph Editors menu > Track View - Dope Sheet > Edit Keys

Main toolbar > Curve Editor (Open) > Modes > Dope Sheet > Edit Keys

Edit Keys displays your animation as a series of keys as boxes on a grid in the Key window. Dope Sheet Editor turns on Edit Keys by default.

Edit Keys is useful for getting a global view of your animation because it displays animation timing for all tracks. Use this mode for key and range editing when you want to view your changes in the context of the total animation.

Use Edit Keys mode to:

- Select and change one or more keys
- Drag range bars to change all animation in multiple tracks

Procedures

To turn on Edit Keys mode, do one of the following:

- If you are in Curve Editor, choose Modes > Dope Sheet.
-  If you are in Dope Sheet > Ranges, click Edit Keys on the Track View toolbar.

To delete keys in Edit Keys mode:

1. Highlight the keys in the Key window.
2. Press DELETE on the keyboard to delete the selected keys.

Snap Frames 

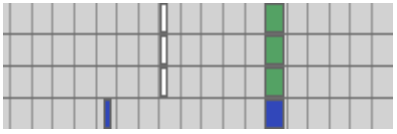
Main toolbar > Curve Editor (Open) > Keys menu > Snap Frames

Graph Editor > Track View - Dope Sheet/Track View - Curve Editor > Keys menu > Snap Frames

With Snap Frames, all key and range bar positions that are changed are forced to absolute frame increments. This includes selection sets of multiple keys. When Snap Frames is on, each key in a selection set will snap to the nearest frame when the selection is moved or scaled. Default=on.

When you use the Time Display format MM:SS:Ticks, which doesn't use frames, Snap Frames snaps to time values that match frame boundaries.

In Dope Sheet > Edit Keys mode, keys are displayed as a box grid. Keys that are snapped to frames are displayed as filled boxes. Keys that have been moved with snap frames off are displayed as narrow rectangles within the grid. When using the Curve Editor, a similar display is available in the track bar.



Subframe keys moved with snap frame off

Warning: Don't turn off Snap Frames unless you have a reason to, such as the need for more precise animation timing. If you do turn off Snap Frames and move keys in Curve Editor you will create sub-frame keys, but they won't be noticeably different in Curve Editor. In some cases, sub-frame keys can lead to animation "popping," or overly fast changes due to keys being too close together.

Procedure

To use Snap Frames:

Snap Frames is on by default, so the following procedure only applies if you have turned off snap frames previously.

1. On the Track View Dope Sheet toolbar, click Edit Keys.

   Note: This isn't an absolute necessity, but it makes it easier to see the snapping action. You can also view the snapping in the track bar in either Dope Sheet or Curve Editor mode.

2. On the Track View toolbar, turn on Snap Frames.
3. In the Track View Key window, select one or more keys.
4. Use Move Keys or Scale Keys to move keys.

   With Snap Frames on, each key in a selection set snaps to the nearest frame when the selection is moved or scaled.

Lock Selection

Track View > Curve Editor or Dope Sheet> Highlight keys > Track View toolbar > Lock Selection

Keyboard > **SPACEBAR**

Lock Selection toggles selection locking on and off. When a selection is locked you can't inadvertently deselect or select anything else. When a selection is locked you don't have to click it to move the keys, you can click anywhere in the window to more or scale the keys.

Note: To use the Spacebar keyboard shortcut, the *Keyboard Shortcut Override Toggle (page 3–872)* must be on.

Procedure

To use Lock Selection in Track View:

1. In either Dope Sheet or Curve Editor Key window, highlight one or more keys.
2. On the Track View toolbar, turn on Lock Selection, and then click and drag inside the Track View Key window.

   All the selected keys move regardless of where inside the Track View Key window you click and drag.

Align to Cursor

In either Dope Sheet or Curve Editor, select keys to align to current frame. > Keys menu > Align to Cursor

Use Align to Cursor to move selected keys to the current time. Select the keys to align using Move, Slide, or Scale.

Align to Cursor is useful for taking a group of scattered keys and moving them to the same time location.

Align to Cursor is available in Dope Sheet and Curve Editor modes.

Align Keys to Cursor ignores the state of Snap Frames, and always uses the exact time set by the time slider.

Procedure

To move all selected keys to the current frame:

1. Open either Curve Editor or Dope Sheet - Edit Keys mode.
2. Drag the time slider to the time where you want the keys aligned.
3.  Highlight one or more keys using Move Keys, Slide Keys, or Scale Keys.
4. From the Keys menu choose Align To Cursor.

   The leftmost selected key in each track is moved to the current time. Multiple selected keys on the same track maintain their relative distance from the leftmost key.

Add Visibility Track

Right-click an object in the viewport. > Curve Editor/Dope Sheet > Tracks menu > Visibility Track > Add

Graph Editors menu > Track View - Dope Sheet/Track View - Curve Editor > Click the object whose visibility you want to key. > Tracks menu > Visibility Track > Add

Add Visibility Track controls when you can see an object. Visibility tracks can only be added to objects. You can create a visibility track by either animating the Visibility parameter in the Object Properties dialog or by selecting an object in the Track View Hierarchy and clicking Tracks menu > Visibility Tracks > Add. In both cases a visibility track, using a Bezier float controller, is created in Track View. The visibility track displays as a child of the object in the Track View Hierarchy.

Note: In the Object Properties dialog, Rendering Control must be set to By Object in order to keyframe the Visibility using the spinner.

Note: Right-click over an object in the viewports and select Properties to display the Object Properties dialog.

Warning: An object without mapping coordinates that is invisible at frame 0 will not ask for UVW Map coordinates at render time. The warning will display at the frame that the object becomes visible. This can stop a render midway, so make sure invisible objects at the beginning of your animation have mapping coordinates if necessary.

On and Off Visibility

When you first assign a visibility track to an object, a Bezier float controller is automatically assigned; this allows gradual visibility. You can make an object appear or disappear suddenly by changing the interpolation of the visibility keys to Step tangency.

Variable opacity is supported in the shaded viewports. For ease of use, the object never completely disappears from the viewports.

Visibility Inheritance

An object can inherit the visibility of its parent (as determined by the parent's Visibility track in the Track View). Use the Object Properties > Inherit Visibility check box to specify the visibility inheritance of an object.

Groups and Visibility

All members of a Group inherit the visibility of the parent when a visibility controller is assigned to the parent. Transparent materials and hidden objects have no effect on this function.

Level of Detail Utility and Visibility

The *Level of Detail utility (page 1–1253)* lets you construct an object that alters its geometric complexity or level of detail based on its size in the rendered image. You create several versions of the same object each with different levels of detail; you group them as one, and then assign the Level of Detail utility. This automatically creates a special LOD controller as a Visibility track.

Procedures

To add a visibility track to objects:

1. In Track View Curve Editor or Dope Sheet Edit Keys mode, highlight the object you want to affect in the Hierarchy list on the controller window.
2. From the Track View > Tracks menu, choose Visibility Tracks > Add.
3. Select the keys then use Controller > Properties to adjust the key values and interpolation.

   A value of 0 creates an invisible object, a value of 1.0 creates a fully visible object. Change the interpolation to step for sudden visibility, rather than gradual fades.

A visibility track appears below the selected object(s). Add and edit keys in the visibility track to control the object's visibility.

To add visibility keys:

1. Once you have a visibility track assigned to the object, highlight the track.
2.  On the toolbar, click Add Keys.
3. On the Track View Key window, click in a visibility track to add a key at that time location.

Example: To create and animate a visibility track using the Object Properties dialog:

1. Create a box.
2. 

Turn on Auto Key.
3. Move the time slider to frame 10.

   This is the frame at which you'll set the box to disappear.
4. Right-click the object and choose Properties from the shortcut menu.
5. In the General tab, set Visibility to **0**, and then click OK.
6. Scrub the time slider.

   The box fades. This creates a visibility track in Track View and adds a key with a value of 0 to the track.

Move Keys (Dope Sheet)

Dope Sheet > Keys: Track View toolbar > Move Keys

Move Keys lets you reposition keys horizontally, within their own track(s), thus changing the times at which they take effect. If no keys are highlighted, you can move any key by dragging it. If multiple keys are highlighted, you can move them all the same distance by dragging one of them.

You can clone highlighted keys by moving them while holding down the Shift key.

Note: You can use any button on the Move Keys flyout to move keys in Dope Sheet mode, but you can move keys horizontally only.

Procedure

To move a selection of keys:

1. On the Keys: Dope Sheet toolbar, click Move Keys if it isn't already highlighted. It should be on by default.
2. Highlight the keys you want to move in the Key window by dragging a selection rectangle around the keys. You can also use Ctrl+click to highlight non-contiguous keys.
3. Position the mouse cursor over one of the highlighted keys, and then drag horizontally to move the keys in time.

   Tip: Press SPACEBAR to lock the key selection. Then you don't have to drag over the selected key, but anywhere in the viewport. This is useful when you have a complex selection set of keys.

Slide Keys

Track View > Highlight keys in the Key window. > Track View toolbar > Slide Keys

Use Slide Keys to move a group of keys (highlighted keys plus all the keys to one end of the animation). The direction that you drag determines which group of keys moves:

- Dragging to the right moves the highlighted keys, plus all keys to the last key of the animation, forward in time.
- Dragging to the left moves the highlighted keys, plus all keys to the first key of the animation, backward in time.

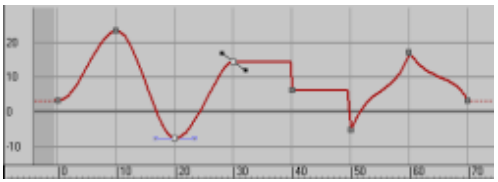
Slide Keys is a way to split the animation at the highlighted keys and spread the ends apart. Slide Keys is available in Edit Keys mode.

You can clone keys and insert them elsewhere in your animation curve while offsetting existing keys by the length of time occupied by the cloned keys by dragging while holding down Shift. Based on the length of your selection, the existing keys to the right of the selection move forward in time (that is, to the right) to allow the new key insertion, whether you drag to the right or the left.

Procedure

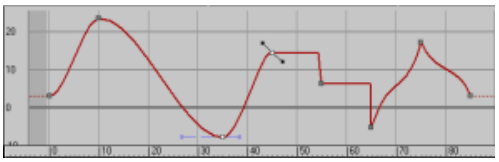
To slide a selection of keys:

1. Click Slide Keys on the Track View toolbar.
2. Highlight one or more keys.

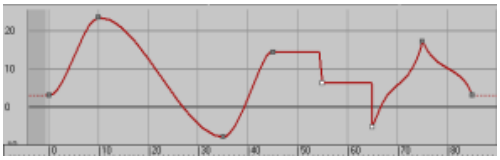


3. Drag to slide the keys to the right.

   The keys following your selection move to account for the offset produced by sliding the keys.

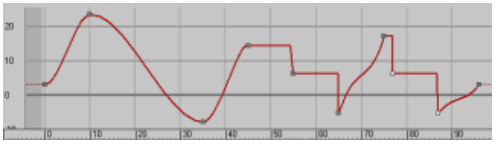


4. Highlight another group of keys.



5. Hold Shift while dragging to slide the keys to the right.

   The keys following your selection move to the right by the length of time occupied by the cloned keys.



Scale Keys - Time 

Track View > Curve Editor or Dope Sheet > Track View toolbar > Scale Keys

Track View > Curve Editor or Dope Sheet > Keys menu > Scale Keys - Time

Scale Keys - Time moves all selected keys proportionally toward or away from the current frame. Use the Scale Keys button in either mode to change the location and amount of time covered by one or more selected keys.

The scale center is defined by the current time set by the time slider. You can scale keys about any moment in time by dragging the time slider before you use Scale Keys.

Procedure

To use Scale Keys:

1. Select an animated object in the viewport, then right-click and choose Curve Editor.

   The following steps work in either Curve Editor or Dope Sheet modes.

2. Drag the time slider to the time you want to use as the scale center.

3.  Click Scale Keys or from the Keys menu choose Scale Keys - Time.

4. Highlight one or more keys.

5. Drag to scale the selected keys or press and hold Shift and drag to add scaled copies of the selected keys.

   Drag away from or toward the current time line for the following results:

   - Dragging away expands the keys from the current time. Expanding the selection increases time between the selected keys and slows that part of the animation.
   - Dragging toward shrinks the keys toward the current time. Shrinking the selection reduces time between the selected keys and accelerates that part of the animation.
   - Dragging through the current time reverses the keys and expands the keys away from the current time.

Tip: You can type-in the Scale Value using the first field in the Key Stats: Track View toolbar.

The scale percentage is displayed in the Show Selected Key Stats field as you drag the selection.

Add Keys (Dope Sheet) 

Track View > Dope Sheet > Track View toolbar > Add Keys

Track View > Dope Sheet > Keys menu > Add Keys

Add Keys inserts a *key (page 3–960)* at the point where you click a curve or a track.

Add Keys is a mode that remains active until you activate another mode. While Add Keys is on, you click in any animation track to add a key at that location in time.

Procedure

To add keys in Track View:

1.  On the Track View toolbar, click Add Keys.
2. Click an animation track to add a key.

   The location where you click sets the time of the key as measured on the time ruler.

   The value of the new key is set by one of the following conditions:

   - Keys added before the first key of a track receive the same value as the former first key.
   - Keys added between two keys receive an interpolated value based on the values of the original keys.
   - Keys added after the last key in the track receive the same value as the former last key.

   If you are unable to add keys to a track, check the following conditions:

   - Only animated tracks can accept keys. If you want to add a key to a track, such as an object parameter, you need to animate it first.
   - The animation controller must be a type that uses keys. Not all controllers use keys. Examples of controllers that do not use keys include procedural controllers such as Noise.
   - The animation controller must be a type that uses keys. Not all controllers use keys. Examples of controllers that do not use keys include Expression controllers, List controllers, and Parametric controllers.

Properties (Track View Key Window)

Track View > Highlight a single track that uses Properties. > Right-click the track. > Choose Properties on the Controllers quad.

Track View > Highlight a single track that uses Properties. > Controller menu > Properties

Properties displays a dialog to change animation values. The type of dialog displayed depends on the type of animation controller the selected track is using. The Properties function is unavailable if the selected track controller does not use properties, or the track selection set is incorrect.

- A track with animation keys display a *Key Info dialog (page 2–304)*. You can change the values in the Key Info fields to change the animation value, time, and interpolation methods of one or more selected keys.
- A track with a parametric controller, such as *Noise (page 2–353)*, displays a Properties dialog. You change the values in the Properties dialog to modify the behavior of the controller over its entire range.

  Access the Properties dialog through the controllers quad of the right-click menu, or on the Controller menu > Properties. The same dialog is also accessible through the Motion panel.

You can also display the Key Info dialog by right-clicking a key in the Key window. Parametric

controller dialogs can also be displayed by right-clicking their range bars.

The Properties button is unavailable in ambiguous cases, for example, when a key and a Parametric controller item are both selected.

Procedures

To display properties for a controller, do one of the following:

- Highlight the track and then choose Controller menu > Properties.
- Right-click the track, and then on the Controllers quad choose Properties.

To display the Key Info dialog for a key:

1. Right-click the key in the Curve Editor Key Window.

   The Key Info dialog is displayed.
2. Change key properties in the dialog.

Track View Utilities

Track View > Utilities menu > Track View Utilities

Track View Utilities displays a dialog listing plug-in Track View utilities. A typical Track View utility displays a modeless dialog (though utilities can be modal) and provides functions that can be applied across multiple selected keys and tracks. You can launch more than one utility by choosing Track View Utilities again.

Available Utilities

- *Randomize Keys (page 2–562)*: Applies random offset values to Time or Value of selected keys or selected time.
- *Create Out of Range Keys (page 2–562)*: Creates keys in the out-of-range time of a track, when the Out of Range type is something other than Constant. Thus, it converts the specified out-of-range area to a keyed area that you can edit and adjust. Select one or more tracks, set the parameters of the utility, and click Apply. In Function Curve display mode, you must select the curve as well as the track.
- *Select Keys by Time (page 2–563)*: Selects keys within a specified start and end range of time. Lets you select a large range of time when using the mouse might be awkward--for example, if keys are not visible in the Track View Key window.
- *Euler Filter (page 2–564)*: Displays a dialog with controls to remove gimbal flipping from animated tracks using Euler rotation.
- *Soft Selection Settings Manager (page 2–525)*: Displays a dialog with controls to adjust the range and falloff of the soft selection of keys in the Dope Sheet and Curve Editor.
- *Current Value Editor (page 2–565)*: Provides transform type-in capability from within the Track View modes. Allows you to choose between absolute and relative value editing. The name of the controller appears above the axis choices. This utility launches a floating Current Value window that works for either Dope Sheet — Edit keys mode or Curve Editor. It doesn't work for Edit Ranges.

Procedure

To select a Track View utility:

1. Open Curve Editor, and then from the Utilities menu, choose Track View Utilities.
2. In the Track View Utilities dialog, choose from the available utilities, and click OK.

Randomize Keys Utility

Track View > Utilities menu > Track View Utilities > Randomize Keys

Randomize Keys applies random offset values to the times and/or values of selected keys.

Procedure

To use Randomize Keys:

1. Highlight one or more keys, or a block of time, in one or more tracks.

   In Dope Sheet > Edit Keys mode, all selected keys are affected. In Edit Ranges mode, the keys in the selected tracks that are within the selected range of time are affected. In Curve Editor mode, the selected keys on selected curves are affected.

2. From the Utilities menu, choose Track View Utilities, and then choose Randomize Keys.

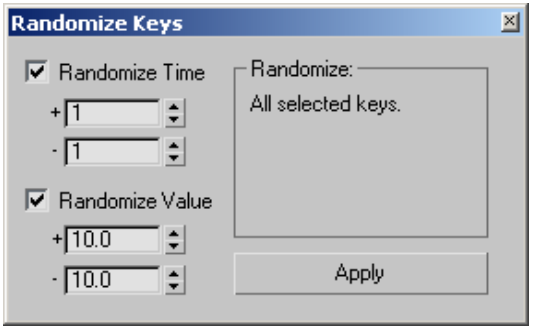
   The Randomize Keys utility is displayed.

3. By default, both Randomize Time and Randomize Value are on. Turn off either if you wish.

4. Set the +/- spinners to specify the desired range of randomized offset.

5. Click Apply.

Interface



Randomize Time—Randomly shifts values in time, based on the spinner settings. For example, if the + spinner is set to 20, and the - spinner is set to 10, the values could shift up as much as 20, and down as much as 10.

+—The amount of random shift in a positive direction.

-—The amount of random shift in a negative direction.

Randomize Value—Randomly shifts values, based on the spinner settings.

+—The amount of random shift in a positive direction.

-—The amount of random shift in a negative direction.

Randomize group box—You can apply randomization either with selected keys or with selected time. When applying to selected keys, this displays the message: "All selected keys." However, when you select a block of time in *Edit Time mode (page 2–566)*, this message is displayed: "Keys in selected tracks that are in the range: n to n."

Apply—Applies the specified random offset values to the selection. You can click this repeatedly for more randomization.

Create Out of Range Keys Utility

Track View > Utilities menu > Track View Utilities > Create Out of Range Keys

Create Out of Range Keys creates keys in the out-of-range time of a track when the out of range type is something other than Constant. It converts the specified out-of-range area to a keyed area that you can edit and adjust.

This utility works both in Curve Editor and Dope Sheet modes.

In Curve Editor mode, you must select the curve as well as the track.

Procedure

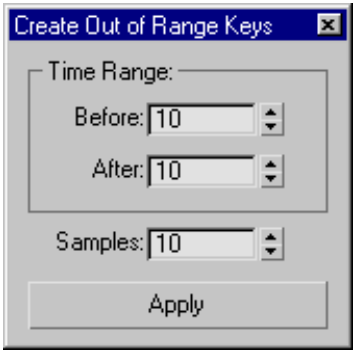
To use Create Out of Range Keys:

1. In the Track View Controller window, select the Position track of an animated object.
2.  On the Track View toolbar, click Parameter Curve Out-of-Range Types button to apply an out-of-range type curve to the selected track.

   Use any curve type except Constant (the default).
3. Click Utilities menu > Track View Utilities. In the dialog box that appears choose Create Out of Range Keys.
4. Set the desired parameters in the Create Out of Range Keys dialog.
5. Click Apply.

   Keys are created in the out-of-range areas of the position track. The track bar expands to the boundaries of the newly created keys.

Interface



Time Range group

Sets the time after the range for generating keys. For example, if your range of keyed animation is from frame 31 to 54, the out-of-range animation might be from 0 to 30 and from 55 to 100.

Before—Specifies the number of frames before the range for generating keys. In the previous example, if you set this option to 20, it would generate keys over frames 10 to 30. When you set this option to 0, no keys are generated before the range.

After—Specifies the number of frames after the range for generating keys. In the first example, if you set this option to 40, it would generate keys over frames 55 to 95. When you set this to option to 0, no keys are generated after the range.

Samples—Specifies the number of keys to be generated for both the before and after time range, based on the settings in Before and After. In the example above, if Samples is set to 20, 20 new keys would be generated over frames 10 to 30, and another 20 new keys would be generated over frames 55 to 95.

Apply—Generates the keys.

Note: Keys are generated only for non-constant out-of-range types. If the area before or after the range is the default, Constant type, no keys are generated in that area.

Select Keys By Time Utility

Track View > Utilities menu > Track View Utilities > Select Keys by Time

The Select Keys By Time utility lets you select keys within a specified start and end range of time. You can select a large range of time, which might be difficult to select using the mouse in the Track View Key window for example.

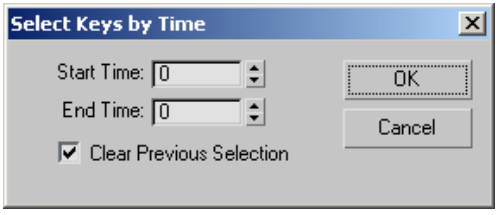
This utility works in Curve Editor and Dope Sheet Edit Keys modes.

Procedure

To use Select Keys By Time:

1. In the Controller window, highlight the tracks you want to work with.
2. On the Utilities menu choose >Track View Utilities, then choose the Select Keys By Time utility from the Track View Utilities dialog.
3. Set a range and click OK.
   - Keys within the range in the highlighted tracks are selected.
   - If you select Clear Previous Selection, all keys are cleared before the keys within the time range are selected.
   - In Dope Sheet, highlight a single track to select all keys in the descendants of the track.

Interface



Start Time—Specifies the start range for selecting keys.

End Time—Specifies the end range for selecting keys.

The Start Time and End Time spinners are activated to match the current time selection.

Clear Previous Selection—Clears all keys before keys within the specified time range are selected.

In Dope Sheet mode, highlight a single track to select all keys in the descendants of the track. Modify Subtree must be on for this to work.

Euler Filter

Track View > Utilities menu > Track View Utilities > Euler Filter

This Track View utility corrects for gimbal flipping (anomalous rotation animation) in objects animated using *Euler rotation (page 2–318)* by processing existing animation keys in selected animated tracks. It's particularly useful for cleaning up rotation artifacts when importing raw motion-capture data. This process does not break any keyed orientations.

By default, the utility modifies only frames with keys on all three tracks (X, Y, and Z). By turning on Add Keys If Needed, you can perform the correction on frames with X, Y, and Z keys and also add keys at frames that contain only one or two keyed axes but require correction, resulting in X, Y, and Z keys at those frames.

The default range for filtering is the current active range for the scene, which automatically appears in the utility. To isolate the operation to a specific subset of the animation, change the Start Time and End Time settings. Any changes made to the utility remain active during the current 3ds Max session.

Tip: The utility works in Dope Sheet mode, but it's much easier to see what it's doing if you use Curve Editor mode.

Procedure

To use Euler Filter:

1. Select an object animated with *Euler rotation (page 2–318)*.
2. Right-click the object, and from the quad menu choose Curve Editor.

   Track View opens in Curve Editor mode, with the animated tracks highlighted.
3. Make sure the rotation tracks to be filtered are highlighted in Track View.

4. From the Utilities menu in Track View, choose Track View Utilities.

   This opens the Track View Utilities dialog.

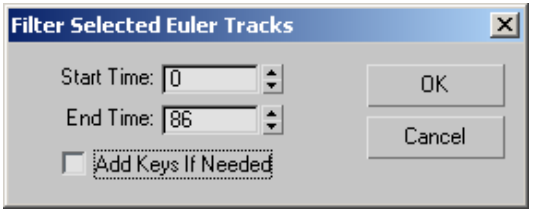
5. In the list of utilities, click Euler Filter and then click OK, or just double-click Euler Filter.

   This opens the Filter Selected Euler Tracks dialog.

6. Change the parameters as necessary, and then click OK.

   The utility adjusts the rotation keys to removed gimbal flipping.

Interface



Start/End Time—The range over which the filtering should occur. Default=current animation range.

Add Keys If Needed—When on, the utility performs the correction on frames with X, Y, and Z keys and also adds keys at frames that contain only one or two keyed axes but require correction, resulting in X, Y, and Z keys at those frames.

OK—Performs the filtering using the current settings.

Cancel—Closes the dialog without performing the filtering.

Current Value Editor

Track View > Utilities menu > Track View Utilities > Track View Utilities Dialog > Current Value Editor

The Current Value Editor gives you a way to use numeric input to affect the values of the keys inside the Track View windows.

Absolute and relative value editing let you increment changes, or apply exact values. The name of the controller whose values are being changed appears above the axis choices.

This utility opens a dockable Current Value window that works for either Dope Sheet - Edit keys mode or Curve Editor. It doesn't work for Edit Ranges.

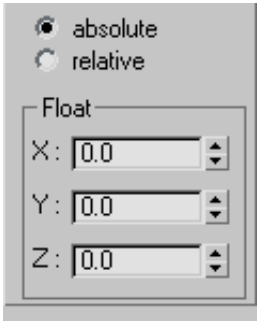
Procedure

To use the Current Value Editor to apply an incremental change:

1. Select the position track of an animated object.
2. Choose Utilities > Track View Utilities > Current Value Editor.
3. Turn on Relative.
4. Change the values for X, Y, or Z.

   The numbers you enter add a relative increment to the existing values.

Interface



Absolute— Applies world space values to the keys. The numbers you enter are the numbers that are applied.

Relative—Increments the values relative to their existing value. The numbers you enter are added to the existing value.

Controller type—Displays the name of the controller about the X, Y, and Z fields.

X, Y, Z—Use these fields to input numbers for relative or absolute value editing.

Edit Time

Use the Edit Time tools to work directly with selected blocks of time in the Dope Sheet Editor, as opposed to working with keys and range bars. A block of time is any contiguous time segment, across one or more tracks, and is independent of key locations.

Enter Edit Time mode by selecting a time segment in the Dope Sheet Editor. Once a time segment has been selected you can insert, cut, copy, paste, or reverse the time segment, including its keys. The time tools are available from the Time menu and the Time: Dope Sheet toolbar. Time tools are unavailable from the Time menu in Edit Ranges mode.

In Edit Time mode, keys and range bars of your animation are there only for reference. You select blocks of time and then apply time-editing functions to your selection.

Edit Time Mode Controls

In Edit Time mode you can perform the following tasks for one or more tracks:

Select Time (page 2–566)

Delete Time (page 2–567)

Cut Time (page 2–567), *Copy Time (page 2–568)*, *Paste Time (page 2–568)*

Reverse Time (page 2–569)

Insert Time (page 2–570)

Scale Time (page 2–570)

Exclude Left End Point (page 2–570)

Exclude Right End Point (page 2–571)

Reduce Keys (page 2–572)

Procedure

To edit time:

1. Select the animated object in the viewport, then right-click and choose Curve Editor.

   This opens Track View navigated to the animated object.

2. On the Modes menu choose Dope Sheet mode.

   The Key window changes from curves to a key spreadsheet.

3. If Edit Keys isn't highlighted on the Keys: Dope Sheet toolbar, click it to turn it on.

   The next steps won't work if you have Edit Ranges on instead of Edit Keys.

4. From the Time menu, choose Select Time.

5. Expand and activate the tracks you wish to alter. For example you could select the Z Position transform track of a bouncing Box object.

6. Drag a time segment out in the Key window. A tooltip displays the selected Start and End frame numbers interactively as you drag.

   This selects the time segment including any keys within it.

7. Perform any of the time tool operations available from the Time Menu or the Time: Dope Sheet toolbar.

Select Time 

Track View > Dope Sheet > Time menu > Select Time

Track View > Dope Sheet > Time: Dope Sheet toolbar > Select Time button

With Select Time you can specify a block of time by dragging in the Key window. You can then apply

any of the following operations to the highlighted block (and any keys contained therein):

- *Delete Time (page 2–567)*
- *Cut Time (page 2–567)*
- *Copy Time (page 2–568)*
- *Paste Time (page 2–568)*
- *Reverse Time (page 2–569)*
- *Insert Time (page 2–570)*
- *Scale Time (page 2–570)*
- *Exclude Left End Point (page 2–570)*
- *Exclude Right End Point (page 2–571)*
- *Reduce Keys (page 2–572)*

Procedure

To select time:

1. In Dope Sheet mode, highlight one or more item labels in the Controller window to specify tracks for time editing.
2.  On the toolbar, click the Select Time button.
3. Drag in the keys window to specify time in the selected tracks.

   A beige bar between two yellow lines appears as you drag across the box grid of the Dope Sheet Key window.

   If a highlighted track does not support time operations, time selection in that track is ignored.

   After specifying a time range, you can perform other time- and key- related operations.

Delete Time 

Track View > Dope Sheet > Specify a time block. > Time: Dope Sheet toolbar > Delete Time button

Use Delete Time to delete a selected block of time and any keys inside the selected block. Keys to the right of the deleted time move to the left.

Note: The deleted block is not copied to the clipboard.

Procedure

To delete a block of time:

1. In Dope Sheet mode, highlight one or more item labels in the Controller window to specify tracks for time editing.
2.  Use *Select Time (page 2–566)* to specify a block of time.
3.  Click Delete Time.

   Keys in the specified time block are deleted, and any keys to the right of the block move leftward.

Cut Time 

Track View > Dope Sheet > Specify a time block. > Time menu > Cut

Track View > Dope Sheet > Specify a time block. > Time: Dope Sheet toolbar > Cut Time button

Use Cut Time to delete a block of time from one or more tracks and place it in the clipboard. Before you can paste time in a track, you must have time in the clipboard. After specifying a block of time, place it in the clipboard with Cut Time or *Copy Time (page 2–568)*.

Procedure

To cut time from tracks:

1. In Dope Sheet mode, highlight one or more item labels in the Controller window to specify tracks for time editing.
2.  Use *Select Time (page 2–566)* to specify a block of time.
3.  Click Cut Time, or, from the Time menu, choose Cut Time.

   The block of time is deleted from the selected tracks and stored in the time clipboard. Keys to the right of the deleted time move left.

Copy Time 

Track View > Dope Sheet > Specify a time block. > Time menu > Copy

Track View > Dope Sheet > Specify a time block. > Time: Dope Sheet toolbar > Copy Time button

Use Copy Time to copy a block of time from one or more tracks to the clipboard, after which you can paste it to other tracks.

You can use the clipboard to copy time, with keys, to different places in the same track or from one track to another. For example, you can copy position keys from one object to another.

Procedure

To copy time from tracks:

1. In Dope Sheet mode, highlight one or more item labels in the Controller window to specify tracks for time editing.
2.  Use *Select Time (page 2–566)* to specify a block of time.
3.  Click Copy Time, or choose Time menu > Copy.

   The block of time is copied to the time clipboard. The original block remains unchanged.

Paste Time 

Track View > Dope Sheet > Specify a time block. > Cut or Copy Time > Specify a different time block. > Time menu > Paste

Track View > Dope Sheet > Specify a time block. > Cut or Copy Time > Specify a different time block. > Time: Dope Sheet toolbar > Paste Time

Use Paste Time to paste a block of time from the clipboard into one or more tracks.

The conditions for pasting time into a track are:

- If the block of time on the clipboard is from a single track, you can paste it into any track using the same type of controller.
- If the block of time on the clipboard is from multiple tracks, you can paste it into a selection of multiple tracks as long as the controllers for the selected tracks include valid track types. For example, position tracks in the clipboard are pasted to position tracks in the selection.

If the above conditions are not met, the Paste Time function has no effect.

Time editing is available only in Dope Sheet mode; it is not available in the Function Curve Editor.

In Edit Time mode, you can copy and paste controller tracks that have no animation keys. When a track has no key, its value at frame 0 is used. You must select a block of time before cutting or pasting; otherwise the time is ignored.

Relative and Absolute Pasting

The values of all the pasted keys are adjusted so that the first pasted key has a value equal to the value at the time of the insertion point. For example, the value of the controller at frame 50 is 10. The clipboard holds three keys spanning 50 frames with values 20, 30, and 40. When you paste the three keys at frame 50, the first key has a value of 20, but the insertion point has a value of 10. 3ds Max subtracts 10 from the pasted key to maintain the value at the insertion point. 3ds Max then subtracts 10 from the remaining pasted keys, resulting in three pasted keys valued at 10, 20, and 30. In addition, any keys after the insertion range are also adjusted by the net change over the range being pasted. The net change is the value of the last key pasted minus the value of the first key pasted. In this case 40-20=20. Every key after the insertion point is increased by 20.

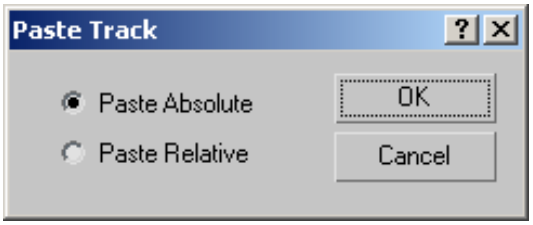
Procedure

To paste time into tracks:

1. In Dope Sheet mode, after using *Cut Time (page 2–567)* or *Copy Time (page 2–568)*, highlight one or more item labels in the Controller window to specify tracks for pasting.
2.  Use *Select Time (page 2–566)* to specify a block of time.
3.  Click the Paste Time button on the Time: Track View toolbar.

   The Paste Track dialog opens.
4. In the Paste Track dialog, choose Paste Absolute or Paste Relative, and then click OK.

Interface



Options on the Paste Track dialog are as follows:

Paste Absolute—Replaces the current animation values with the values in the clipboard. Use this method when you want to replace one animated effect with another.

Paste Relative—Adds the animation values in the clipboard to the current animation values. Use this method when you want to layer animation onto an existing effect.

Reverse Time 

Track View > Dope Sheet > Specify a time block. > Time menu > Reverse

Track View > Dope Sheet > Specify a time block. > Time: Dope Sheet toolbar > Reverse Time

Reverse Time flips the order of keys within the selected time. You can reverse time by scaling a selection past its left edge, but this also changes the position of the selection and the remaining keys around it. Use Reverse Time to reverse keys within a designated block of time.

Procedure

To reverse time:

1. In Dope Sheet mode, highlight one or more item labels in the Controller window to specify tracks for time editing.
2.  Use *Select Time (page 2–566)* to specify a block of time.

3.  Click the Reverse Time button, or choose Time menu > Reverse.

   The position of the selected block of time does not change, but order of keys within the block is reversed.

Insert Time 

Track View > Dope Sheet > Time: Dope Sheet toolbar > Insert Time button

Use Insert Time to interpose time into highlighted tracks. Inserting time adds time at a selected point in your animation, making existing keys slide out of the way.

Procedure

To insert time into tracks

1. In Dope Sheet mode, highlight one or more item labels in the Controller window to specify tracks for time editing.
2.  On the Dope Sheet toolbar, click the Insert Time button.
3. Drag in the Key window to insert time into the tracks.

   Tip: To insert time globally, highlight the World track, turn on the Modify Child Keys button, then insert time into the World track.

Scale Time 

Track View > Dope Sheet > Time: Dope Sheet toolbar > Scale Time button

Scale Time scales a block of time. You can scale down to fit into less time, or expand it to fill more time. Scale Time doesn't use the Track View time slider as the scale origin reference; it always scales from the first frame of the key selection.

Procedure

To scale time:

1.  In the Dope Sheet Editor, on the Time: Dope Sheet toolbar, click Scale Time.
2. In the controller window, click item labels to highlight tracks for time editing.
3. Drag out a time block in the Key window, or use an existing block.
4. Move your cursor over the active time block in the Key window. The cursor changes to show you when you can scale.
5. Do one of the following:
   - Drag to the right within the selection to expand time from the left edge of the selection. All keys to the right of the selection slide right as the selection expands.
   - Drag to the left within the selection to reduce time towards the left edge of the selection. All keys to the right of the selection slide left as the selection shrinks.
   - Drag past the left edge of the selection to reverse time and expand it with a negative scale factor. Keys inside the selection, and keys to the right of the selection, can overlap keys to the left of the selection.

Exclude Left End Point 

Track View > Dope Sheet > Extras: Dope Sheet toolbar > Exclude Left End Point button

Use Exclude Left End Point to exclude the beginning key in a selected block of time.

If you paste the same block of time repeatedly, one block following the other, you can create a looping

segment in your animation. To create a smooth looping animation, you need to exclude either the first or last key of the copied block to prevent keys from doubling up at the ends.

The key to be excluded must be at the exact start time of the copied block of time.

Note: By default, the Extras: Dope Sheet toolbar doesn't appear in the Dope Sheet Editor. You'll find a method for displaying it in the following procedure.

Procedure

To animate a loop by copying and pasting keys:

1. In Dope Sheet mode, highlight one or more item labels in the Controller window to specify tracks for time editing.
2.  Use *Select Time (page 2–566)* to specify a block of time.

   At least one track in the block should start with a keyframe. Also, for a smooth loop, the first and last frame of the selection should be the same.
3. Right-click an empty section of the toolbar area and choose Show Toolbars > Extras: Dope Sheet.

   The Exclude Left End Point and Exclude Right End Point buttons appear.
4.  On the Track View toolbar, click Exclude Left End Point.
5.  On the Track View toolbar, click Copy Time.

   The selection is copied to the clipboard, minus the first frame.
6. Click to define an insertion point in the keys window.

   The last position key should be your insert point.
7.  Click Paste Time.

   The Paste Time dialog appears.
8. In the dialog choose Paste Absolute or Paste Relative. For an animated loop choose Paste Absolute, then click OK.

   The time, including keys, is pasted.
9. Play your animation to observe the effect.

Exclude Right End Point 

Track View > Dope Sheet > Extras: Dope Sheet toolbar > Exclude Right End Point button

If you paste the same block of time repeatedly, one block following the other, you can create a looping segment in your animation. To create a smooth looping animation, you need to exclude either the first or last key of the copied block to prevent keys from doubling up at the ends.

The key to be excluded must be at the exact end time of the copied block of time.

Note: By default, the Extras: Dope Sheet toolbar doesn't appear in the Dope Sheet Editor. You'll find a method for displaying it in the following procedure.

Procedure

To animate a loop by copying and pasting keys:

1. In Dope Sheet mode, highlight one or more item labels in the Controller window to specify tracks for time editing.
2.  Use *Select Time (page 2–566)* to specify a block of time.

At least one track in the block should end with a keyframe. Also, for a smooth loop, the first and last frame of the selection should be the same.

3. Right-click an empty section of the toolbar area and choose Show Toolbars > Extras: Dope Sheet.

   The Exclude Left End Point and Exclude Right End Point buttons appear.

4.  On the Track View toolbar, click Exclude Right End Point.

5.  On the Track View toolbar, click Copy Time.

   The selection is copied to the clipboard, minus the first frame.

6. Click to define an insertion point in the keys window.

   The last position key should be your insert point.

7.  Click Paste Time.

   The Paste Time dialog appears.

8. In the dialog choose Paste Absolute or Paste Relative. For an animated loop choose Paste Absolute, then click OK.

   The time, including keys, is pasted.

9. Play your animation to observe the effect.

Reduce Keys 

Track View > Curve Editor or Dope Sheet > Highlight the keys to reduce in the Key window. > Keys menu > Reduce Keys

Track View > Curve Editor > Highlight the keys to reduce in the Key window. > Keys: Track View toolbar > Reduce Keys

Use Reduce Keys to decrease key density. Animating with inverse kinematics, or creating any complex animation, can result in many keys, which can make editing the animation difficult. In the case of applied inverse kinematics, the software generates a key on nearly every frame. Often, the same animation can be produced with fewer keys. Having fewer keys in a track makes it easier to change your animation.

Reduce Keys analyzes the pattern of keys in a block of time and creates a new pattern of fewer keys that produces nearly the same animation. You specify how closely the new animation matches the original.

Note: The Reduce Keys button, depicted above, appears on the Keys toolbar only in Curve Editor, but you can add it to a Dope Sheet toolbar with *Customize User Interface (page 3–792)*.

Procedure

To reduce keys:

1. Select an animated object in the viewport.
2. Right-click the object and choose Curve Editor or Dope Sheet from the quad menu.
3. In the Hierarchy list, highlight the tracks whose keys you want to reduce. Reduce Keys works only on highlighted tracks.
4. Optionally, in the Key window, specify a time range within which to reduce keys. Highlight a key at either end of the range to reduce, or drag a selection rectangle around the keys you

want to reduce. Alternatively, in Dope Sheet mode, highlight specific tracks and then use *Select Time (page 2–566)* to designate a block of time within which to reduce keys. If no keys are highlighted in a highlighted track, Reduce Keys works on all keys in that track.

You can use any of the above methods to specify different ranges for different highlighted tracks.

5. From the Keys menu, choose Reduce Keys.

   The Reduce Keys dialog appears. This dialog has a single Threshold parameter.

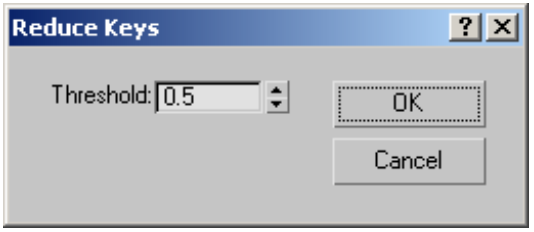
   Raising the Threshold setting will increase the number of keys that are reduced. The higher the threshold, the greater the reduction.

6. Click OK to reduce the keys.

   Observe the results. If you are left with too few keys, press Ctrl+Z to undo, then lower the threshold and reduce the keys again. If too many keys remain, increase the threshold and reduce keys again.

   Tip: It will take repeated experimentation to find exactly the right threshold for your particular animation.

Interface



Reduce Keys—Displays the Reduce Keys dialog.

Threshold—Sets a threshold value. Higher values will result in fewer keys.

OK—Accepts the Threshold setting and reduces keys as follows:

- Keys are reduced only in highlighted tracks.
- If any of a track's keys are highlighted, reduction is performed only within the indicated range (that is, between the leftmost and rightmost highlighted key for each track). If no key is highlighted, reduction is performed on all keys in the track.

Edit Ranges

Track View > Dope Sheet > Track View toolbar > Edit Ranges

Edit Ranges displays all tracks as range bars. This mode is useful for quickly scaling and sliding complete animation tracks.

- Drag the range bar of an animation track to change all animation in that track.
- Drag range bars in higher-level tracks to change all animation in multiple tracks.

Note: You cannot access individual key values in this mode.

The Ranges: Dope Sheet toolbar contains tools for working with Ranges. Right-click an empty area adjacent to the Dope Sheet toolbar and choose Show Toolbars > Ranges: Dope Sheet to display the toolbar. Save your layout after you do.

Besides the Edit Ranges button, the Ranges: Dope Sheet toolbar shows the following buttons:

Position Ranges (page 2–574)

Recouple Ranges (page 2–574)

Procedure

To drag the ranges of an object and all of its linked descendants:

1.  On the Track View toolbar, click Edit Ranges.

2.  On the Track View toolbar, click Modify Subtree.
3. In the Track View Key window, drag an Object range bar or the World range bar.

   With Modify Subtree on, a range bar displays in the Objects track. The Objects branch is the default parent of all objects in the scene.

   Dragging a parent-object range bar with Modify Subtree on affects all tracks subordinate to the object and all tracks of all of its linked descendants.

Position Ranges 

Track View > Dope Sheet > Ranges: Dope Sheet toolbar > Position Ranges

Position Ranges mode allows you to position range bars independently from their associated keys and produce special effects.

The process of adjusting a range bar independent of its keys is called *decoupling the range*. Positioning a range bar so it matches the first to last keys of a track is called *recoupling the range*.

Note: The Ranges: Dope Sheet toolbar doesn't appear in Track View by default. To open it, right-click an empty area adjacent to the Dope Sheet toolbar and choose Show Toolbars > Ranges: Dope Sheet. Save your layout after you do.

Decouple a Range Bar

You decouple a range bar for two reasons:

- You want some keys at the start or end of an animation range not to take effect.

  Keys outside of the range bar are ignored during animation playback. The keys outside the range still affect interpolated values inside the range, but the applied Out-of-Range type is used to animate time outside the range bar.
- You want to add extra time before the first key or after the last key that is not affected by the applied Out-of-Range type.

  Time beyond the first or last key of the track, but still within the range, uses the constant value of the nearest key within the range.

Procedure

To decouple a range bar from its keys:

1. Right-click an empty area adjacent to the Dope Sheet toolbar and choose Show Toolbars > Ranges: Dope Sheet.
2.  On the Track View toolbar, click Position Ranges.
3. In the Track View Key window, drag the entire range bar left or right of the keys, or drag either end of the range bar to make it longer or shorter than the keys.

   Using *Out-of-Range types (page 2–551)* in conjunction with this procedure allows you to change the behavior of an animated loop. Positioning either end of a range bar changes the loop.

Recouple Ranges 

Track View > Dope Sheet > Ranges: Dope Sheet toolbar > Recouple Ranges

Recouple Ranges resizes the range bar so that the beginning and end of the range matches the first and last keys in the track. This is useful to quickly realign the range and the keys after editing.

Note: The Ranges: Dope Sheet toolbar doesn't appear in Track View by default. To open it, right-click an empty area adjacent to the Dope

Sheet toolbar and choose Show Toolbars > Ranges: Dope Sheet. Save your layout after you do.

Procedure

To recouple a range:

1. Right-click an empty area adjacent to the Dope Sheet toolbar and choose Show Toolbars > Ranges: Dope Sheet.
2.  On the Track View toolbar, click Position Ranges.
3. Select one or more item labels in the Hierarchy list to select tracks to be recoupled.
4.  On the Track View toolbar, click Recouple Ranges.

   The range bars for the selected tracks are positioned to match up with the first and last keys in their track.

   Position Ranges mode allows you to decouple the range bar from the animation keys. Recouple Ranges is a quick way to realign the range bars.

Editing Tracks: Copying, Pasting, and Handling Instances and References

Copying and Pasting Items

You can copy and paste geometry, lights, materials, and animation controllers between items in the Track View Hierarchy list of the controller window. You can copy these categories of items in Track View:

Containers—Items with multiple branches that completely define something in your scene. Container items that can be copied include:

- Material Parameters containing the Basic Parameters for a material definition.
- Material Maps containing the entire set of maps and map parameters assigned to a material.
- Map definitions containing a single map type with its associated parameters and coordinates.
- Map Coordinates containing the map XYZ and UVW coordinate offset, tiling, and angle settings.
- Map Parameters containing the parameters for a specific map type.
- Objects, on levels below the transforms, containing creation parameters for an unmodified object.
- Modified Object, containing modifiers applied to an object and the object-creation parameters.
- Modifiers containing modifier parameters.

Controllers—These control actual animated values for each parameter. Keep in mind that when you copy containers, you are actually copying and pasting groups of related controllers. Details about copying and pasting single controllers are presented in *Assigning Controllers (page 2–292)*.

Whether you copy objects, materials, or controllers, the same basic principles apply. Here are restrictions for using Copy and Paste:

- You can copy only a single highlighted item.
- You can paste a copied item only into another item of the same type. An exception to this restriction involves pasting Object and Modified Object containers.
- You can paste a copied item into a selection of multiple items only if all the items are of the same type.

- When pasting items, you can choose to make an instance or a copy of the pasted item.
- You cannot copy/paste actively linked objects.

Copying Items

You copy an item by highlighting it in the Controller window Hierarchy list, and then right-clicking and choosing Copy from the quad menu. If Copy is not available the selected item is not a valid copy source. 3ds Max disables Copy when the selected item is not a valid copy source or if multiple items are selected.

Pasting Items

Pasting items involves a few more choices than copying. You select one or more items from the Hierarchy list. If Paste is available the selected items are valid paste targets. The program disables Paste if the selected targets are not all the same type or if they do not match the type of item in the clipboard.

You can use Paste as follows:

- Paste into a single target item.
- Paste into multiple target items.
- Paste copies or instances.
- Choose to convert other instances in the scene automatically.

Clicking Paste on the quad menu displays the Paste dialog, with controls for determining how the Paste operation will be carried out.

Copy—Pastes the item in the clipboard as an *independent copy*. The target item will have no connection to the copied source item.

Instance—Pastes the item in the clipboard as an *instance of the copied source item*. The target item will be an instance of the source item. Any change you make to either item affects the other.

Replace All Instances—Controls whether existing instances of the target item are also converted to the paste source or left as they are.

Making Instances Unique

You can convert instanced items to unique items by clicking *Make Unique (page 2–550)* on the Controller menu. If the selected item is not an instance, or if a selection of multiple items does not contain similar instances, Make Unique is disabled.

Copying and Pasting Objects

In the Track View controller window you can use the Hierarchy list to copy and paste objects. This allows you to replace the geometry of one animated object with the geometry of another object. This is similar to the functionality provided by XRef objects and scenes where you have to option of defining *proxy* or stand-in objects.

Copying and pasting objects is achieved by highlighting the Object track for the source object, copying it, then pasting it to the Object track of the target object. The Object track is found beneath the Transforms and Modified objects entries, and appears with parentheses that define the geometry type. For example, Object (Box) or Object (Editable Mesh) are two label names that might appear on the correct track for object cut and paste.

Using these two container types, you can copy geometry between objects.

To copy one object to other objects:

1. In the Curve Editor controller window, find the object that you want to copy from.
2. Navigate to the Object track that is beneath the Transform or Modified Object entries. The name will be something like Object

(PivotDoor) or Object (Editable Mesh), for example.

Note: If the object is modified, the order of the name of its Object track in the Hierarchy list is reversed; for example, Object (Sphere) becomes Sphere (Object).

3. Highlight the object track you wish to copy, then right-click and choose Copy from the quad menu.
4. Navigate to the object track for the target object, highlight it, and then right-click and choose Paste. Set options in the Paste dialog, then click OK.

   The object is replaced with the new one in the viewport as well as in the Hierarchy list of the controller window.

   Warning: The original target has been completely replaced by the pasted object. Use *Save Selected (page 3–392)* on the target object to save a copy of that object, if you think you'll need to use it again sometime.

Here are some examples of how you might use object copy and paste:

- Paste an Object to another Object to replace one simple object with another. For example, you animate the transforms of a box object and want to replace the box with a torus.
- Paste an Object into a Modified Object to replace a complex, modified object with a simple stand-in object. For example, you model a complicated vehicle and you want to replace it with a box while you animate its transforms. Be sure to save the Modified Object to a file (using *Save Selected (page 3–392)*).
- Paste a Modified Object into another Modified Object to replace one complex modified object with another. For example, you animated a flight of jets and want to replace them with helicopters.
- Paste a Modified Object into an Object to replace simple stand-in objects with complex modified objects. For example, you have animated a box and now want to replace it with a complex model of a jet.

Copying and Pasting Object Modifiers

You can also copy and paste modifiers below the Modified Object container. Copy and paste modifiers to do the following:

- Copy modifiers within the modifier stack of a single object.
- Copy modifiers between objects.

Unlike pasting other items, modifiers do not replace the highlighted item when you click Paste. Instead, the pasted modifier is inserted above the selected item.

Making Instance and Reference Controllers and Objects Unique



You can convert instanced and referenced objects into unique objects by clicking Make Controller Unique on the Track View toolbar or by choosing Make Unique on the Controller menu. The result of clicking Make Unique depends on how the object was instanced or referenced, and which object container you select in the Hierarchy list.

When you make an instance or reference object, you cause the data flow to branch on its way from the master object to two or more named objects.

- In Track View, these branches usually occur at Modified Object containers and sometimes at the base Object container.
- In Track View, instances and references are displayed in bold face.

- In the modifier stack, these branches display as a horizontal line, called a *derived object* line.

For descriptions of object data flow and instance objects, see *Understanding Object Data Flow (page 1–494)*.

You can make an object unique by selecting a Modified Object or an Object container and clicking Make Unique on the Track View toolbar. Your result depends on whether or not the data flow branches at the selected container.

- If the data flow does not branch anywhere below the selected container, nothing happens.
- If the data flow does branch at or below the selected container, the data flow above the branch is split from the current data flow as a unique object. The data flow from the selected container to the master object is copied and attached to the new unique object.

Making a Selection of Multiple Instances Unique

If you select multiple objects that are instances of each other, or share instanced modifiers, you can choose how to make them unique. After you click Make Unique, the Make Unique dialog appears, with the question, "Do you want to make the selected items unique with respect to each other?"

- Click Yes to make each object in the selection completely unique.
- Click No to leave the objects in the selection as instances, but to make them unique from other objects not in the selection.

Curve Editor

Main toolbar > Curve Editor (Open) > Highlight animated tracks. > Curves are displayed in Key Window.

Graph Editors > Track View - Curve Editor > Highlight animated tracks. > Curves appear in Key window.

Right-click any animated object. > Curve Editor > Curves appear in Key window.

The Curve Editor displays animation as function curves, charting the controllers' changes of value over time.

Highlighting a controller label displays keys for that item as vertices on the function curve. You can then highlight keys and change their values.

You can add keys to function curves for tracks that are not yet animated. The curves appear as straight lines. When you add a key to a function curve, a controller is created for that track.

Not all controller types display function curves. For example, a TCB Rotation controller doesn't display a function curve, while an Euler XYZ rotation controller shows individual curves for X, Y and Z. Some controllers display individual tracks for each axis, while other controllers combine the three axis values into a single curve.

The Track View tools for working with function curves, found on the Curve Editor *toolbars (page 2–535)* and menus, include the following:

Move Keys (Curve Editor) (page 2–579)

Scale Keys (page 2–580)

Slide Keys (page 2–558)

Scale Values (page 2–581)

Add Keys (Curve Editor) (page 2–581)

Draw Curves (page 2–582)

Reduce Keys (page 2–572)

Show All Tangents (page 2–531)

Show Tangents (page 2–582)

Lock Tangents (page 2–583)

Apply Ease Curve / Apply Multiplier Curve (page 2–584)

Remove Ease/Multiplier Curve (page 2–585)

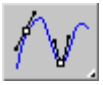
On/Off (Curves) (page 2–585)

Ease Curve Out-of-Range Types (page 2–585)

Multiplier Curve Out-of-Range Types (page 2–586)

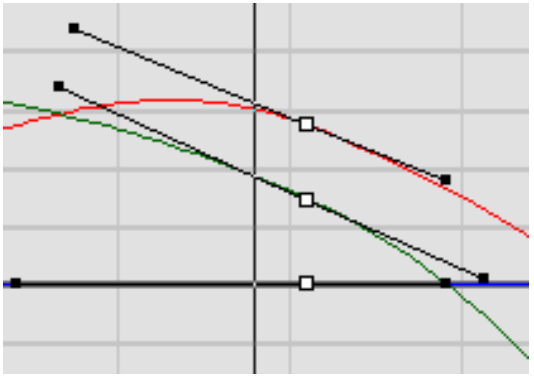
Freeze Non-Selected Curves (page 2–587)

Bezier Tangent Handles



Custom tangents

You can use Bezier handles in the Curve Editor to change the shape of the trajectory through a key. To display handles on a key, right-click a key in the Track View Key Window and choose the Custom Tangents option. You can assign the Custom Tangent type only to items using a Bezier controller. This will display tangent handles on the keys that you can manipulate horizontally and vertically.



Procedure

To display function curves:

1. Select an animated object in a viewport.
2. Right-click the active viewport and choose Curve Editor.

   The function curves for the object's animation are displayed in the Key window of the Function Curve Editor.
3. To display additional curves at the same time, from the Track View menu choose Settings > Manual Navigation.
4. Navigate the controller window until you see the additional tracks, and then hold down the Ctrl key and click them.

   The curves for all highlighted tracks are displayed together in the Key window.

Move Keys (Curve Editor)





Track View > Curve Editor > Track View toolbar > Move Keys flyout > Move Keys, Move Keys Horizontal, or Move Keys Vertical

With the Move Keys flyout buttons in the Curve Editor, you can move one or more highlighted keys to change keys' timing, values, or both.

You can clone keys by holding down Shift while dragging. The new keys are added to the existing keys on the curve.

You can also clone keys to replace existing keys by holding down Shift+Alt while dragging.

Tip: When you select keys to clone and replace others, you can include the key before the first selected key to preserve the tangent of your first copied key.

Procedure

To move a selection of keys:

1. In the Hierarchy list, highlight one or more tracks.

   Only keys in highlighted tracks appear in the Key window.

   By default, the Move Keys button is yellow, showing it is already on.
2. In the Track View Key window, highlight one or more keys.
3. Drag the highlighted keys to a new location, or press Shift and drag the highlighted keys to add copies at the new location.
4. Highlight another group of keys.
5. Press Shift+Alt, and drag the highlighted keys towards existing keys to replace them with copies.

Interface


Move Keys—Moves keys in both horizontal and vertical axes, to change both timing and values.


Move Keys Horizontal—Constrains movement to the horizontal axis, to change timing.


Move Keys Vertical—Constrains movement to the vertical axis, to change values.

Scale Keys 

Track View > Curve Editor or Dope Sheet > Track View toolbar > Scale Keys

Track View > Curve Editor or Dope Sheet > Keys menu > Scale Keys - Time

Use Scale Keys to scale keys horizontally, changing the horizontal location and amount of time covered by selected keys. The horizontal location of the keys represents the frame in time at which the key action takes place.

The scale center for scale time is the current time set by the time slider. You can scale keys about any point in time by dragging the time slider before you use Scale Keys.

Procedure

To scale a selection of keys horizontally:

1. Right-click an animated object and choose Curve Editor.
2. In the Hierarchy list, highlight a track.
3. Scrub the *time slider (page 3–701)* to the frame to use as the scale center.
4.  On the Track View toolbar, click Scale Keys.
5. Highlight one or more keys.
6. In the Track View Key window, drag any of the highlighted keys horizontally to scale, or press Shift and drag to add scaled copies of selected keys.
7. The scaling effect depends on the direction you drag with respect to the time slider location:
   - Drag away from the time slider location to expand keys' scaling.
   - Drag toward the time slider location to reduce keys' scaling.

- Drag past the time slider location to reverse the keys' order and expand the keys away from the current time.

Scale Values 

Main Track View > Curve Editor > Track View toolbar > Scale Values

Track View > Curve Editor > Keys menu> Scale Values

Scale Values is used to proportionally increase or decrease the vertical distance between selected keys on the function curves in Curve Editor. It uses the *Scale Origin Line (page 2–506)* as a movable reference point for the scale operation. Practically speaking, scaling values always creates spatial adjustments, rather than temporal. To scale in time, use *Scale Time (page 2–570)* instead.

Note: Scale Values is not available in Dope Sheet mode.

Procedure

To scale key values:

1. Select an object in the viewport, then right-click and choose Curve Editor.
2. Highlight one or more tracks, if necessary, and then highlight the keys you want to scale.
3.  On the Curve Editor toolbar, click Scale Values.

   The orange Scale Values Origin line appears at the 0 location on the vertical scale.
4. Drag the Scale Values Origin line to where you want to scale values from or to.
5. Drag any highlighted key vertically to scale all of them.

   Tip: If you use *Lock Selection (page 2–555)* after highlighting the keys, you can drag anywhere in the Key window to create the scale operation.

Add Keys (Curve Editor) 

Track View > Curve Editor > Track View toolbar > Add Keys

Track View > Curve Editor > Keys menu > Add Keys

When Add Keys is active, you can click anywhere on a function curve to add a key at that location on the curve. Add Keys stays active until you click another mode.

If the curve is part of a multicurve controller, vertices are added to each curve at the time location where you clicked. For example, adding a key to the red X curve of a TCB Position track simultaneously adds keys to the green Y curve and the blue Z curve.

Tip: You can move an existing key while in Add Keys mode simply by dragging it.

Adding Keys to a Multicurve Controller

The value of each key is set by one of the following conditions:

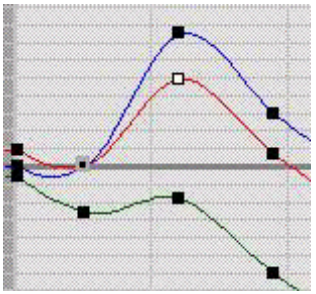
- Keys added between keys receive the interpolated value shown by the function curve at that location.
- Keys added before the first or after the last key on the curve receive the same value as the nearest key.
- You can change a key value and/or time while adding it by dragging after clicking the curve.

Procedure

To add keys to a function curve:

1. In the Hierarchy list, highlight a track.
2.  On the Track View toolbar, click Add Keys.
3. Click the function curve (Curve Editor) where you want to add the key.

If the highlighted curve is part of a multicurve controller, vertices are added to each curve at the time location where you clicked. For example, if you are using a Bezier Position controller, adding a key to the red X curve of a Position track simultaneously adds keys to the green Y curve and the blue Z curve.



Draw Curves

Track View > Curve Editor > Keys: Track View toolbar > Draw Curves button

Use this tool to draw new curves, or revise existing ones by sketching directly on the function curve in the Curve Editor Key window.

The speed with which you draw determines how many keys are added to the curve. If there are too many keys, use Reduce Keys to make fewer keys on the curve.

Procedure

To draw a curve:

1. On the Curve Editor toolbar, click Draw Curves.
2. Select the track of the curve you want to draw on.
3. Move your cursor onto the curve and draw.

   If there are too many points after drawing curves, delete or *reduce (page 2–572)* keys.

Show Tangents

Track View > Curve Editor > Curves: Track View Toolbar > Show Tangents button

Show Tangents lets you view adjust the tangents of highlighted vertices in Curve Editor.

Custom Tangent

You can assign the Custom Tangent type to any controller that displays keys on a curve. Highlight the key on the curve and then click the Set Tangents To Custom icon on the Key Tangents: Track View toolbar to display the handle controls.

The tangent handles can be either continuous or discontinuous. To make a handle discontinuous, hold down the Shift key when you move it.

- Continuous handles behave like a single lever with its pivot point at the key location. Moving either end of the handle also moves the opposite end, maintaining the continuity of the tangent. The function curve, and the animation it represents, moves smoothly through a key with continuous tangent handles.
- Discontinuous handles behave like leaves on a hinge, with the hinge pin at the key location. Each end moves independently of the other.

To convert a discontinuous tangent to a continuous tangent, change the tangent type from Custom to Smooth and then back to Custom again. This technique is performed at the key level and can affect unselected vertices if they belong to the same key as one of the selected vertices.

Horizontal Bezier Handle Control

When using the Custom Tangent type, a tangent handle can be moved horizontally in length away from its keyframe to affect the shape of the function curve. This affects the amount of

influence the keyframe has over the interpolation between keyframes.

This feature can be used to aid in the slow-in and slow-out workflow of traditional animation. For example, suppose you're animating a ball bouncing up and down on the world's Z axis. On frame 1, the ball is on the ground; on frame 8, the ball is in the air; on frame 16 the ball is back on the ground. You could adjust the tangents for frame 8 so that their length pulls more of the curve towards its keyframe and away from the keyframes on frame 1 and 16. When rendered, the ball will be more visible on screen near the top of its bounce and blurrier near the points where it touches the ground.

Procedures

To assign a Custom tangent type:

1. Open the Curve Editor.
2. In the Hierarchy list, highlight a track that uses any controller that displays keys on the curve.
3. Highlight one or more keys on the curve.
4.  On the Key Tangents: Track View toolbar, click Set Tangents to Custom.

   The handles are now all Custom and available for editing.

   Tip: If a highlighted key has Auto Tangency set, you can just move its handle to instantly convert the key to Custom. For selections of multiple keys, choose Lock Tangents, then move any handle to convert all the keys to Custom tangency.

To edit the horizontal length of a Bezier handle:

1. Assign a Custom Tangent type to the key you want to edit on the function curve, or highlight a key that already has one.
2.  Make sure Show Tangents is on. The toggle is located on the Track View toolbar.
3. Highlight the key on the function curve.

   Tangent handles appear on the highlighted keyframe.
4. Move one of the handles horizontally away from its key.

Interface

Show Tangents—Displays the tangents of highlighted vertices, whereupon you can adjust the tangents by dragging the handle endpoints.

By holding down the Shift key, you can edit the tangents handles separately.

Lock Tangents 

Track View > Curve Editor > Select keys on curves. > Curves: Track View Toolbar > Lock Tangents

Lock Tangents lets you manipulate handles of multiple vertices simultaneously. When Lock Tangents is active, dragging a tangent handle affects the handles of all highlighted keys. When Lock Tangents is inactive, dragging a tangent handle affects only that handle's key.

Procedure

To drag discontinuous tangent handles:

1. Highlight one or more keys in the Curve Editor. Hold down the Ctrl key to add to the selection set.
2.  Click Lock Tangents to set one of the following modes:
   - **Off:** Dragging a tangent handle affects only the single key displaying the handle.

- **On:** Dragging a tangent handle affects both handles of all selected keys.

Apply Ease Curve / Apply Multiplier Curve

Track View > Curve Editor > Curves menu > Apply Ease Curve or Apply Multiplier Curve

Apply - Ease Curve and Apply - Multiplier Curve (available from the Curves menu in the Curve Editor) create an ease curve or multiplier curve track. These tracks let you modify or intensify an existing animation track without changing the original track.

- A multiplier curve shifts the value of the original track up or down.

  Adjusting the multiplier curve increases or decreases the "intensity" of the original track.

- An ease curve shifts the time of the original track left or right.

  Adjusting the ease curve changes time. The value of the ease curve at a particular frame is a frame value from the original track. For example, if the ease curve is 0 at frame 0 and 10 at frame 10, the original track plays at its original speed. If the ease curve value at frame 10 increases to 20, the original track plays to frame 20 by frame 10: it has been sped up by a factor of two.

Once you apply an ease or multiplier curve, you can edit its keys, ranges, and properties just like any other animation track.

One use of ease and multiplier curves is to blend between the effects of different controllers. This can be an aid to nonlinear editing.

Procedure

To apply an ease curve or a multiplier curve:

1. In the Curve Editor Hierarchy list, highlight the tracks you wish to ease or intensify.

   Not all controller items can receive ease or multiplier curves. For example, the Path Constraint position controller cannot receive an ease or multiplier curve, but its subordinate Percent controller can. If the highlighted controller is incapable of receiving ease or multiplier curves, nothing happens when you apply the curve.

2. From the Curves menu, choose Apply - Ease Curve or Apply - Multiplier Curve.

3. In the Controller window expand the highlighted tracks to view the ease curve.

4. Hold down the Ctrl key and click the ease curve and controller tracks.

5. Adjust the keys on the ease curve and watch the result in the controller curve.

   Use a Multiplier curve to intensify an effect, or an ease curve to make less of an effect.

   Tip: Use Step interpolation to use the ease curve to turn tracks on and off.

Interface

Apply - Ease Curve—Applies a curve to alter the timing of the selected function curve. Ease curves vary the timing of a superior function curve. A normal function curve charts an animated parameter value over time. An ease curve charts changes to the timing of a function curve over time.

Apply - Multiplier Curve—Applies a curve to scale the value of the selected function curve. The value of a multiplier curve is a scale factor applied to the value of its superior function curve.

You apply ease and multiplier curves in Curve Editor mode. Once you apply an ease or multiplier

curve, you can edit its keys, ranges, and properties, as with any other animation track.

Remove Ease/Multiplier Curve

Track View > Curve Editor > Curves menu > Remove

Remove (Ease/Multiplier Curve) lets you delete an ease or multiplier curve at any time.

Procedure

To delete an ease or multiplier curve:

1. In the Hierarchy list, select one or more ease or multiplier curve tracks.
2. On the Track View menu bar, choose Curves > Remove.

   The selected Ease or Multiplier curve is removed.

   Tip: You can use *On/Off (Curves) (page 2–585)* to turn off the ease or multiplier affect without removing the curve.

On/Off (Curves)

Track View > Curve Editor > Curves menu > On/Off

Use On/Off when you want to work with the original function curve rather than a curve altered by ease and multiplier curves. A disabled curve has no effect on its assigned track, but can be enabled to restore the effect.

Procedure

To toggle the active state of an ease or multiplier curve:

1. In the Hierarchy list, highlight one or more ease or multiplier curves.
2. From the Curves menu choose On/Off.

   If you also have the parent track selected (the transform that has the ease applied), you will see the difference in the curve when you turn the ease or multiplier off.

Ease Curve Out-of-Range Types

Track View > Curve Editor > Curves menu > Ease Curve Out-of Range Types

Ease Curve Out-of-Range Types displays a dialog that lets you assign out-of-range types for the current *ease curve (page 3–932)*.

Ease and multiplier curves use the same out-of-range types as other controller tracks. See *Parameter Curve Out-of-Range Types (page 2–551)*.

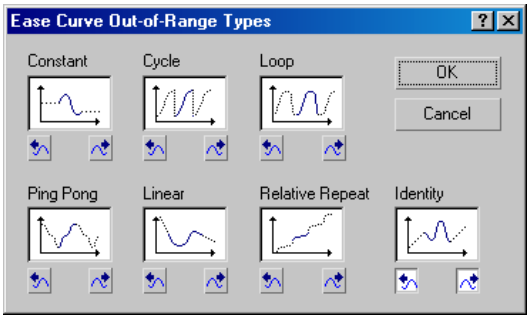
You will usually want to use the same out-of-range type for an ease curve that you use for its superior function curve.

Procedure

To set the out-of-range type for an ease curve:

1. Highlight one or more ease curve items in the Hierarchy list.
2. From the Curves menu choose Ease Curve Out-of-Range Types.
3. Do one of the following.
   - Click an out-of-range type image to apply that type to both the incoming and outgoing curves of the ease curve.
   - Click the left arrow button below an out-of-range type image to apply it to only the incoming curve of the ease curve.
   - Click the right arrow button below an out-of-range type to apply it to only the outgoing curve of the ease curve.

Interface



Constant—Holds the values at the ends of the range for all frames. Use Constant when you want no animated effect before or after the range. The time value at start of the range is held for all frames before the range, and the time value at the end of the range is held for all frames after the range. Constant is the default out-of-range type for multiplier curves.

Cycle—Repeats the same animation as within the range. Use Cycle when you want an animation to repeat exactly.

Loop—Repeats the same animation as within the range but interpolates between the last key and the first key of the range to create a smooth loop. Use Loop with an extended range bar to produce smoothly repeating animation.

Ping Pong—Alternates between a forward and backward repeat of the animation within the range. Use Ping Pong when you want your animation to alternate back and forth.

Linear—Projects the animation value along a line tangent to the ease or multiplier curve at the end of the range. use linear when you want the animation to enter and leave the range at a constant velocity.

Relative Repeat—Repeats the same animation as within the range but offsets each repetition by the value at the end of the range. Use Relative Repeat to create animations that build on each other as they repeat.

Identity—Projects the ease curve along a line with a slope of 1.0 from either end of the range. This causes time to flow at a normal, constant rate outside the range of the ease curve. Use Identity when you want the ease curve to be a localized, non-repeating effect.

Identity for an ease curve produces an effect similar to constant for other types of curves. Identity is not available for multiplier curves; it is the default Out-of-Range type for ease curves.

Multiplier Curve Out-of-Range Types

Track View > Curve Editor > Curves menu > Multiplier Curve Out-of Range Types

Multiplier Curve Out-of-Range Types displays a dialog that lets you assign out-of-range types for the current *multiplier curve (page 3–978)*.

Ease and multiplier curves use the same out-of-range types as other function curves.

You will usually want to use the same out-of-range type for a multiplier curve that you use for its superior function curve.

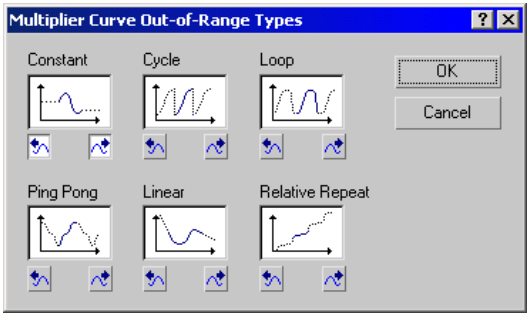
Procedure

To set the out-of-range type for a multiplier curve:

1. Highlight one or more multiplier curve items in the Hierarchy list.
2. From the Curves menu choose Multiplier Curve Out-of-Range Types.
3. Do one of the following.
   - Click an out-of-range type image to apply that type to both the incoming and outgoing curves of the multiplier curve.
   - Click the left arrow button below an out-of-range type image to apply it to only the incoming curve of the multiplier curve.

- Click the right arrow button below an out-of-range type to apply it to only the outgoing curve of the multiplier curve.

Interface



Constant—Holds the values at the ends of the range for all frames. Use Constant when you want no animated effect before or after the range. The time value at start of the range is held for all frames before the range, and the time value at the end of the range is held for all frames after the range. Constant is the default out-of-range type for multiplier curves.

Cycle—Repeats the same animation as within the range. Use Cycle when you want an animation to repeat exactly.

Loop—Repeats the same animation as within the range but interpolates between the last key and the first key of the range to create a smooth loop. Use Loop with an extended range bar to produce smoothly repeating animation.

Ping Pong—Alternates between forward and backward repetition of the animation within the range. Use Ping Pong when you want your animation to alternate back and forth.

Linear—Projects the animation value along lines tangent to the ease or multiplier curve at the limits of the range. Use Linear when you want the animation to enter and leave the range at a constant velocity.

Relative Repeat—Repeats the same animation as within the range but offsets each repetition by the value at the end of the range. Use Relative Repeat to create animations that build on each other as they repeat.

Freeze Non-Selected Curves

Track View > Curve Editor > Display menu > Show Non-Selected Curves > Freeze Non-Selected Curves

Freeze Non-Selected Curves makes all non-highlighted curves uneditable. Default=On. Freeze Non-Selected Curves is a selection modifier and has the following properties:

- Non-highlighted curves display as dashed lines.
- You can manipulate vertices on highlighted curves only in the Edit window.
- You highlight other curves by clicking their icons in the Hierarchy list.

  Turn Freeze Non-Selected Curves off when you want to be able to edit multiple curves without having to select them all. Turn Freeze Non-Selected Curves On when you have multiple vertices on multiple curves all in the same place.

Procedure

Example: To freeze non-selected curves:

1. Animate an object's position.
2. From the Graph Editors menu, choose Track View - Curve Editor. On the Options menu, make sure Auto Expand > Selected Objects Only/Transforms/XYZ Components are all on.

   The X/Y/Z position tracks are all highlighted, and their curves appear in the Key window. You can drag keys on any curve.
3. On the Display menu, make sure Hide Non-Selected Curves is on. This is the default setting.

4. In the Hierarchy list, click the X Position track icon.

   Now only the red X Position curve is visible.

5. From the Display menu, choose Show Non-Selected Curves.

   The other two curves reappear, and are editable.

6. From the Display menu, choose Freeze Non-Selected Curves.

   The Y Position and Z Position curves remain visible, but display as dashed lines.

7. Position the mouse cursor over a vertex on the Y Position or Z Position curve.

   The cursor doesn't change to a crosshair, indicating that you can't modify the vertex.

8. In the Hierarchy list, click the Y Position track icon and then Z Position track icon.

   All three remain visible, but in each case the two non-highlighted tracks display as dashed lines, and are not editable.

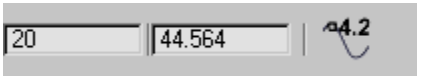
Status Bar/View Controls

Track View > Status Bar/View Controls

Use the Status Bar/View Controls to control the Track View Key window and view or enter frame and key values on the Track View status bar.



Track Selection: Track View toolbar



Key Stats: Track View toolbar



Navigation: Track View toolbar

Track Selection: Track View Toolbar

Zoom Selected Object (page 2–588)

Select By Name (Track View) (page 2–589)

Track Sets List (page 2–590)

Track Sets Editor Dialog (page 2–591)

Key Stats: Track View Toolbar

Key Time Display (page 2–594)

Value Display (page 2–594)

Show Selected Key Statistics (page 2–595)

Navigation: Track View Toolbar

Pan (Track View) (page 2–595)

Zoom Horizontal Extents (page 2–595)

Zoom Value Extents (page 2–596)

Zoom Track View Key Window (page 2–596)

Zoom Region (Track View) (page 2–597)

Zoom Selected Object 

With Track View already open to navigate to a particular object, select the object in the viewport. > Navigation: Track View toolbar > Zoom Selected Object

Use Zoom Selected Object to place the currently selected object at the top of the Hierarchy list in the controller window.

In Dope Sheet mode, Zoom Selected Object affects the Hierarchy list as well as the Edit window, but in the Curve Editor, it affects only the Hierarchy list. This function scrolls the hierarchy to place

the track of the currently selected object at the top of the list. If the object is a child of a closed hierarchy, the hierarchy is opened to display the selected object. If more than one object is selected in the scene, the uppermost selected object in the Hierarchy list is placed at the top of the list. If no objects are selected, nothing happens. This function works only on object tracks.

Procedure

To place the selected object at the top of the Hierarchy list:

1. Select an object in the viewports.
2.  Open Track View, and then click Zoom Selected Object.

   The object is placed at the top of the Hierarchy list.

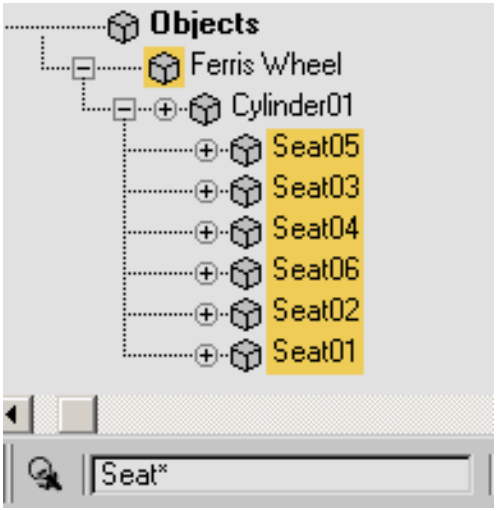
   This option is handy to quickly find an object in a hierarchy, such as a character's finger, and place it at the top of the Hierarchy list.

Select By Name (Track View)

Track View > Curve Editor or Dope Sheet > Track Selection: Track View toolbar > Select By Name field



With Select By Name you can highlight tracks in the Hierarchy window by entering the track names (including optional wildcards) in the editable field. If the Controller Type display is active in the Filters options, you can specify by controller type as well.



Select By Name using wildcard

Notes

- The name selection is not case-sensitive, except for text within parentheses.
- Only those tracks that are open and displayed can be selected.
- You can use wildcard characters. For example:

  Seat selects all objects with Seat in the beginning of the name.

  (Cone01)Position selects only the Position track under Cone01.

  (Cone)Position* selects the Position track under all objects that begin with "Cone"
- You can include multiple names in your selection if you separate them with commas; for example: *Cone,Box,Sphere*.

Procedures

To highlight all the tracks in a scene:

- Enter ***** in the Select by Name field in the Track Selection toolbar.

  All the tracks are selected at once.

Example: To highlight and display Position tracks for three objects:

1. Create three boxes.

   By default, the boxes are named Box01, Box02, and Box03.
2. On the main toolbar, click Curve Editor (Open).
3. In the Track View Hierarchy list, expand the hierarchies for the three box objects.
4. In the Select By Name field, type **(Box*)p**.

   All the Position tracks are highlighted.
5. On the Track View toolbar, right-click Filters and choose Selected Tracks Only in the dialog.

   The Position tracks for the three boxes are displayed at the top of the Hierarchy list.

Example: To highlight multiple tracks using comma-separated name selections:

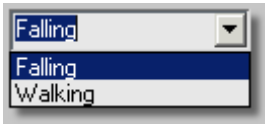
1. Create a box.
2. On the main toolbar, click Curve Editor (Open).
3. In the Track View Hierarchy list, expand box01's hierarchy to see all Transform children tracks.
4. In the Select By Name field, type **x pos,y rot,scale**.

   The box's X Position, Y Rotation, and Scale tracks are highlighted. The commas placed between each name act as Boolean addition ("OR") operators, allowing multiple names in one selection.

Track Sets List

Track View > Track Selection toolbar

You can assign a name to the current track selection, and then later reselect those tracks by choosing the name of the respective track set from the list.

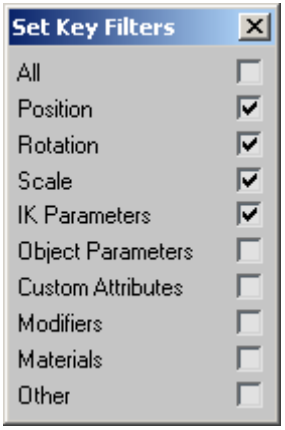


Track Sets List


You can also create and edit track sets with the *Track Sets Editor dialog (page 2–591)*.

Using Key Filters with Track Sets

If you are animating track sets using Key Filters in *Set Key mode (page 3–718)*, the filters dictate which track set members are keyable, because they take precedence over track sets.



Only the track set members specified in the Key Filters list are keyable.

Procedures

To assign a name to a track set:

1. Highlight one or more tracks in the Track View Controller window.
2. Click in the Track Set List field on the *Track View toolbar (page 2–535)*.
3. Enter a name for your track set. The name can contain any standard ASCII characters,

including letters, numerals, symbols, punctuation, and spaces.

Note: Names are case-sensitive.

4. Press Enter to complete the track set.

   You can now highlight another combination of tracks and repeat the process to create another track set.

To retrieve a track set:

1. In the Track Sets List, click the arrow.
2. On the list, click a name.

   The respective tracks are highlighted in the Track View Controller window.

To edit a track selection:

You can change a track selection using track sets and the Ctrl and Alt keys.

1. Highlight one or more tracks in the Track View Controller window.
2. Open the Track Set List by clicking the arrow.
3. Press and hold Ctrl, and then click a name in the list.

   The track set members are added to the current track selection.

4. Press and hold Alt, and then click a name (either the same or another) in the list.

   The track set members are subtracted from the current track selection.

To edit track sets;

-  On the Track View toolbar, click Edit Track Set to display the *Track Sets Editor dialog (page 2–591)*.

Track Sets Editor Dialog

Curve Editor Track Selection toolbar > Edit Track Set

The Track Sets Editor dialog is a *modeless dialog (page 3–973)* that lets you create and edit animation track groups called track sets. This feature facilitates working on multiple tracks at a time, because you can recall your tracks without the need to select them individually.

Tip: As your projects grow in size, you can also use track sets to help organize your animations by regrouping them under reference dates or versions such as *Walk_Animation_v5.0* or *Baseball_Hit_2005/07/01*.

Within the Track Sets Editor dialog, you can browse track set listings, add or remove tracks to/from a track set, highlight track set members based on selections from the Hierarchy list, and vice-versa.

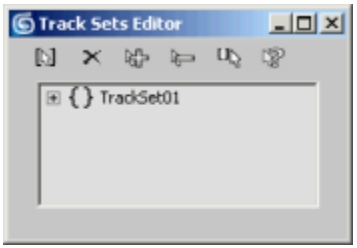
See also

Track Sets List (page 2–590)

Procedures

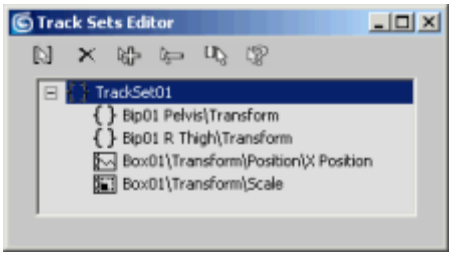
To create a track set from selected tracks:

1. In the Track View Controller window, highlight tracks to group as a track set.
2.  Click Edit Track Set.
3.  In the Track Sets Editor dialog, click Create A New Track Set.



4. Expand the new track set.

   A new track set appears containing your selected tracks.



To edit the contents of a track set:

This procedure follows from the previous procedure, and focuses on adding and removing tracks to/from your track set.

1. In the Track View Controller window, highlight a few tracks to add to your track set. Make sure the tracks are not already part of it.
2.  In the Track Sets Editor dialog, highlight your track set or any of its members, and then click Add The Track View Selection To The Current Track Set.

   The highlighted tracks are now members of your track set.
3.  Click Remove The Track View Selection From The Current Track Set.

   The highlighted tracks are now removed from the track set.

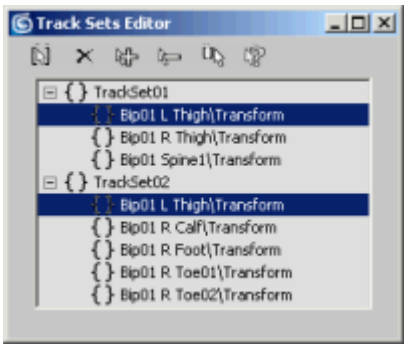
To locate tracks in the Track View or the Track Sets Editor:

This procedure is a continuation of the previous procedures and explains how you can use the Track Sets Editor to quickly locate your tracks in the Hierarchy list and vice-versa. This is useful for managing scenes comprised of many tracks that span across multiple track sets.

1. Create two track sets each containing tracks. Make sure one track is common to both track sets.
2. In the Track View Controller window, highlight the common track.

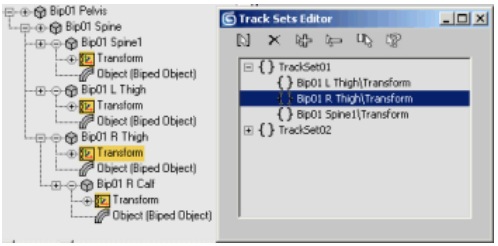
    In the Track Sets Editor dialog, click Highlight Tracks Selected In The Track View.

   The common track is highlighted in both track sets.

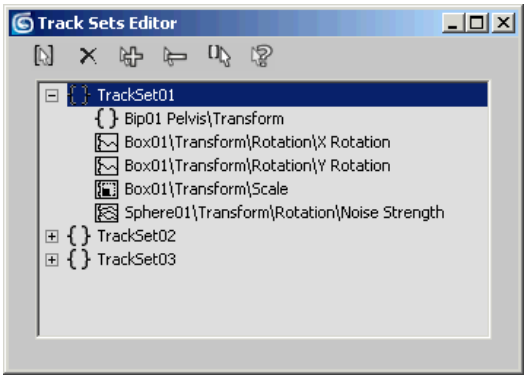


3.  Highlight another track set member and click Select The Current Track Or Track Set In The Track View.

   The respective track in the Controller window is highlighted.



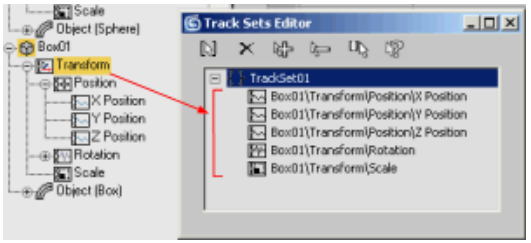
Interface



The Track Sets Editor dialog window displays all current track sets. To expand or collapse the track list for a set, click the plus (+) or minus (-), button next to the track set name.

Create A New Track Set—Adds a new track set comprising any highlighted tracks in the Hierarchy list.

Only *leaf (page 3–951)* tracks can be part of a track set. If you highlight a non-leaf parent track when the track set is created, all its children leaf tracks are added instead.



A track set created from a non-leaf parent track contains all children leaf tracks.

Note: If no tracks are highlighted, an empty set is created.

Delete Track Sets Or Tracks—Removes a highlighted track set or track set member.

Note: Deleting a track set member does not delete it in the Track View.

Add The Track View Selection To The Current Track Set—Combines a highlighted track from the Controller window to the current track set.

Only leaf tracks can be added to a track set. If you highlight a non-leaf parent track to be added to the track set, all its children leaf tracks are added instead.

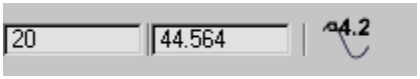
Remove The Track View Selection From The Current Track Set—Deletes a track set member from the current track set. The track itself must be highlighted in the Controller window.

Highlight Tracks Selected In The Track View—Highlights a track set member based on the current Track View selection. If a track is shared between track sets, all track set members are highlighted.

Select The Current Track Or Track Set In The Track View—Highlights a track in the Controller window based on the current track set member selection.

Key Time Display

Track View > Curve Editor or Dope Sheet > Key Stats:
Track View toolbar > Key Time Display (field on left)



The Key Time editable field (the field on the left) displays the frame number (position in time) of the selected key. You can enter a new frame number or enter an expression to move the key to a frame.

Using Expressions in the Time and Value Fields

You can enter any expression in the time and value fields. The variable *n* stands for the key time or value. An expression can be as complex as you want, and can involve any function the *expression controller (page 2–320)* uses. For example, you could enter an expression such as: n/2+sin(n)*40^n.

The expression parsing in the time field works only when *Time Configuration (page 3–725)* > Time Display is set to Frames. The expression parser doesn't understand SMPTE or Frame:Ticks.

Procedures

To move a key to a specific frame:

1. In Track View, highlight a key.
2. In the Key Time Display, enter the frame number you want to move the key to.
3. Press Enter.

   The key (and, in Dope Sheet mode, any associated keys) moves to the new time.

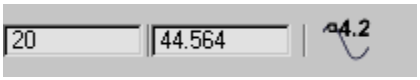
To move keys ahead 10 frames:

1. In the Track View Key window, highlight keys.
2. In the Key Time field, type **n+10** .

   Selected keys are incremented by 10 frames.

Value Display

Track View > Curve Editor or Dope Sheet > Key Stats:
Track View toolbar > Value Display (field on right)



Value Display (the field on the right) displays the value, or position in space, of a highlighted key. This is an editable field. You can enter a new number or expression to change the value of selected keys.

Procedures

Example: To move an object in space using Value Display:

1. In the Track View (either mode) Hierarchy list, expand the Position track of an animated object. In Curve Editor mode, if Show Non-Selected Tracks is not active, highlight a track, if necessary, so the keys are visible.
2. Highlight an X, Y, or Z Position key.

   The Value Display shows the X, Y, or Z position key spatial value.
3. Enter the value you want in the Value Display and press Enter.

   The new value takes effect. If the time slider is at or near the same frame as the key, the change is visible in the viewports.

Example: To add 10 to the values of specific keys:

1. In Track View, either mode, highlight the keys to edit.
2. In the Value Display field, enter **n+10**.

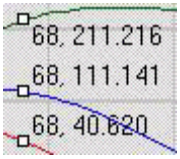
   Each key's value is incremented by 10. Simply entering 10 would give each key a value of 10.

Show Selected Key Statistics 

Track View > Curve Editor > Key Stats: Track View toolbar > Show Selected Key Stats

Show Selected Key Statistics displays the statistics represented by the currently selected keys in the Key window of the Curve Editor.

The frame number and value are displayed to the right of the key in the Key window. For example, 68, 40.620 (frame=68, value=40.620).



Procedure

To show selected key statistics:

1. On the Key window of the Curve Editor, select any key or set of keys.
2. On the Key Stats: Track View toolbar, click Show Selected Key Stats.

   The key statistics are displayed next to each key.

Pan (Track View) 

Track View > Curve Editor or Dope Sheet > Navigation: Track View toolbar > Pan

With Pan you can click and drag the Key window to move it left, right, up or down. Pan remains active until you right-click to cancel or click another option. The Pan button is yellow while active.

If you have a middle mouse button, pressing it will instantly give you the Pan capability in Track View, as in the rest of the software.

Drag horizontally in all modes to slide the view forward and backward in time. Drag vertically to slide the view up and down in value.

Zooming and Panning with a wheel mouse

You can use the wheel on a wheel mouse to zoom and pan the Track View Key window. Drag with the wheel to pan, and roll the wheel to zoom. (If you have a standard three-button mouse, use Ctrl+Alt+middle mouse button to zoom.)

Procedure

To pan in the Key window:

1.  On the Navigation: Track View toolbar, click Pan.
2. Position the mouse cursor over the Key window, and then press and drag to pan.

Zoom Horizontal Extents  

Track View > Curve Editor or Dope Sheet > Navigation: Track View toolbar > Zoom Horizontal Extents

Zoom Horizontal Extents is a flyout containing the Zoom Horizontal Extents and Zoom Horizontal Extents Keys buttons.

Use Zoom Horizontal Extents to show only the *active time segment (page 3–904)*, and use Zoom Horizontal Extents Keys to show all keys, including those outside the active time segment.

Procedure

To display the active time segment in the Key window:

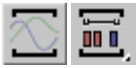
-  On the Track View status bar, select Zoom Horizontal Extents.

The active time segment (lighter background color) is centered in the Key window.

Interface


Zoom Horizontal Extents—Adjusts the magnification of the Track View Key window horizontally so that all of the active time segment is visible at once.


Zoom Horizontal Extents Keys—Scales the Track View Key window horizontally to display the full range of all animation keys. Depending on your animation, this view could be larger or smaller than the active time segment.

Zoom Value Extents 

Track View > Curve Editor > Navigation: Track View toolbar > Zoom Value Extents

Zoom Value Extents is a flyout containing the Zoom Value Extents and Zoom Value Extents Range buttons.

In Function Curves mode, Zoom Value Extents adjusts the size of the Track View Key window vertically so that you can see the full height of the curve. Zoom Value Extents Range also adjusts the window vertically but it zooms only to the height of the keyframes that are currently in view.

Note: In Dope Sheet mode, Zoom Value Extents and Zoom Value Extents Range serve no function.

Procedures

To zoom to the height of the curves:

1. In the Hierarchy list, highlight the tracks to display.
2. 
On the Track View status bar, click Zoom Value Extents.

   The Key window scales vertically to display the function curves.

   This is particularly useful if you are zoomed in and need to see all the keys again. There is no undo for view changes that occur in the Key window, so the Zoom Extents button can serve the function of resetting your view.

To zoom to the height of the keyframes that are currently visible:

1. In the Hierarchy list, highlight the tracks to display.
2. 
On the Track View status bar, click Zoom Value Extents Range.

   The Key window scales vertically to the height of the keyframes in view.

   This is useful if you want to focus on working with one section of your animation curve.

Zoom Track View Key Window   

Main toolbar > Curve Editor or Dope Sheet > Navigation: Track View toolbar > Zoom

In Track View, the Zoom controls are available from a three-button flyout. You can zoom your view of time horizontally (Zoom Time), vertically (Zoom Values), or in both directions simultaneously (Zoom). Drag to the right or upward (Zoom Values) to increase magnification, or to the left or downward (Zoom Values) to decrease magnification.

Zooming occurs around the cursor location.

Note: All three zoom modes work in Curve Editor, but only Zoom and Zoom Time work in Dope Sheet.

Tip: If you have a three-button mouse, you can also zoom uniformly by using Ctrl+Alt+middle mouse button or by scrolling with your mouse wheel in the Curve Editor.

Tip: You can also zoom exclusively in time (left and right) or in value (up and down) by holding down Ctrl or Shift respectively while scrolling with your mouse wheel in the Curve Editor.

Procedure

To zoom in on time:

1.  On the Navigation: Track View toolbar, choose Zoom Time from the flyout.
2. Click and drag to the right in the Track View Key window.

   Your view of time expands horizontally in the Key window.

Interface

 **Zoom**—Zooms your view of time and values simultaneously. Zoom remains active until you right-click to cancel or click another option. The Zoom button is yellow while active.

 **Zoom Values**—In Curve Editor, zooms the contents of the keys window vertically. Drag upward to increase magnification, or downward to decrease magnification.

 **Zoom Time**—Zooms the contents of the Key window horizontally.

Zoom Region (Track View) 

Track View > Curve Editor or Dope Sheet > Navigation: Track View toolbar > Zoom Region

Zoom Region lets you drag a region in the Key window to scale that region to fill the window. Zoom Region remains active until you right-click to cancel or click another option. The Zoom Region button is yellow while active.

In Curve Editor mode, both time and value are scaled to fit the Key window. In Dope Sheet mode, only time is scaled to fit the Key window.

Procedure

To zoom a region of the Key window:

1.  On the Navigation: Track View toolbar, click Zoom Region.
2. Drag a window in the Key window.

   The region zooms to fill the Key window.

   Tip: There is no undo for view changes in the Track View windows. Use the Zoom, Zoom Region, and Zoom Extents tools to navigate the Key window.

Managing Track View Windows

New Track View

Graph Editors menu > Track View > New Track View

New Track View opens a new, untitled *Track View (page 2–501)* window. A scene in 3ds Max can have up to 13 Track View windows. You can choose

these by name from the Graph Editors > *Saved Track Views (page 2–599)* submenu. New Track View is disabled if the scene already has 13 Track View windows.

You can name a Track View window by typing in the editable field at the right end of the Track View toolbar.

Procedures

To create a new Track View window:

- Click Track View > New Track View.

  The software opens a new, untitled Track View window, regardless of whether you have closed other Track View windows.

To name a Track View window:

- Type the name in the edit box at the right end of the Track View toolbar.

  As soon as you type the name in, it is stored in the Saved Track Views list.

Delete Track View

Graph Editors menu > Delete Track View

Delete Track View displays the *Delete Track View dialog (page 2–598)*, which lets you delete one or more stored *Track View windows (page 2–501)*.

Procedure

To delete a Track View window:

1. From the Graph Editors menu, choose Delete Track View.

   The Delete Track View dialog opens.
2. In the Delete Track View dialog, highlight one or more Track View windows.
3. Click OK to delete the Track View windows you chose.

Delete Track View Dialog

Graph Editors menu > Delete Track View

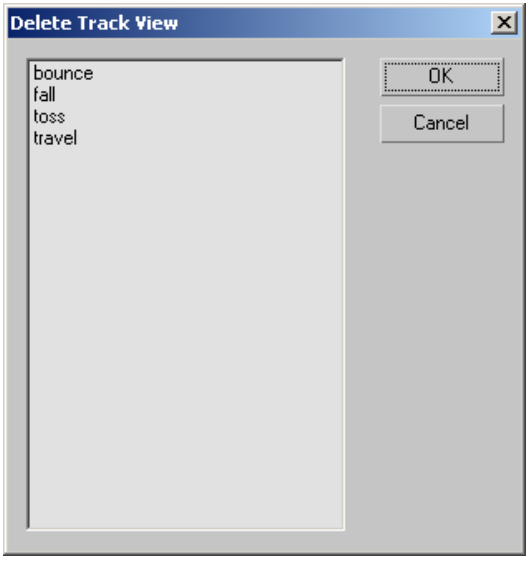
The Delete Track View dialog lets you delete one or more stored *Track View (page 2–501)* windows by choosing their names from a list.

Procedure

To delete a Track View window:

1. Choose Track View > Delete Track View.
2. In the Delete Track View dialog, choose the name of one or more Track View windows.
3. Click OK to delete the highlighted Track View windows.

Interface



Drag a selection, use Shift+click to select additional adjacent Track View windows, or Ctrl+click to toggle a selection.

Saved Track Views

Graph Editors menu > Saved Track Views > Choose the named Track View from the list.

Saved Track View lets you recall various Track View windows that you save by naming them.

It is different from saved Track View layouts, which are recalled by right-clicking the toolbar and choosing Load layout.

Saved Track View windows are saved with the MAX file.

Procedure

To save a Track View window:

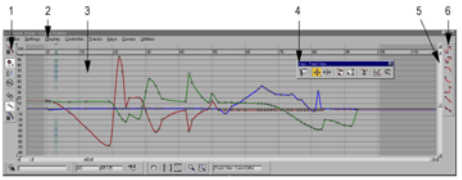
1. In either Dope Sheet or Curve Editor, enter a name in the Name: Track View field at the upper right.

   Tip: If the Name field is not visible, right-click the toolbar and choose Show Toolbars, then choose Name: Track View from the list. This toolbar can float or dock, as you like.

2. Close Track View.
3. To recall the saved Track View window, from the Graph Editors menu, choose Saved Track Views, and then the desired window name.

Track View Customization

The Track View user interface is fully customizable. There are many ways you can change the UI to suit your needs.



A customized Track View user interface:

1. *Toolbar docked left*
2. *Menu bar*
3. *Key window (Hierarchy list hidden)*
4. *Floating toolbar*
5. *Scroll bar*
6. *Toolbar docked right*

You can hide or display UI elements as you need to use them. Customization tools are available from the Track View menu bar or from several Track View toolbars. You can float or dock the toolbars to the top, bottom, left, or right of the application window, for example, floating a toolbar so it is closer to where you are working.

To customize the Track View UI, right-click the menu bar or the blank space to the right of the Track View toolbars. A right-click menu is displayed, with commands for the tools you need to perform your customization.

Track View Toolbar Right-Click Menu

You can use commands on the *toolbar right-click menu (page 2–603)* to perform most Track View UI customization tasks. Move your cursor over the blank space to the right of the Track View toolbars and right-click to display this menu.

Show UI Elements

Toggles on and off the display of Track View user interface elements. For example, you can hide or

show the menu bar, scroll bars, the controller and keys windows, and the time ruler.

Menu Bar

The menu bar gives access to most of the tools also found on the toolbars. The menu bar can be displayed or hidden in the Curve Editor, Dope Sheet, or track bar. The menu bar is context-sensitive, so it changes its display depending on which Track View mode is active.

Scroll Bar

Standard Windows scroll bars can be displayed or hidden for the Track View window. These are useful for navigating complex or lengthy windows.

 Use the Pan hand button to scroll when these are not visible. The middle mouse button can also be used to activate the pan hand to scroll this window.

Windows

There are two side-by-side display areas in Track View. They are the Controller window and the Key window. The Controller window contains the Hierarchy list with labeled tracks. The Key window displays the keys either as curves in the Curve Editor, or as boxes or ranges on a grid in Dope Sheet mode. Use Show UI Elements, available from the toolbar right-click menu, to hide either window.

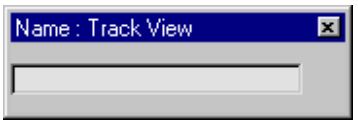
Time Ruler

A movable time ruler is provided. To measure time, you can move the time ruler close to the keys or curves when working zoomed in. Default=on.

Track View - Curve Editor Toolbars

You can display *Curve Editor toolbars (page 2–535)* for quick access to tools.

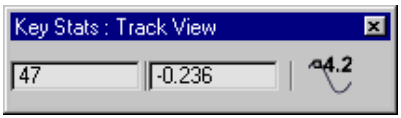
The Curve Editor toolbars that you can display or hide, dock or float, from the right-click menu include the following:

- 

Name toolbar—Enter a name in this field to create a named Track View window. You can recall the saved Track View window by choosing Graph Editors menu > *Saved Track View (page 2–599)* submenu. This option is also available in Dope Sheet mode.

- 

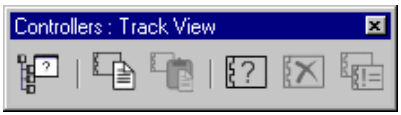
Navigation toolbar—Use these tools to *pan (page 2–595)* and *zoom the Key window (page 2–596)*, or *zoom to the extents of the time (page 2–595)* or *the values (page 2–596)*. This option is available in both Track View modes.

- 

Key Stats toolbar—Track View provides *tools for the display and type-in transform of key values (page 2–594)*. This option is also available in Dope Sheet mode.

- 

Key Tangents toolbar—Track View gives you *quick access to different tangency types of keys (page 2–536)*. This toolbar is most useful when floating near the curves. This option is also available in Dope Sheet mode.

- 

Controllers—Track View has *tools to copy, paste,*

assign, delete, and make controllers unique (page 2–540). This option is also available in Dope Sheet mode.

- 

Track Selection—Track View gives you *tools to select tracks by typing in the name of an object (page 2–588)*. This option is also available in Dope Sheet mode.

- 

Keys—Track View provides the tools to *move (page 2–579)*, *scale (page 2–580)*, and *slide keys (page 2–558)*, *scale values (page 2–581)*, *add keys (page 2–581)*, *draw curves (page 2–582)*, and *reduce keys (page 2–572)*. This toolbar also has the Filter button, which lets you limit the display in the Controller window. A *similar option is available in Dope Sheet mode (page 2–538)* with different and fewer controls.

- 

Curves—Track View provides *tools to select and lock keys and curve handles (page 2–536)*, as well as *to make tracks keyable (page 2–531)*, and to apply *parameter curve out-of-range types (page 2–551)*.

- 

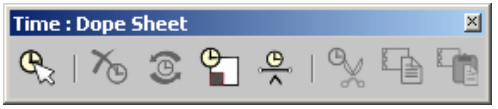
Tools—Track View gives you tools to create or remove Note and visibility tracks, snap frames, lock selection, properties, and Track View utilities. This toolbar is hidden by default.

Dope Sheet Toolbars

There are 11 toolbars that can be displayed when in Dope Sheet mode. In addition to the ones listed above for the Curve Editor, you can also choose to display the following:

- 

Keys—Dope Sheet displays *the tools to move, add, and scale keys (page 2–538)*, *the Filters button (page 2–541)*, and the buttons that let you choose between *Edit Keys (page 2–554)* and *Edit Ranges (page 2–573)* options for Dope Sheet.

- 

Time—Track View provides *tools to select and work with time (page 2–539)*. Time can be selected, inserted, cut, copied, pasted, scaled, reversed, or removed. This toolbar works only in Dope Sheet mode.

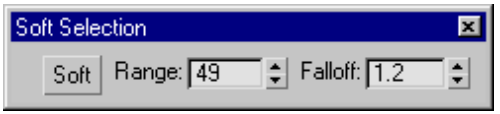
- 

Ranges—Here are the tools to *edit (page 2–573)*, *position (page 2–574)*, and *recouple ranges (page 2–574)*.

- 

Display—Dope Sheet holds tools to *modify subtree (page 2–528)* and *modify child keys (page 2–529)*, *lock selection (page 2–555)*, *snap frames (page 2–554)*, and *display keyable icons (page 2–531)*.

Soft Selection Toolbar



Both the Curve Editor and Dope Sheet Edit Keys mode take advantage of soft selection of keys. If

you choose Keys menu > Use Soft Select, then choose *Soft Selection Settings (page 2–525)*, the Soft Selection Settings toolbar is displayed. This toolbar is accessible only from the Keys menu, not from the right-click menu. If you want this menu to stay visible, *save a Track View layout (page 2–602)*.

Track Bar

You can dock the Track View windows below the time slider and track bar by choosing Dock > Bottom from the toolbar right-click menu. You can dock multiple Track View windows using this technique.

 You can also display a Track View window below the viewports by clicking the Open Mini-Curve Editor button at the upper left corner of the track bar. Using this method, the time slider and track bar are replaced with the Track View window. As with Dope Sheet and Curve Editor modes, you use the toolbar right-click menu to display or hide all Track View UI elements.

Procedures

To save a Track View layout:

1. Customize the Track View layout as you want.
2. When you want to save it, right-click the Track View menu bar or the blank area at the right of the Track View toolbar, and choose Save Layout As.
3. Enter a name you will remember.

To reload a saved Track View layout:

1. Right-click the Track View menu bar or the blank area at the right of the Track View toolbar and click Load Layout.
2. Choose the layout you want to display in Track View.

To save a Track View window:

You can save a particular Track View window, displaying to a particular set of tracks, by using the Name: Track View toolbar.

- Enter the name for the window in the Name field. When you close the window, you can reload it by choosing Graph Editors > Saved Track Views. Then choose the named window from the list.

  Tip: Named Track View *windows* are saved with a file. *Named layouts* are independent of the file.

To restore the default Track View user interface:

If you have customized your Track View user interface and want to return to the original installed version, do the following:

1. In the *\Program Files\Autodesk\3ds Max 9\ui* folder, locate the file *trackview.bak*.
2. Copy the file and rename the copy **trackview.ini**.
3. In Windows Explorer or My Computer, right-click and choose Properties. Make sure the Attributes is not set to Read-only.

   Tip: You can also right-click the Track View title bar and choose Load Layout > Default to return to the default layout. This works only if you have not saved over the default.

To add a new menu bar name to a custom Track View layout:

1. From the Customize menu, choose Customize User Interface.
2. On the Menus panel, click New to create a new menu. Enter the name of the menu in the New Menu field and click OK.
3. Using Notepad, open *ui/trackview.ini*, and navigate to the layout section. To affect the Curve Editor toolbar, choose the section named *[default]*, or to affect the Dope Sheet toolbar, choose the section called *[Dope Sheet layout]*.

4. Change the value of the "Menu Name" entry to the name of the menu you created in step 2.

   Save the layout after you are done and the changes will be persistent.

   Tip: You can also make the Dope Sheet Editor point at a renamed menu bar by using MAXScript:

   In the MAXScript Listener, type**trackviews.current.ui.menubar = "Dope Sheet - Menu Bar"**, substituting your menu bar name, and then press the Enter key.

   Note: There is also a script that lets you choose menus from a list. Choose Customize menu > Customize User Interface, then on the Menus panel, set Group to Main UI and Category to Track View, and choose Load Menu Bar from the Action list. Select the Track View menu bar you want to customize from the drop-down list on the right. Drag the action to the right window and the menu bar will be customized.

Track View Toolbar Right-Click Menu

You can use commands on the right-click toolbar menu to perform most Track View UI customization tasks. Move your cursor over the blank space to the right of the Track View toolbars and right-click. This shortcut menu's customization tools are as follows:

- **Dock—**Snaps the toolbars into position horizontally above or below the window, or vertically along the left or right window edge. Once toolbars are horizontally docked, they can be reordered by dragging and dropping them as you like. Also used to dock the window beneath the viewports.
- **Float—**Floats the selected toolbar. You can float toolbars anywhere you like. You can also drag the double bars at the left of the toolbar to float the toolbar over the viewport.
- **Hide—**Turns off the display of the selected user interface element. Once something is hidden, you can redisplay it using Show UI Element or Show Toolbar on this menu.
- **Load Layout—**Loads a saved default, named, Function Curve, Dope Sheet, or track bar layout.
- **Save Layout—**Saves your customized layout, creating a new default layout. When you open a new Track View, this saved layout will be displayed.

  Note: To restore the default 3ds Max layout, see *To save a Track View layout: (page 2–602)*.
- **Save Layout As—**Saves your customized layout with a name you choose. When you've customized your layout, save it with an easily recognizable name, then you can return to it whenever you choose. Use Load layout to retrieve a saved layout.
- **Show UI Elements—**Toggles the display of Track View user interface elements. For example, you can hide or show the menu bar, scroll bars, the controller and Key windows, and the time ruler.
- **Add Toolbar—**Lets you display additional toolbars that you can create using commands the Customize User Interface pull down from the main menu bar.
- **Delete Toolbar—**Deletes the selected toolbar.
- **Show Toolbars—**Lets you choose which toolbars you want to display or hide.

  Note: Some toolbars are hidden by default.

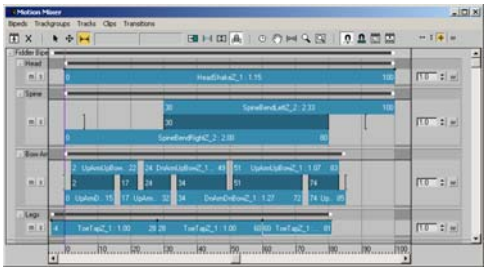
Motion Mixer

The Motion Mixer allows you to combine motion data for biped and non-biped objects.



The Motion Mixer takes its design from the world of audio. When a song is recorded in a studio, each instrument is played and recorded separately. Each recording is called a *track*. The tracks are then put together in a sound mixer so they play simultaneously, or overlap one another. During the mixing process, the mixer operator can change the length or speed of a track, increase or reduce volume, move a track to another place in the song, or cause a track to fade in or out.

The Motion Mixer works in a similar way. For any object, you can add multiple *tracks (page 3–1023)* to the mixer, each holding a separate series of motion clips (*BIP files (page 3–916)*, XAF files). You can trim clips to use only part of a motion, make the clips play slower or faster, or create transitions from one clip or set of clips to another.



You can also use the Motion Mixer to animate some body parts with one set of clips, and other body parts with other motions. For example, suppose you have two clips, one where the biped runs with its arms pumping by its sides, and another where the biped stands and cheers with its arms in the air. You can mix the leg and hip motions from the running motion with the arms from the cheering motion to make an animation of a biped cheering as he runs across a finish line.

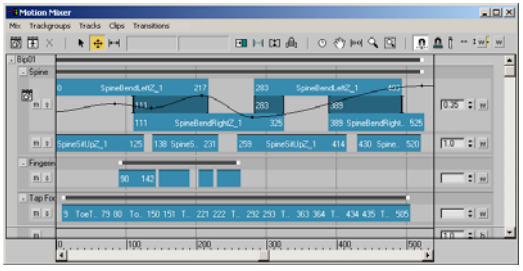
See also

Using the Motion Mixer (page 2–604)

Motion Mixer Interface (page 2–628)

Using the Motion Mixer

The Motion Mixer allows you to mix motion files (*BIP (page 3–916)* and *XAF (page 3–476)* files) for any biped or non-biped objects. These motion files are also called *clips*.



The Motion Mixer is comparable to an audio mixer. Motion clips can be cross-faded, stretched, layered, and finally, mixed down to a single clip.

The Motion Mixer works by placing motion files on *tracks (page 3–1023)*. With the Motion Mixer, you can:

- Transition or fade between motions.
- Move motions in time.
- Trim a motion so only part of it is used.
- Vary the speed of a motion over time.
- Use animation from selected biped or non-biped body parts within a motion clip.
- Keep planted feet from sliding during foot-based transitions.

 To see a biped motion in the Mixer applied to the biped in the scene, you must turn on Mixer Mode from the *Biped rollout (page 2–936)*.

In the Motion Mixer, motions are imported and mixed for any objects in the scene. Once motions are mixed, you can pass the mix on to other biped (if it contains BIP files) and non-biped objects (if it contains XAF files), or save the mix to a MIX file.

Blending Motions in the Mixer

The Motion Mixer provides these types of motion blending.

- You can blend motion from one clip to another with a *transition (page 3–1027)*. Transitions in the Mixer are similar to those used in the *motion flow system (page 2–1026)*. The optimization feature can automatically find the best timing for a transition between two clips. When you make a transition between foot-based clips (clips where IK constraints keep the feet planted at certain times), you can cause the transition to focus on one foot or the other. See *Working with Transitions (page 2–616)*.
- In biped animation, if a transition between foot-based clips causes the foot to slide or pop slightly during the transition, you can fix this problem with a *mixdown (page 3–973)* (flattening all tracks into one track). See *Exporting Animation to the Biped (page 2–624)*.
- You can change the *weight (page 2–619)* of clips or tracks so they affect the animation more at some times than others. See *Adjusting Track Weight (page 2–619)*.
- You can use only the part of a clip's motion that affects specified body parts. See *Filtering Mixer Tracks (page 2–612)*.
- When motions applied to a biped's upper body vary greatly from motions on its lower body, the Mixer automatically compensates for discrepancies in balance. See *Adjusting Biped Balance in the Mixer (page 2–622)*.

  Note: Balance parameters are available only when mixing biped objects.

Transitions Between Foot-Based Clips

The Motion Mixer excels in its ability to maintain smoothness of motion over transitions between foot-based clips. Foot-based clips use IK constraints to keep one or both feet locked down over the course of the motion.

For bipeds, the Mixer provides tools for maintaining foot positions during transitions:

- Computing the influence of balance from the upper to the lower body. Without this, layered motion tracks will appear to be artificially combined since out since the resulting motion will be plagued by the isolation of dynamics and improper balance.
- Mixdowns allow for precise elimination of the foot sliding with the aid of knee limit filtering to unnatural knee poppimg artifacts. In short, the goal of the mixer is to assemble seamless, coherent character motion using

building blocks composed from other motion fragments.

For non-biped rigs and other 3ds Max objects, the Mixer provides velocity blends and sub-blends to maintain smooth and coordinated motion.

Comparing the Motion Mixer and Motion Flow

Both the Motion Mixer and *Motion Flow (page 2–1026)* can be used to create transitions between a consecutive series of BIP files. However, there are a few significant differences between the two features:

- With the Motion Mixer, several BIP files can be used at the same time to affect biped motion. Compare with Motion Flow, where only one BIP file, or a transition between two consecutive BIP files, affects the biped at any given time.
- With the Motion Mixer, you can specify that a BIP file will affect only one part of the body. This feature is not available in Motion Flow.
- In the Motion Mixer, you can use a mixdown to automatically lock feet during transitions between biped foot-based motions. Motion Flow does not perform this function.
- A Motion Flow network can be used in conjunction with the Crowd feature, giving the bipeds a choice of motions over the course of the simulation. The Motion Mixer does not generate motion for biped crowds.
- A Motion Flow network can be used to randomly generate different motion scripts for one or more bipeds. This feature is not available in the Motion Mixer.

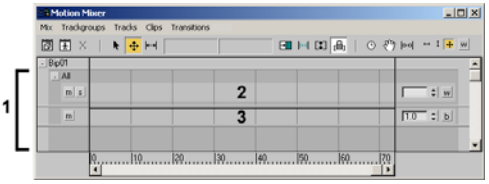
Use the Motion Mixer when you want to create a specific animation for an object from several clips, or you want to specify that some motions be applied only to specific body parts. If you want to work with biped crowds or create random motion, use Motion Flow.

If you're creating a straight series of transitions between motions, you can use either one. However, the Motion Mixer has additional tools for improving foot-based transitions.

Opening the Motion Mixer

You open the Mixer by either clicking the Mixer button on the Motion panel > *Biped Apps rollout (page 2–935)* or choosing Motion Mixer from the *Graph Editors menu (page 3–682)*.

When you open the Mixer while a biped is selected, a *trackgroup (page 3–1024)*, a clip track, and a *balance track (page 3–914)* are assigned to that biped. If a non-biped object is selected, the Mixer opens empty.



1. Trackgroup

2. Clip track

3. Balance track

These items are discussed in detail in the sections that follow.

Procedures

To access the Motion Mixer from the Motion panel:

1. Select a biped.
2.  Access the Motion panel.
3. On the *Biped Apps rollout (page 2–935)*, click the Mixer button.

   The Motion Mixer opens, and displays the selected biped and two tracks. To learn about these tracks, see *Adding Tracks to the Mixer (page 2–607)*.


Accessing the Mixer also turns on Mixer mode for the biped. Mixer mode must be on for the biped in order to see the mixed motion on the biped.

To display the Motion Mixer in a viewport:

1. Activate the viewport in which you want to display the Motion Mixer.
2. Right-click the viewport label, and choose Views > Extended > Motion Mixer from the pop-up menu.

   Note: When the Motion Mixer is displayed in a viewport, its menu bar is not accessible. However, you can access each menu by right-clicking the corresponding element in the Mixer.

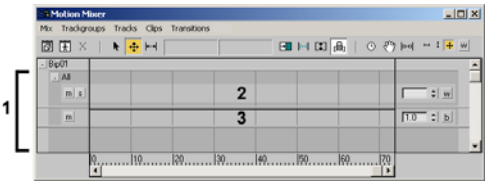
To remove the Motion Mixer from a viewport:

- Right-click a blank area of the Motion Mixer toolbar, and choose another view from the selections under Views on the pop-up menu.

Adding Tracks to the Mixer

In the Motion Mixer, you begin the mixing process by opening the Mixer, then adding objects to mix, and then loading motion clips to clip *tracks (page 3–1023)*. One clip track is created automatically if a biped is selected while you open the Mixer, but you'll probably need more to hold all the motions you want to use with your mixed objects.

When you open the Mixer while a biped is selected, a *trackgroup (page 3–1024)*, a clip track, and a *balance track (page 3–914)* are assigned to that biped. If a non-biped object is selected, the Mixer opens empty.



1. Trackgroup

2. Clip track

3. Balance track

This section discusses clip tracks. To learn about trackgroups, see *Filtering Mixer Tracks (page 2–612)*. For information on the balance track, see *Adjusting Track Weight (page 2–619)*.

Adding and Merging Non-Biped Object to a Mix

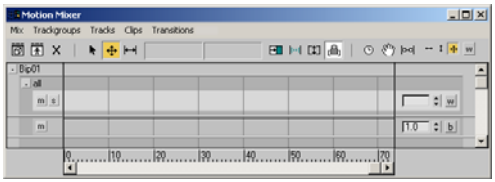
When you add a non-biped object to a mix, its keyable controllers are replaced by mixer controllers, which store the object's original animation as well as the Mixer animation.

If you merge a non-biped object containing Mixer controllers into a new scene, it comes in with its mix name preserved from the original file, but without any mixed animation.

Types of Clip Tracks

There are two types of tracks to which you can add clips:

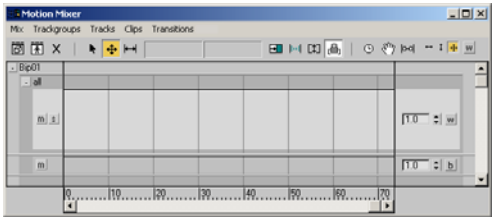
Layer—A track for a series of motions that do not need transitions between them. By default, a Layer track is created for a biped when you open the Mixer.



Layer track

Transition—A tall track with room for two rows of clips. You can put one motion file on the top

tier and another on the bottom, and place a transition in the area where the two clips overlap. Transitions on these tracks are similar to *Motion Flow transitions (page 2–1034)* between clips.



Transition track

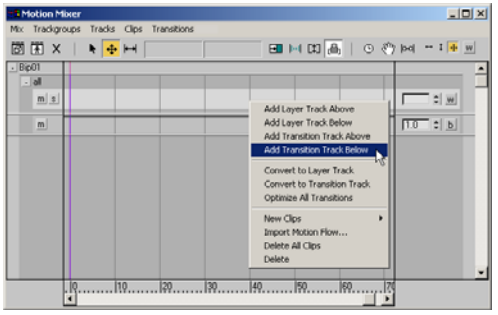
Each biped in the Mixer is automatically given a *balance track (page 3–914)*, used for adjusting balance between clips in Layer and Transition tracks. Clips cannot be added to the balance track.

Note: Balance tracks are only available for mixed biped objects.

Procedures

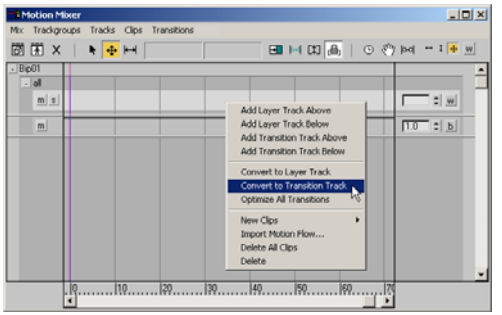
To add a Transition or Layer track to a trackgroup:

- Right-click an existing track and choose one of the options for adding tracks from the pop-up menu.



To convert a clip track from one type to another:

- Right-click an existing track, and choose one of the options from the pop-up menu for converting tracks.



To add a biped to the Motion Mixer:

1.  On the Motion Mixer toolbar, click Add Biped.
2. From the Bipeds dialog, choose the biped(s) you want to add to the Motion Mixer.

   Several bipeds can be added to the mixer at once. When you click OK, highlighted bipeds are added to the mixer. Each is automatically assigned a trackgroup, a clip track, and a balance track.

To add a non-biped object to the Motion Mixer:

1.  On the Motion Mixer toolbar, click Add Max Objects.
2. From the Max Objects To Mix dialog, choose the object(s) you want to add to the Motion Mixer.

   Tip: You can save time if you first select your objects, and then open the dialog. This way, the objects are automatically highlighted.
3. Enter a mix name or leave the default one.

   When you click OK, the highlighted objects are added to the mixer under a single object mix . A track is automatically assigned to the object.

Importing Clips to the Mixer

You can bring motion clips (BIP and XAF files) into the mixer with any of these methods:

- Choose a file directly from the folder in which it resides.
- Choose a file from those you have placed in the *Reservoir (page 2–626)*, a storage area for BIP and XAF files.
- Import motion from another biped in the scene. See *Working with Clips in the Mixer (page 2–611)*.
- Import a motion flow script.

All clips imported for a single object are assigned a random, unique color. For example, all clips for one biped might be red, while those for another biped are blue. The color distinction makes it easier to work with multiple bipeds and non-biped objects in the Motion Mixer.


Biped rollout > Mixer Mode must be turned on in order to see the motion on the biped in viewports. This option is turned on by default when you open the Mixer from the *Biped rollout (page 2–936)*.


The Motion Mixer might not display the entire active segment when it is first opened. To set the Mixer display to the extent of clips in the Mixer, click Set Range on the Motion Mixer toolbar.

Preparing BIP and XAF Files

To use a motion file with the Motion Mixer, it must first be saved as either a BIP or XAF file. Many BIP files come with 3ds Max; you can use these, or you can make your own. To find out how to create your own BIP files, see *Loading and Saving BIP Animation (page 2–920)*. For more information of XAF files, see *Saving and Loading Animation (page 3–472)*.

Importing XAF Files

Non-biped object assets are loaded into the Mixer using the *Load XML Animation File dialog (page 3–474)*, from which you pick the desired XAF motion file to import. If the motion file needs an XMM mapping file, you can either choose an existing one or open the *Map Animation dialog (page 3–478)* to create a new file.

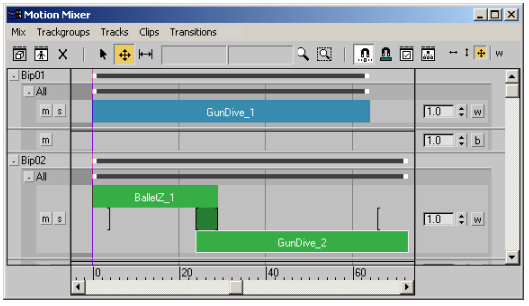
Tip: If you load multiple motion clips at the same time, you can pick a mapping file only once if all clips use the same.

Note: The Relative/Absolute and Replace/Insert options are unavailable when you import XAF files.

Note: Loading XAF clips from the *Reservoir (page 2–649)* follows the same workflow.

Interpreting Clip Names

When you load a clip into the Mixer, the clip appears on the track as a solid bar, with its name displayed on the clip.




Other information can appear on the clip, such as the start and end frames, and the *clip scale (page 2–615)*. You can change the information that displays on the clip by clicking Preferences on the Motion Mixer toolbar, and changing the Clips options on the *Mixer Preferences dialog (page 2–651)*.

You can use the same clip numerous times in the Mixer. Each version of the clip displays the name

of the BIP file, followed by a number, such as 1 or 2. The name might also have the letter Z between the name and number. These suffixes tell you how the clip is interpreted by the Mixer.

When you load a biped clip from a file, you have the option of setting the lowest starting foot height to Z=0 (the construction plane). If you leave this option turned on, the clip's motion will be moved in space so the biped's feet fall on the construction plane on the first frame of the clip. The letter Z appears after the clip name to indicate this setting.

When you load a clip into the Mixer for the first time, the clip name is followed by the number 1. If you clone or load the same clip to another part of the Mixer, the clip will display the same number or a different number depending on whether the new clip is an *instance (page 3–958)* or an *adaptation (page 3–905)* of the original clip. Instances are versions of the same clip used with the same object, or different objects of the same size. Adaptations are versions of the clip used with objects of different sizes.

The Mixer displays instanced clips with the same numbers, while adaptations receive sequential numbers. For example, if you clone a clip with the number 1 to a track for the same object, the new clip's name is also followed by the number 1. If you clone the clip to a object with a different size, the new clip is an adaptation, and the displayed name is followed an incremental number: in this case, the number 2. If you then clone the adaptation to another track for the same object, the new clip is an instance of that clip, and will be displayed with the number 2. These numbers can help you determine which clips are instances of one another, and which are adaptations.

The ability to tell instances from adaptations becomes important when you replace a clip in the Mixer with another clip. During this operation, you have the option of replacing only the selected clip, only the instances of the clip, or all instances and adaptations of the clip.

Note: Operations in the Mixer, such as cloning clips and editing clip timing, do not affect the clip sources (BIP files).

Procedures

To import BIP clips from a file:

1. Select a biped, and open the Motion Mixer. See *Adding Tracks to the Mixer (page 2–607)*.
2. In the Motion Mixer, highlight the track into which you want to bring the motion file. Be sure to highlight a blank area of the track and not any existing clips on the track.
3. From the Motion Mixer menu, choose Tracks > New Clips > From Files. You can also right-click the track and choose New Clips > From Files from the pop-up menu.
4. Navigate to the folder where your BIP files are stored, and select one or more BIP files.

   You can select as many BIP files as you like. These motions will be placed consecutively on the selected track.

   Files imported into the Mixer in this way are automatically placed in the Reservoir.

To import clips from the Reservoir:

In order to import clips from the Reservoir, you must first place them in the Reservoir. See *Using the Reservoir (page 2–626)*.

1. Highlight a blank area of the track into which you want to bring the motion(s).
2. From the Motion Mixer menu, choose Tracks > New Clips > From Reservoir. You can also right-click the track and choose New Clips > From Reservoir from the pop-up menu.
3. Select one or more BIP or XAF files.

To import clips from a motion flow script:

Motion flow scripts can be imported to a track in the Mixer. Both clips and transitions from the script are imported to the track.

Before you can import motion flow to a track, the biped must have at least one motion flow script assigned to it in Motion Flow mode. See *Creating a Motion Flow Script (page 2–1030)*.

1. Highlight a blank area of the track into which you want to bring the Motion Flow clips and transitions. You can choose either a transition or layer track.
2. From the Motion Mixer menu, choose Tracks > Import Motion Flow. You can also right-click the track and choose Motion Flow from the pop-up menu.
3. Choose a script from which to import motion.

   The clips and transitions from the script are imported to the Mixer and displayed the track. If the track was a layer track, it is converted to a transition track to accommodate the transitions imported from the script.

To see the motions in the Motion Mixer on the biped in the scene:

1. Select the biped for which you want to see the current motion from the Motion Mixer.
2.  Click Mixer Mode on the Biped rollout.
3. Play the animation or drag the time slider to see the motion in viewports.

Working with Clips in the Mixer

Once you have created tracks and imported clips to the Motion Mixer, you can adjust your animation mix by moving, cloning, and replacing clips. Within the Motion Mixer, you can:

- Move clips in time within the same track.
- Move clips to other tracks.
- Clone clips.
- Slide all clips in time on a track.
- Replace a clip with motion from a biped, a motion flow script, or another clip.

To move a clip in time within a track:

1.  Click Move Clips on the Motion Mixer toolbar.
2. Drag the clip and move it to another location on the track.

   A clip cannot be moved in such a way that it overlaps another clip. If you move a clip on top of another clip, it will return to its original location when you release the mouse.

Procedures

To move all clips horizontally on one track:

1. Select one or more clips. You can select multiple clips on different tracks for this operation.
2.  Click Slide Clips on the Motion Mixer toolbar.
3. Move the selected clips to slide clips to the left or right.

   When you slide a selected clip to the right, all unselected clips to the right of the clip move to the right. When you slide a clip to the left, all unselected clips to the left move to the left.

To move a clip to another track:

1.  Click Slide Clips on the Motion Mixer toolbar.
2. Move one clip on a track to slide all clips on the track.

To replace one clip with another:

After a clip is placed on a track, you can replace the clip with another clip. The new clip is scaled to fit the time of the original clip.

1. Select the clip(s) you want to replace.
2. Right-click a selected clip and choose Load Source > From Files from the pop-up menu. Select a BIP file to replace selected clips.
3. On the Mixer Clip Source Options dialog, choose whether to replace only the selected clips, or to also replace *instances (page 3–958)* (versions of the clip on same-sized bipeds) or *adaptations (page 3–905)* (versions of the clip on all bipeds).

To clone a clip:

Clips can be cloned within the same track, across tracks, and even across bipeds in the Mixer. A cloned clip retains the Z=0 status of the original clip (if this option was selected for the original clip), and any changes made to the original clip's timing. Cloning a clip to a different-sized biped creates an *adaptation (page 3–905)* of the original clip.

1. Make sure there's enough space on the track where you want to place the cloned clips. Move clips as necessary to make room.
2. Select the clip(s) you want to clone.
3. Hold down the Shift key on your keyboard, and drag the selected clips to the new location.

   If there isn't enough room where you attempted to placed the cloned clips, the operations is cancelled and clips are not cloned.

To replace a clip with motion from a biped in the scene:

All motions used in the Motion Mixer must be saved as BIP files. When you copy animation from a biped in the scene, you must save this motion as a BIP file.

1. Place a clip in the track using any method.
2. Select the clips you want to replace with biped motion from the scene.
3. From the Motion Mixer menu, choose Clips > Load Source > From Biped. The *Copy Biped Animation to Clip dialog (page 2–629)* appears.
4. Highlight the biped from which you want to copy the animation.
5. Enter a filename for the new clip.
6. Click Copy to replace clips with the biped motion.

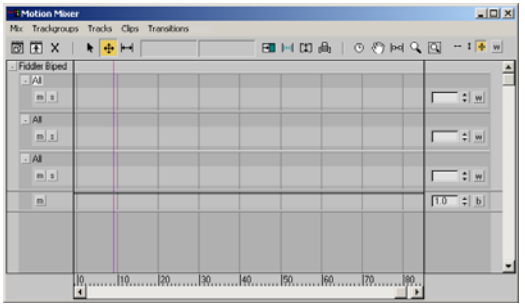
   The name of the clip in the track is replaced with the new filename.

Filtering Mixer Tracks

Each track in the Motion Mixer is part of a *trackgroup (page 3–1024)*. Each trackgroup can be filtered so it uses motions from one set of body parts and not another. For example, you might want to use the arm motions from one motion file with the legs from another. For this purpose, you would need two separate trackgroups.

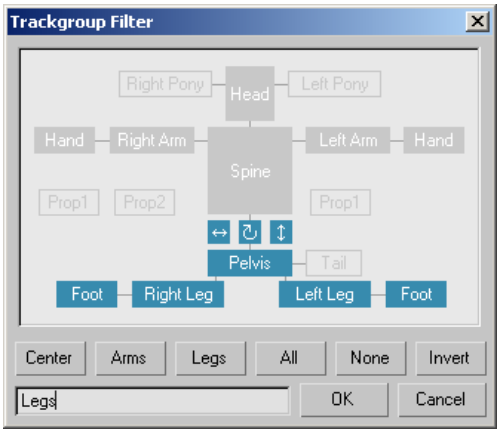
There are two steps involved in filtering with trackgroups:

- Create additional trackgroups as needed in the Mixer.



- *Filter* each trackgroup (set it to use some body parts but not all), and name the trackgroup

appropriately. This is accomplished with the Trackgroup Filter dialogs (biped and non-biped objects).

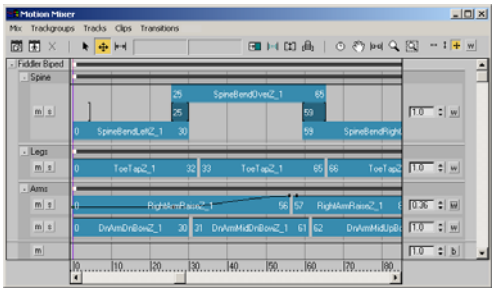


The Trackgroup Filter dialog for biped objects.



The Trackgroup Filter dialog for non-biped objects.

When a trackgroup is created, it is given one Layer track by default. However, you can add more tracks to each trackgroup, making it possible to mix motions for each set of body parts. For example, you could create a trackgroup for the arms, and have two tracks within this trackgroup. Then you could put two completely different sets of motions on each track, and use weighting to use one track or the other at different times during the motion. See *Adjusting Track Weight (page 2–619)*.



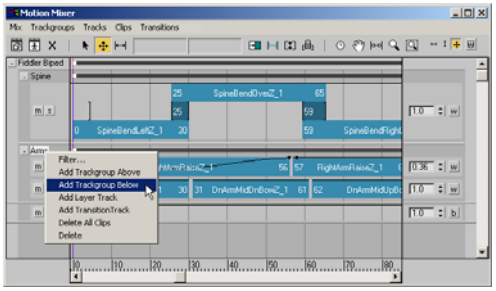
Trackgroups for Spine, Legs and Arms. The Arms trackgroup has two tracks.

Note: Create only one trackgroup for each filtered set of body parts. For example, don't use two trackgroups that both filter the spine. If you do so, the trackgroup highest in the Motion Mixer is used, and the other trackgroup is ignored. Instead, use multiple tracks within the trackgroup to achieve the animation you want.

Procedures

To add a new trackgroup for the biped:

1. Highlight the name of an existing trackgroup.
2. Right-click the trackgroup name and choose Add Trackgroup Above or Add Trackgroup Below.

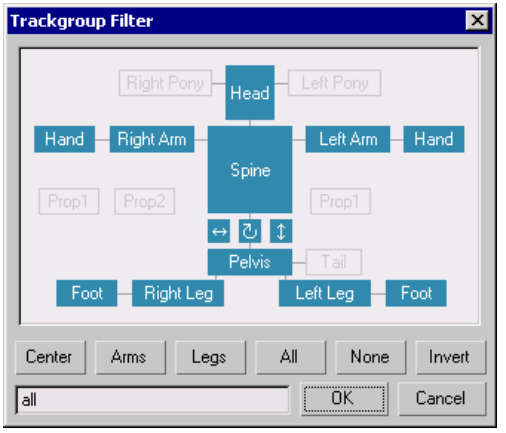


A new trackgroup and Layer track are added above or below the existing trackgroups. You can add more tracks to the trackgroup by right-clicking the existing track and choosing one of the options from the pop-up menu.

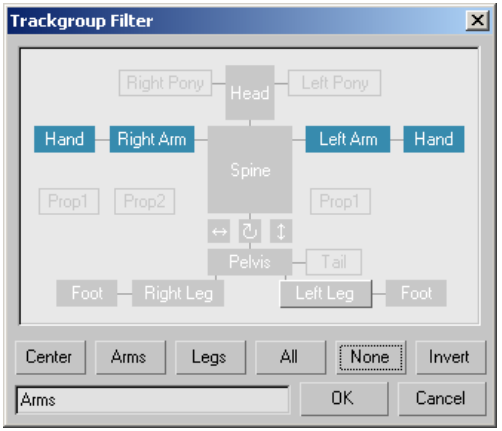
To filter and name a trackgroup:

1. Highlight the trackgroup label at the upper left corner of the trackgroup you want to filter. By default, the trackgroup label is All, denoting that the trackgroup currently affects all biped parts.
2. From the Motion Mixer menu, choose Trackgroups > Filter.

   The Trackgroup Filter dialog appears. By default, all body parts are selected.

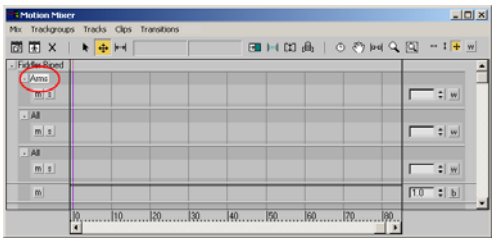


3. Click on body parts in the diagram to select or deselect them. You can also use the buttons below the diagram to help make the selection. For example, if the trackgroup will affect only the arms, click None to deselect all biped parts, then click the four arm boxes to select the arms.
4. Enter an appropriate name for the trackgroup at the bottom of the Trackgroup Filter dialog, such as **Arms**.



5. Click OK to close the dialog.

   The new trackgroup name appears at the upper left of the track display.



Adjusting Clip Timing

In the Motion Mixer, you can shorten clips or change their timing interactively within the track display. You can change the timing of motion clips in a number of ways:

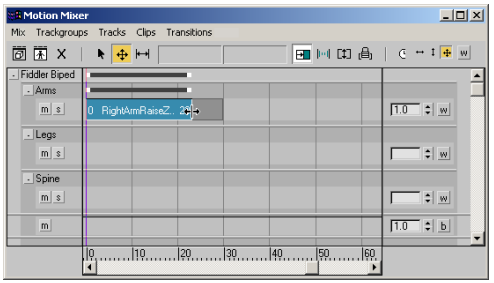
- Change the length of a clip without changing its speed by trimming the clip at the start or end.
- Change the speed of an entire clip by stretching it out or shrinking it.
- Cause varying changes in speed throughout the clip with time warps. See *Adding Time Warps (page 2–620)*.

Tip: When changing clip timing, it can be helpful to display start and end frame numbers and the clip time scale. To do this, click Preferences on the Motion Mixer and choose the appropriate options in the *Mixer Preferences dialog (page 2–651)*. As you choose options, the Motion Mixer display is changed interactively.

Procedures

To change the length of a clip without changing its speed:

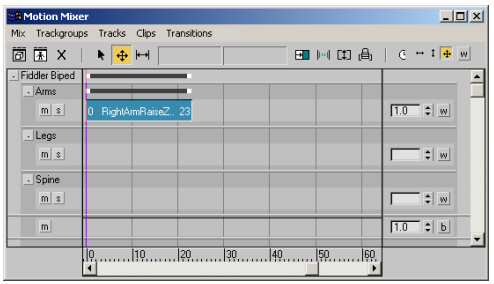
1. On the Motion Mixer toolbar, click Trim Clips.
2. Drag either end of a clip to trim its length.



   A gray bar appears on a trimmed end of a motion clip.

3. Turn off Trim Clips.

   The gray area disappears.



Tip: At any time, you can view original clip lengths in gray by turning on Trim Clips.

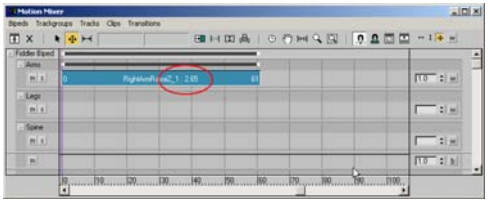
To remove scaling or trimming from a clip in the Motion Mixer, select the clip, right-click and choose Remove Scale/Trim.

To change the speed of an entire clip:

You can cause a motion to slow down or speed up by changing the length of the clip in the Mixer. Lengthening a clip will slow down the motion, while shortening the clip will speed it up.

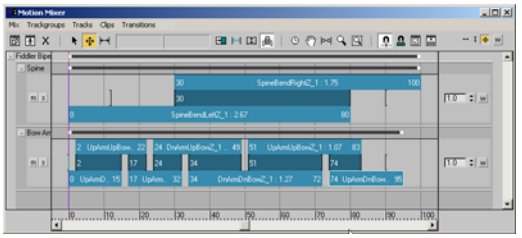
1.  On the Motion Mixer, click Move Clips.
2. Drag either end of the clip to extend or shorten it in time.

   Tip:  To see the ratio of the original motion time to the new time, click Preferences. In the Mixer Preferences dialog, turn on Scales. This displays the ratio on the clip next to its name.



Working with Transitions

When two or more clips are placed on a *transition track (page 3–1027)*, you can create transitions between any two clips. A transition creates a constant, gradual change from one clip to another.



Clips and transitions on Transition track. Transitions are shown as dark areas between clips.

To learn how to create a transition track, see *Adding Tracks to the Mixer (page 2–607)*.

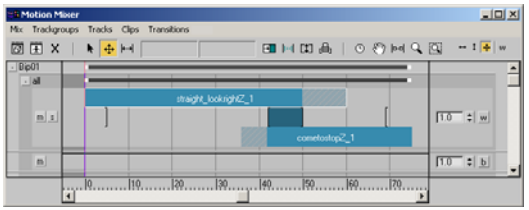
If you want to make a less linear blend between clips than the blend created by a transition, see *Adjusting Track Weight (page 2–619)*.

Adjusting Transitions

Transitions are automatically created between clips when you load the clips into a Transition track. You can adjust the start or end point of a transition by dragging either end of the transition in the Mixer.

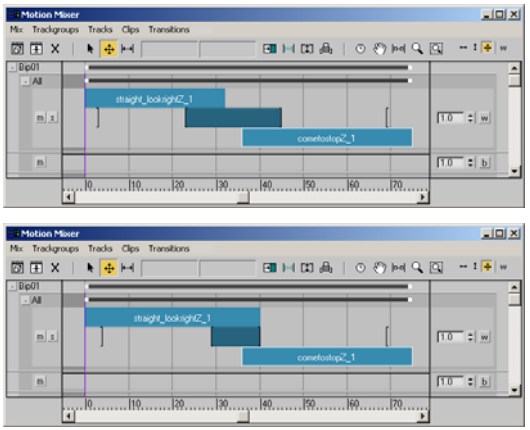
When a transition appears with hatchmarks, this indicates that the transition is invalid. This can occur when a clip's brackets are reversed, with the out-bracket before the in-bracket. You can correct this problem by dragging each end of the transition to reverse the ends.

A useful transition can occur anyplace where two clips overlap. The unused part of the clip on either side of the transition appears with hatchmarks to indicate that that part of the clip will not be used in the mix. This type of transition is often used for foot-based clips, when the best place for a transition is not always at the edge of a clip.



Hatched parts of the clips are not used in the mix.

The Motion Mixer will allow you to extend a transition into an area where two clips do not overlap. However, this transition will most likely result in undesirable motion in the mix.



Transitions extending into areas where clips don't overlap can create unpredictable results.

For the best results, make sure all your transitions cover only areas where two clips overlap completely.

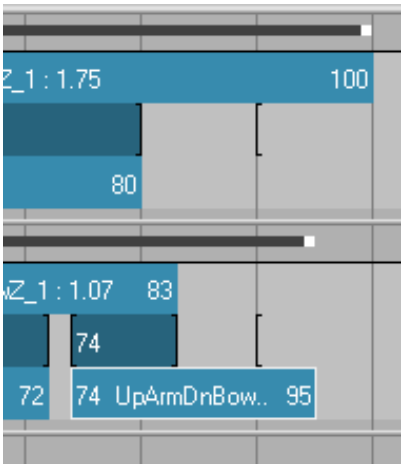
To aid in making your transitions as smooth as possible, you can select transitions and use the Transitions menu > Optimize option. This feature searches the two clips to find the start and end time that will result in the smoothest transition, and adjusts the transition accordingly.


If you want to snap the ends of a transition to the ends of clips, turn on Snap Clips on the Motion Mixer toolbar before dragging the transition edges.


When you move or scale a clip on a transition track, by default its transitions are also moved or scaled. To prevent transitions from changing when you move or scale a clip, turn on Lock Transitions on the Motion Mixer toolbar.

Using Transition Brackets

On a transition track, transition brackets appear around each transition, and at the start or end of each clip that is not surrounded by a transition.



Transition brackets

When you load a clip onto a transition track where clips already exist, the clip creates a transition between the last clip on the track and the new clip. The transition is created at the location of the last clip's transition bracket.

You can move a transition bracket to create the transition automatically exactly where you want it before loading the next clip. This feature can make it very quick and easy to load a long series of clips into a transition track with correct transitions.

Working with Foot-Based Transitions

Much like Motion Flow, transitions in the Mixer between foot-based motion clips (clips where IK constraints keep the feet planted at various times) must be carefully placed and adjusted to make the animation look as natural as possible. The workflow depends on whether you are animating a biped, or 3ds Maxobjects such as a rigged character.

Foot-Based Transitions with Biped

For the most natural transitions between foot-based motions, follow this workflow:

1. Find a good place to start and end the transition.

2. Create the transition with this timing. You can also optimize the transition to fine-tune it.
3. Change the transition focus.
4. Perform a mixdown.

Study both clips to find a start time and range for the transition that will result in smooth motion. Look for times when each foot has the same planting or free status in both clips. For example, a time period when both clips have the left foot planted and the right foot free is a good time range for a transition.

You can fine-tune your choice of transition time by optimizing the transition. This features searches each clip, either in its entirety or near an existing transition, to find the best timing for the transition.

You can improve a transition between foot-based clips by changing the transition focus on the *Transition Editor dialog (Biped Object) (page 2–636)*. In the example described above, the left foot would be chosen as the transition focus, because this would keep the left foot planted while the animation moves around it.

When transitions cause a planted foot to slide or pop slightly, or force a leg to straighten unnaturally, you can use a *mixdown (page 3–973)* to correct these problems. See *Exporting Animation to the Biped (page 2–624)*.

Foot-Based Transitions with 3ds Max Objects

For foot-based motions, follow this overall workflow:

1. Find a good place to start and end the transition.
2. Create the transition with this timing.
3. Choose the parent object to be a velocity blend object.
4. Choose the model's feet to be sub-blend objects.
5. Optimize the transition.

Procedures

To create a transition between two clips:

1. Right-click an existing track to insert a Transition track or convert the current track to a Transition track.
2. Right-click the track and choose New Clips > From Files from the pop-up menu. Choose at least two clips to load.

   The clips are placed on the transition track with transitions between them.
3. Drag the transition to move it, or drag its ends to resize it.
4. Before adding the next clip, move the last clip's transition bracket to the next desired transition location, and then load the clip.

To optimize a transition:

With foot-based clips, you can use the *Optimize command* to help you find the best timing for a transition. The Optimize command exhaustively searches the specified time period for times when the foot positions match most closely.

Tip: A good workflow is to look at the clips individually to find the best approximate time for a transition, then use Optimize to fine-tune the timing.

1. Create a transition between clips at approximately the right location on the track.
2. Right-click the transition and choose Optimize from the pop-up menu.

   The *Transition Optimization dialog (page 2–641)* appears.
3. Choose Search Near Existing Transition and click OK.

   The best transition is calculated, and the transition changes position and/or length in the Mixer to reflect the change.

Tip: If you aren't sure where you want the transition to take place, you can choose Search Entire Clip. This option can place the transition anywhere in the two clips. Check the motion afterward to ensure the transition takes place at an appropriate time.

To change the focus on a foot-based transition with Biped animation:

1. Scrub the time slider to the middle of the transition.
2. Zoom in on the feet in any viewport.
3. Right-click the transition and choose Edit from the pop-up menu. Alternately, you can select the transition and choose Transitions > Edit from the Motion Mixer menu.

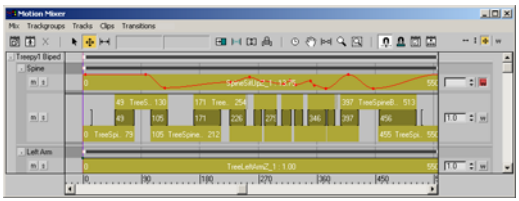
   The transition editor appears. In addition, two stick figures appear in viewports to show the bipeds' positions in the two clips. The yellow stick figure shows the biped's position in the first clip, while the red figure shows the second.
4. On the dialog that appears, choose a different Transition Focus from the list.

   As you change the Transition Focus, you will see the foot positions change. The Transition Focus sets a center or focal point for the transition, changing the biped's position slightly depending on the setting. On foot-based transitions, try different settings to bring the yellow and red skeletons closer together on a planted foot.

   Note: Transitions involving planted feet often result in one or both feet sliding when they should stay planted. After setting up the transition, you can correct this problem with a mixdown. See *Exporting Animation to the Biped (page 2–624)*.

Adjusting Track Weight

Each track has one or more *weight curves (page 3–1033)* that you can use to mix motion from several tracks to varying degrees. The weight curve determines how much a specific track contributes to the motion within the trackgroup.



With weight curves, you can create several full or partial transitions between all the tracks in a trackgroup. Compare with transitions on a *transition track (page 3–1027)*, which can only create a full transition between two clips.

Weight curves are useful for:

- Creating transitions between two tracks in a trackgroup.
- Creating random motion.
- Experimenting with transitions quickly and easily.

Each track's weight curve appears as a thin black line across the top of the track.

Tip: If you can't see the weight curve, click Preferences to display the *Mixer Preferences dialog (page 2–651)*. Turn Weight Curves off and on and watch the display to see the weight curve appear.

 You access the weight curve by clicking the Weight Mode button at the right end of the track. Clicking Weight Mode on the Motion Mixer toolbar enables weight curve editing for all tracks. When Weight mode is on, you can work only with the weight curve, and not with clips.

You edit a weight curve using nodes (points) on the curve. Clips on layer tracks automatically have

a node at each end of each clip. A transition track has no nodes by default.

Weight Curve Evaluation

By default, each track's weight curve is set to a value of 1.0 across the entire track, which uses 100% of the track's motion in the mix. When you change the curve to dip to a value below 1.0, the motion on the track is mixed with the motion on the track below it. For example, if you set the weight curve for two tracks in a trackgroup to 0.5, the two clips will be mixed equally.

Several tracks can be stacked in a trackgroup with varying weight curves on each one. For each frame in the animation, the Motion Mixer evaluates the curve on the topmost track. If its weight curve is less than 1.0 at that frame, the Mixer evaluates the next track down, and so forth. The Motion Mixer adds up the weight curve values until the sum reaches 1.0, and disregards tracks below that point.

For Transition tracks, one weight curve is used for the entire track. On Layer tracks, weight curves are set individually for each clip.

Procedure

To adjust a weight curve:

1. Open the Mixer, add a biped to the Mixer, and create two tracks for a trackgroup. See *Filtering Mixer Tracks (page 2–612)*.
2. Place different motion clips on the tracks.
3.  For the topmost track in the trackgroup, click Weight Mode to the right of the track display.

   The weight curve turns red to indicate it can be edited.

   You have to do the next two steps only if the weight curve has no nodes, or would like to create more nodes.
4. Move the cursor over the line until a small arrow appears. Click to set a node.
5. Move the cursor and click to set more nodes on the weight curve.

   Tip: On a transition track, you would ordinarily set a node at each end of the track, and at intervals in between. The number of nodes you set depends on the type of curve you want to make. Make as few nodes as possible to get the curve required for your animation.
6. Move the cursor over an existing node until a cross cursor appears.
7. Click and drag to move the selected node downward.

   The Motion Mixer mixes the motions from the topmost track with the track below it on that frame. The motions are mixed to a greater or lesser degree depending on how low you set the weight curve value for the topmost track.

Adding Time Warps

Select a biped. > Motion panel > Biped Apps rollout > Mixer > Motion Mixer menu bar > Clips > Add Time Warp

Graph Editors menu > Motion Mixer... > Clips > Add Time Warp

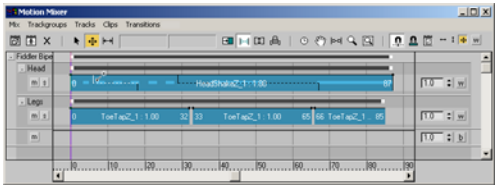
You can cause varying changes in speed throughout a clip with a time warp. Adding a time warp to a clip allows you visually to squash and stretch time over different parts of the clip. You can use a time warp to:

- Cause the object to do some parts of the motion quickly and others slowly.
- Cause a particular motion within the clip to occur at a specific time.

You use a time warp by picking a time within the motion clip and dragging that time's motion to another time in the same clip. The overall time of the clip is unchanged, so the object takes the same

amount of time to do the entire motion, but does some parts quickly and others slowly.

Dashed lines appear on a time-warped clip in the Mixer to indicate the motion speed at different points in the clip.



Time warp added to topmost clip.



Warping a clip's time requires two steps. First a time warp must be added to the clip with the Clips > Add Time Warp menu option. Then you turn on Editable Time Warps to see the dashed lines and adjust the clip's timing.

Procedures

To prepare to warp a clip's time:

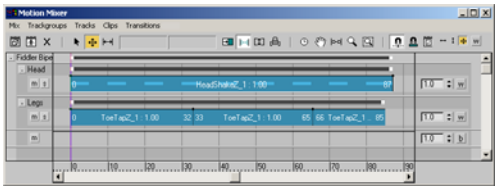
1. Open the Motion Mixer, add a biped to the Mixer, and load a clip on a track. See *Adding Tracks to the Mixer (page 2–607)*.
2. Scrub the time slider to find a motion you would like to occur earlier or later in the clip. Note the frame number. This frame number is the *original time*.
3. Scrub the time slider to find the frame on which you would like the motion to occur, and note the frame number. This frame number is the *warped time*.

   With a time warp, you will cause the motion to occur at the desired frame by moving the original time to the warped time within the clip.

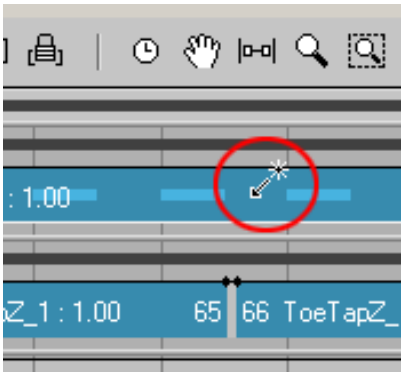
To warp a clip's time:

1. In the Motion Mixer, zoom in on the original time for the time warp. This is the frame number you noted in step 2 of the previous procedure.
2. In the Motion Mixer, select the clip.
3. From the Motion Mixer menu, choose Clips > Add Time Warp.

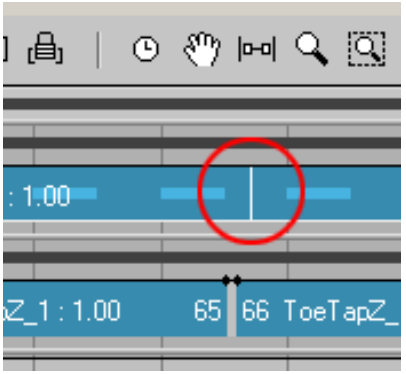
   The appearance of the clip does not change when a time warp is added.
4.  On the Motion Mixer toolbar, click Editable Time Warps.



   A series of evenly spaced dashes appears on the clip. The lengths of the dashes indicate the relative speed of the clip at various points. Right now, all dashes are evenly spaced to indicate the clip is playing at its original speed from beginning to end.
5. Move the cursor over the clip's horizontal center line until an arrow cursor appears.



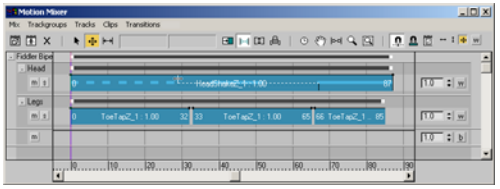
6. At the approximate original time, click the clip to set a time warp bar.



A time warp bar actually consists of two parts, a top and a bottom.

7. Click the top half the bar. The top half turns white to indicate it is selected.
8. Drag the top half of the bar to the left or right to set it to the warped time noted in step 3 of the previous procedure. This will cause the motion at the original time to occur at the warped time.

   Tip: As you drag the top bar, watch the Frame value at the top of the Motion Mixer window to see the top bar's new frame number.



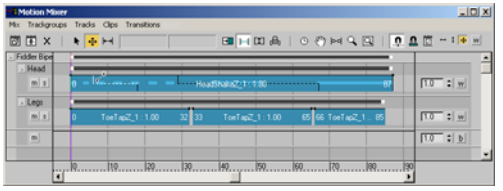
The lengths of the dashed lines change to indicate the new timing. Where lines are shorter, the motion will play faster. Where lines are longer, the motion will slow down.

9. Select and move the bottom bar to change the original time.

   Either bar can be moved to change the effect of the time warp. The length of separation between the two bars sets the degree of time-warping that will occur. When both bars are aligned, no warping occurs.

To add multiple time warps to a clip:

1. You can time-warp other parts of the clip by adding more time-warp bars with the arrow cursor, then moving the new top bars. When you add a second time warp bar after editing the first, the top bar appears where you click, but the bottom bar appears in a different location. The bottom bar's location corresponds to the original time of the top bar before the clip was time-warped.



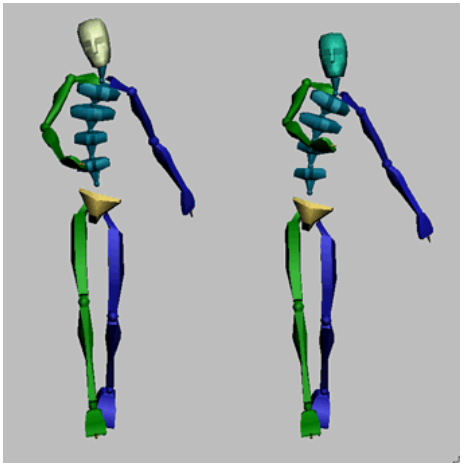
   Note: You cannot move one time warp's top bar past another on the clip.

2.  When you have finished warping the clip's time, turn off Editable Time Warps.

   When Editable Time Warps is turned off, the bars and the dashed line between them remain displayed on a clip to indicate it has been time-warped.

Adjusting Biped Balance in the Mixer

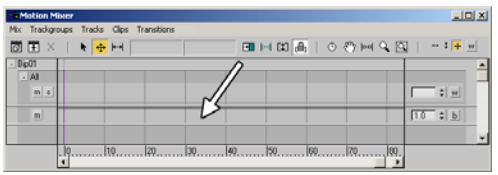
When you use different motions on the upper and lower parts of the biped, you can create a situation where the balance in the two clips do not match one another. For example, if the arms are waving wildly in the motion used for the upper body, the hip motion should compensate to some degree to keep the biped in balance. A straight mix of this arm motion with another hip/leg motion will most likely not match up in terms of balance.



Bipeds with same set of upper and lower body motions. Biped on the left has balance compensation, biped on the right does not.

By default, the Mixer compensates for differences in upper and lower body motion by making slight alterations to the spine and pelvis motions. If the biped bends over at the waist, for example, the pelvis will be moved to compensate for the weight shift, and the spine rotation will be lessened to help the biped keep its balance. Balance compensation is intended to make the biped's motion look as natural as possible.

In the Mixer, you can control the degree of automatic balance compensation using the balance track. One balance track is automatically created for each biped as it is added to the Mixer.



Balance track

Tip:  If the balance track for a biped is not displayed, click Preferences on the Motion Mixer toolbar, and turn on Balance Curves on the *Mixer Preferences dialog (page 2–651)*.

The balance track has a weight curve for adjusting the degree of automatic balance compensation between upper and lower body trackgroups. By default, the weight value of 1.0 across the balance track provides the maximum degree of compensation. You can reduce the weight curve at various points to lessen the degree of automatic balancing performed by the Mixer.

Note: Each biped can have only one balance track, and the balance track cannot be deleted.

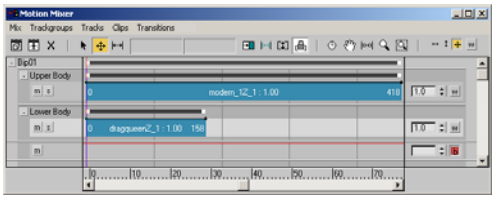
For fine adjustments to balance compensation, you can also change the parameters on the *Balance Parameters dialog*, available from the Mix menu.

Note: Balance Curves and related parameters are not available for non-biped mixes.

Procedures

To adjust balance using the balance track:

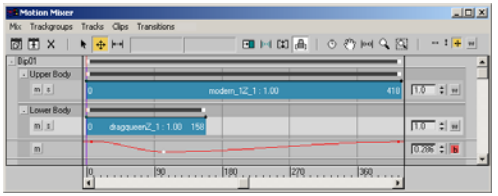
1. Add a biped to the Mixer, and create at least two trackgroups for the biped. See *Filtering Mixer Tracks (page 2–612)*.
2. Filter one trackgroup to use motion only from the spine, arms and head. See *Filtering Mixer Tracks (page 2–612)*. This trackgroup will hold the upper body motion.
3. Filter a different trackgroup to apply only to the legs, pelvis and COM tracks. This trackgroup will hold the lower body motion.
4. Add a clip with a great deal of upper body motion to the upper body track. See *Importing Clips to the Mixer (page 2–609)*.
5. Add a clip with very different lower body motion to the lower body track.
6. Click the Balance Weight Mode button at the far right of the balance track.



The weight curve becomes visible at the top of the balance track.

Tip: If you can't see the weight curve, click Preferences to display the Mixer Preferences dialog, and turn on Weight Curves.

7. Add nodes to the weight curve, and adjust the nodes. For information on how to add nodes and edit weight curves, see *Adjusting Track Weight (page 2–619)*.



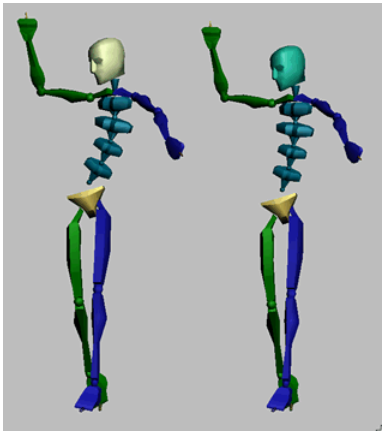
When the weight is set to 0.0, the Mixer will not adjust the spine and pelvis motion to compensate for differences in the upper and lower body motion. Values between 0.0 and 1.0 will adjust the balance to some degree.

To fine-tune balance compensation on the pelvis and spine:

The values on the Balance Parameters dialog can be used to make subtle adjustments to the biped's balance.

1. In the Motion Mixer, select the biped by clicking its name at the upper left corner of its trackgroups.
2. From the Motion Mixer menu, choose Mix > Balance Parameters.

   The Balance Parameters dialog appears.
3. To adjust the degree of horizontal balance compensation on the pelvis, change the Lateral Ratio parameter. Lower values make more forward/backward motion on the pelvis, while higher values use more side-to-side motion to compensate.
4. To adjust the degree to which spine rotation from the lower body motion is propagated on the upper body, change the Propagation parameter. Higher values rotate the spine links to better follow the COM and pelvis motion.



Bipeds with same upper and lower body motion. Biped on left has Propagation set to 0.0, biped on right with Propagation at 1.0.

Exporting Animation to the Biped

Once you've worked with the Motion Mixer to create an animated sequence for the biped, there are several ways to work with the finished mix.

If you want to work with the mix on the current biped (rather than a another biped in the scene), you have several options:

-  If you're satisfied with the animation as it is, you can turn on Mixer Mode on the *Biped rollout (page 2–936)* to display the mix on the biped, and simply render the scene.

- If transitions between different foot/leg motions have created small pops or jerks near the transition area, you will want to create a *mixdown (page 3–973)* for the biped and check the mixdown before copying it to the biped. A mixdown can correct many continuity problems with feet and legs. You can correct transitions and recompute the mixdown as many times as you like. When the mixdown is satisfactory, you can copy it to the biped.

If you want to use the motion with a different biped, you can copy the mixdown to the current biped, save the motion as a BIP file, and load the new BIP file onto the other biped.

The mix can also be saved as a MIX (*.mix*) file. A *MIX file (page 3–973)* can be loaded onto other bipeds and edited as needed in the Motion Mixer.

Fixing Foot-Based Transitions with a Mixdown

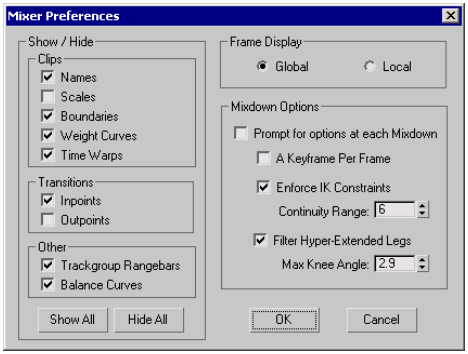
When you create a transition between two motions with planted feet (feet with IK constraints), the motion might not match perfectly between clips. When this occurs, a planted foot might slide or pop, or a leg might stretch unnaturally during or near the transition period.

A mixdown can correct two types of problems during transitions between foot-based motions:

- A foot sliding as it transitions from one planted motion to another.
- A leg extending completely straight (hyper-extending) unnaturally to reach a planted position.



The mixdown options on the *Mixer Preferences dialog (page 2–651)* can help correct these problems. To access this dialog, click Preferences on the Motion Mixer toolbar.



To have an effect on a mixdown, these options must be set before the mixdown is computed. If Prompt for options at each Mixdown is turned on, you will be prompted to set these options each time you compute a mixdown. Otherwise, the mixdown uses the settings on the Mixer Preferences dialog.

The options in the Mixdown Options group determine if and how the mixdown affects transitions between clips with planted feet. Enforce IK Constraints affects planted feet, while Filter Hyper-Extended Legs affects unnaturally straight legs. Turn on Prompt for options at each Mixdown if you want to be asked for these settings each time you do a mixdown. For more information on these options, see *Mixer Preferences Dialog (page 2–651)*.

Tip: You should not rely on the mixdown process to correct extreme problems with clips or transitions. For the best results, adjust transitions before a mixdown to minimize sliding feet and hyper-extending legs. See *Working with Transitions (page 2–616)*.

Procedures

To perform a mixdown and copy it to the biped:

1. Select the biped in the Motion Mixer by clicking the biped name at the upper left of the trackgroup set.

2. From the Motion Mixer menu, choose Bipeds > Compute Mixdown.

   If Prompt for options at each Mixdown is turned on in the Mixer Preferences dialog, you will be prompted for mixdown options. Click OK to start the mixdown.

   A progress bar at the bottom of the Mixer shows the progress of the mixdown process. After a few moments, a new track called Mixdown is created as the last track in the biped's mix.

   You can turn the track on and off by clicking the track. When the track is solid, the Mixdown track is on (active), and the biped performs the animation in the mixdown. When the track appears with hatchmarks, the track is off (inactive), and the biped performs the animation in the raw mix.

3. Click the Mixdown track to make it active.
4. Scrub the time slider or play the animation to check the transitions. If they need correction, turn off the Mixdown track, correct the transitions, and compute the Mixdown again. Repeat until the transitions are correct.
5. With the Mixdown track active, choose Bipeds menu > Copy Mixdown to Biped.
6.  Turn off Mixer Mode if it is on.

   The mix from the Motion Mixer is visible on the biped in viewports even though the biped is no longer in Mixer mode.

7.  To save the animation as a BIP file, click Save File on the Biped rollout.

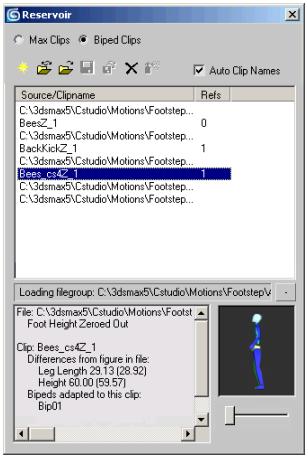
To save Motion Mixer data to a MIX file:

A MIX file saves the current state of the Motion Mixer data for a selected biped, including all clips, transitions, trims, scaling and time warps. It is not necessary to collapse clips before saving a MIX file.

1. Select the biped in the Motion Mixer by clicking the biped name at the upper left of the track set.
2. From the Motion Mixer menu, choose Bipeds menu > Save Mix File. Enter a filename for the *.mix* file, and click Save.

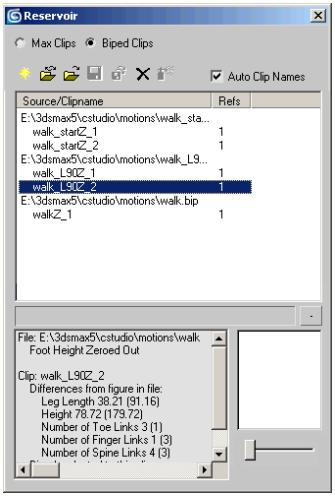
Using the Reservoir

The Reservoir serves as a storage facility for *motion clips (page 3–975)* (BIP and XAF files) that you use with the Motion Mixer. You can load clips directly into the Reservoir, and all clips that you load directly in the Motion Mixer also show up in the Reservoir.



For each motion file used in the Motion Mixer, the Reservoir list shows the source (disk path and name of the file), and, under the source, shows the name of each clip or set of clips derived from that file.

The Motion Mixer can apply multiple *adaptations (page 3–905)* (occurrences of the clip for different-sized bipeds) to different bipeds. In the Reservoir, the various clip adaptations are listed under the clip name.



If you click a clip in the list, the window under the list displays information about the clip, including structural differences between the figure it's applied to and the figure from which the file was originally saved. A second window displays a graphical view of the biped animation; drag the slider below the window to move through the animation.

Note: The preview window is available only for BIP files.

Tip: For best results when loading animation clips that might require remapping into the Motion Mixer, load the clips directly into Mixer tracks, rather than going through the Reservoir. This always gives you the opportunity to remap the animation data.

Saving Clip Adaptations

The Reservoir lets you save a unique version of an adapted clip to a new BIP file. This new file will contain the new biped's size data, so will not have to be adapted when it is loaded into the Mixer.

If you are mixing long BIP files on a biped of a size other than the size used to create the BIP file, when you reload the MAX file or load a MIX file, you might have to wait while the Mixer recalculates the adaptation for the biped. You can save loading time by saving adapted clips to new BIP files and replacing the original clip in the Mixer.

If you save an instanced clip from the Reservoir under a new file name, all clips in the Reservoir are renamed likewise.

Procedures

To access the Reservoir:

-  On the Motion Mixer toolbar, click the Reservoir button.

  The Reservoir opens.

To add clips to the Reservoir:

1. Choose between Max Clips and Biped Clips.
2.  On the Reservoir toolbar, click the New Entry button.

   The Open dialog appears.
3. Find the directory containing the clip files to load, highlight the desired file or files, and then click the Open button.

   The file loads and appears as a Source in the Reservoir list. You can then add it to a track in the Motion Mixer using the latter's New Clips > From Reservoir command.

To replace a clip in the Reservoir:

Use the Load File command to replace a highlighted source or clip.

1. In the Reservoir list, click a single item.

   The software permits replacing only one source or clip at a time.
2.  On the toolbar, click Load File.

   The Open dialog appears.

3. Find the directory containing the clip files to load, highlight the desired file, and then click the Open button.

   The file loads and replaces the highlighted item in the Reservoir list and the Motion Mixer. If you replace a source item, all of its clips are replaced as well. If you replace a clip, all of its references are replaced.

To save an instanced clip to a new clip:

Once you apply a clip instance to a different biped in Motion Mixer, the software automatically fits the motion to the biped. You can then save the modified motion from the Reservoir.

1. In the Reservoir list, highlight the clip whose motion you want to save.
2.  On the Reservoir toolbar, click the Save File button. Use the Save As dialog to specify a name and location for the motion file, and then click the Save button.

   The motion file is saved as fitted to the biped to which the clip is applied.

   Note: The motion file is saved at its original length, whether or not you changed the length of any instances of the clip in the Motion Mixer.

To save multiple instanced clips from the Reservoir:

1.  On the Reservoir toolbar, click Batch Save.

   This operation does not depend on which items are highlighted in the Reservoir list.
2. In the Save Reservoir Items dialog, turn on items you want to save. For each source/clips group, you can save the source or any combination of clips, but not both.

   Optionally turn on Reset Names After Saving.
3. Click Save to save the specified items.

To remove clips not used in the Motion Mixer from the Reservoir:

If you delete all instances of a clip from the Motion Mixer, it appears in the Reservoir with a Refs value of 0.

-  On the Reservoir toolbar, click Clean to remove these clips from the Reservoir.

  The motion file is saved under the new name.

Motion Mixer User Interface

Select a biped. > Motion panel > Biped Apps rollout > Mixer

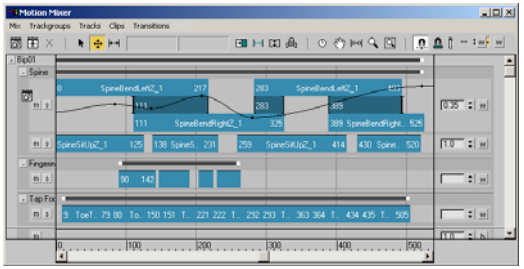
Graph Editors menu > Motion Mixer...

The Motion Mixer is used to mix motion clips that are assigned to bipeds and non-biped objects. The concept of the Motion Mixer is based on audio mixing done in the music and sound industries. Like an audio mixer, Motion Mixer lets you add a series of *tracks (page 3–1023)* for each object in a scene. To each track you add *clips (page 3–975)* in the form of BIP or XAF files. You can also add *transitions (page 3–1027)* to smoothly link clips together. Clips and transitions can be *filtered (page 3–939)* so their motion affects only specified biped parts.

The entirety of the clips, transitions and other clip settings for a single object in the Motion Mixer is called a *mix (page 3–972)*. As each object has its own mix, there can be several mixes in the Motion Mixer at any given time. You can save a mix as a *MIX file (page 3–973)*.

For more information and procedures describing how to use the Motion Mixer, see *Using the Motion Mixer (page 2–604)* and the topics that follow it.

Interface



The Motion Mixer interface contains the following:

- The *Motion Mixer menu bar (page 2–629)*, with five menus.
- The *Motion Mixer toolbar (page 2–642)*, containing button-activated commands for adding bipeds, transitions, and other settings for creating a mix.
- The *Motion Mixer editor (page 2–646)*, showing a graphical interface where the mix is built.

The Motion Mixer is normally opened as a floating window. It can also be opened in a viewport, like the Asset Browser and MAXScript Listener.

See also

Using the Motion Mixer (page 2–604)

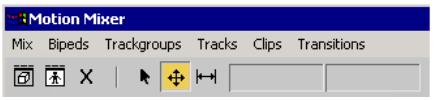
Motion Mixer Menus

Select a biped. > Motion panel > Biped Apps rollout > Mixer

Graph Editors menu > Motion Mixer...

The Motion Mixer interface contains five menus located at the top of Motion Mixer window.

Interface



Mix menu

The Mix menu provides tools for managing your bipeds and non-biped objects when they are selected. If an object is already added to the Motion Mixer, right-clicking its name displays the menu. Use commands on this menu to load and save MIX files, adjust track colors and balance parameters, and activate and deactivate the Mixer mode.

Edit Max Mix Object...—Opens the Edit Max Mix Object dialog, which lets you modify the content of the highlighted non-biped object mix. This dialog is comparable to the *Max Objects To Mix dialog (page 2–642)* with the exception of the Bake Removed Nodes option.

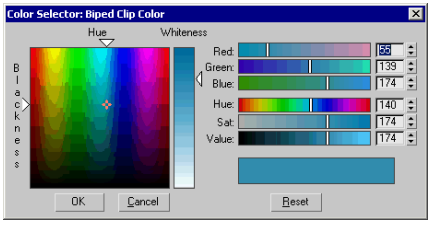
Note: This dialog is only available for non-biped object mixes.



Bake Removed Nodes—When on, animation keys from removed objects are baked to the main timeline. Default=off.

Add Trackgroup—Adds a *trackgroup (page 3–1024)* to the top of the selected object's mixes.

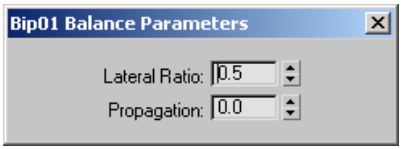
Track Color...—Allows you to change the color of the clips in the selected object's mixes. Choosing this command displays a Color Selector, where you can specify the clip color.



Balance Parameters...—Opens the Balance Parameters dialog. These values influence the effects of balance compensation on the biped's

pelvis and spine when you use different motions on the upper and lower parts of the biped. See *Adjusting Biped Balance in the Mixer (page 2–622)*.

Note: This dialog is available only for biped object mixes.



- **Lateral Ratio—**Decreasing the Lateral Ratio to 0.0 will cause balance compensation to use only forward/backward motion on the pelvis. Increasing Lateral Ratio to 1.0 will adjust the pelvis with only side-to-side motion. Range=0.0 to 1.0; Default=0.5.
- **Propagation—**Affects the degree to which spine links are rotated to follow COM and pelvis rotation. When set to 0.0 (the default), spine rotation is not influenced by lower body motion. When set to 1.0, all spine links except the topmost one use the lower body motion's spine animation to some degree to better follow the COM and pelvis motion. The lowest spine link is rotated the most, and each subsequent spine link is rotated less and less. The topmost spine link is not rotated at all so arm motion will be preserved.

  In order for this setting to affect spine rotation, the lower body motion must have some spine animation as part of the motion clip.

  A Propagation setting above 0.0 is appropriate when the upper body motion causes the spine to bend over a great deal during a large part of the animation. Range=0.0 to 1.0; Default=0.0.

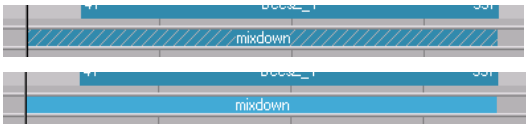
Delete—Clears the selected object(s) from the Mixer.

Load Mix File—Opens a dialog where you can select a *MIX file (page 3–973)* to load into the Mixer.

Save Mix File—Opens a dialog where you can save the current object's *mix (page 3–972)* to a MIX file.

The following Mixdown-related tools are available only for biped object mixes.

Compute Mixdown—Performs a *mixdown (page 3–973)* on the selected biped(s). A Mixdown track appears at the bottom of each biped's mix. Clicking the Mixdown track deactivates or activates it. When the Mixdown track is active, each biped inherits this single track's animation rather than the transition and layer tracks in the mix. When the Mixdown track is inactive, the biped uses the *raw mix (page 3–972)* instead.



An inactive Mixdown track (above) and an active Mixdown track (below)

A mixdown can make automatic adjustments to transitions between planted foot motions. Creating a Mixdown track makes it possible to copy the mixdown to the biped so it can be saved as a BIP file. See *Exporting Animation to the Biped (page 2–624)*.

Delete Mixdown—Removes the mixdown from the mix of the selected biped(s).

The next three menu options are mutually exclusive; only one can be turned on at a time.

Copy Mixdown to Biped—Copies animation from the Mixdown track to the biped. The animation remains on the biped even after Mixer mode is turned off, and the animation can be saved as a BIP file.

Effect Raw Mix—When this option is turned on, Mixer mode is turned on for the biped. Default=active.

Effect Mixdown—When this option is turned on, the mixdown track is activated. This option is available only if a mixdown has been computed.

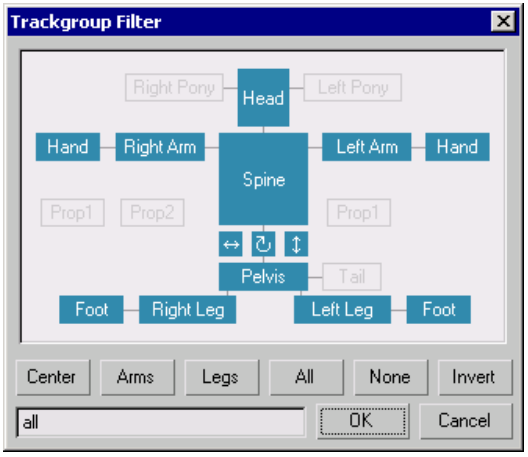
Effect Biped—When this option is turned on, Mixer mode is turned off for the biped. Click Effect Raw Mix to turn Mixer Mode on again.

Trackgroups menu

When a trackgroup is selected, you can access the commands in the Trackgroup menu. Alternately, you can right-click a trackgroup label to open the menu.

The Trackgroups menu lets you administer the bipeds in the mixer. Each biped you add to the Motion Mixer gets its own trackgroup, a selection of biped parts.

Filter...—Opens the Trackgroup Filter dialog where you can customize the object body parts that are included in a trackgroup mix. This command is only available when a single trackgroup is selected. For more information, see *Trackgroup Filter Dialog (Non-Biped Object) (page 2–645)*, *Trackgroup Filter Dialog (Biped Object) (page 2–645)*, or *Filtering Mixer Tracks (page 2–612)*.



The Trackgroup Filter dialog for biped object mixes.



The Trackgroup Filter dialog for non-biped object mixes.

Add Trackgroup Above—Adds new trackgroups above the currently selected ones.

Add Trackgroup Below—Adds new trackgroups below the currently selected ones.

Add Layer Track—Adds a Layer track at the top of the selected trackgroups. Layer tracks hold clips to be mixed without transitions.

Add Transition Track—Adds a Transition track at the top of the selected trackgroups. Transition tracks hold clips that are mixed with transitions.

Clips in Transition tracks will automatically spawn transitions when they overlap.

Delete All Clips—Clears the selected trackgroups of all clips and transitions.

Delete—Clears the selected trackgroups from the Motion Mixer.

Tracks menu

The commands in the Tracks menu are only active when you have a track, or tracks, selected in the Mixer. These menu commands are also available when you right-click a track in the Mixer window.

Add Layer Track Above—Adds a new Layer track above the currently selected tracks.

Add Layer Track Below—Adds a new Layer track below the currently selected tracks.

Add Transition Track Above—Adds a new Transition track above the currently selected tracks.

Add Transition Track Below—Adds a new Transition track below the currently selected tracks.

Convert to Layer Track—Converts the selected Transition tracks to Layer tracks.

Convert to Transition Track—Converts the selected Layer tracks to Transition tracks.

Optimize All Transitions—All transitions on the track are optimized.

New Clips > From Files—Opens a dialog where you can choose one or more BIP or XAF files to add to the selected track. If a selected BIP file has been saved for this release of 3ds Max , a thumbnail is displayed in the Motion Preview where you can move the slider and get an idea of what the motion looks like.

Any clip you add to a track using this method is automatically added to the *Reservoir (page 2–649)*.

Note: If you choose a BIP file that is older, you will receive a warning that the file is obsolete and should be resaved. Until the BIP file is resaved, it will not display in the Motion Preview.

New Clips > From Reservoir—Opens a Reservoir Files dialog where you can choose one or more BIP or XAF files listed in the Reservoir.

Import Motion Flow—Allows you to import clips and transitions from a motion flow script to the selected track. You can import a script only if it is active on the currently selected biped, and if it has been loaded or saved as part of an *MFE file (page 3–972)*. If scripts are available for import, the Motion Flow Scripts dialog displays a list of available scripts. The imported script replaces the currently selected track with a *transition track (page 3–1027)* containing all the clips and transitions in the script.

Note: This dialog is available only for biped object mixes.

Note: If the selected track contains clips and transitions when you choose this option, these clips are removed from the track when the motion flow script is imported.

Delete All Clips—Deletes all the clips residing on the selected track.

Delete—Deletes a selected track or tracks.

Clips menu

The Clips menu is active when a clip is selected or when you right-click a clip in the Mixer. Some menu commands are not available when multiple clips are selected.

Add Time Warp—Applies the ability to be time warped to the selected clips. A time warp allows you to alter clip speeds. It gives you the ability to stretch or squeeze time for parts of selected clips. For more details and procedures, see *Adding Time Warps (page 2–620)*.

Remove Time Warp—Removes the ability for clips to be Time Warped and removes any existing Time Warping from selected clips.

Tile View—Tiles the selected clip along the width of the Motion Mixer. The number of duplicated clips depends upon the clip's length and the number of frames displayed across the Motion Mixer. For example, if a clip is 100 frames long and the number of frames across the Motion Mixer is 500, the clip will tile five times.

Tile Range—Tiles the selected clip along the range of the active time segment. The number of duplicated clips depends upon the clip's length and the number of frames in the active time segment. For example, if a clip is 100 frames long and the number of frames in the active time segment is 500, then the clip will tile five times.

Remove Trim/Scale—Removes the trim/scale from the selected clip or clips.

Delete—Deletes a selected clip or clips from a track.

Load Source > From File—Opens a dialog where you can choose a new BIP or XAF file to replace selected clips. After selecting the new clip, the *Mixer Clip Source Options dialog (page 2–634)* is displayed.

Load Source > From Biped—Opens the *Copy Biped Animation To Clip dialog (page 2–635)*, where you can choose a biped from which to copy animation to selected clips.

Note: This tool is available only for biped object mixes.

Load Source > From Reservoir—Opens the *Reservoir File Groups dialog (page 2–635)*, where you can choose a clip from the Reservoir to replace selected clips.

Copy to Biped—Puts the clip's animation onto the base state of the selected biped, evident when not in Mixer or any other mode. This option is available only when a single clip is selected.

Note: This dialog is available only for biped object mixes.

Collapse—Saves a collapsed version of the selected clip to a new BIP file. Collapsing removes all scaling, trims, and time warps from the clip, and prompts for a new BIP file name for the collapsed version of the clip. The original selected clip is replaced by the new clip. Collapsing a clip does not affect the selected clip itself.

Note: This dialog is available only for biped object mixes.

Transitions menu

Select a transition track to make commands on this menu active. Right-clicking a transition also opens this menu.

Edit—Opens either the *Mixer Transition Editor dialog (Biped Object) (page 2–636)* or the *Mixer Transition Editor dialog (Non-Biped Object) (page 2–638)*, based on the kind of mix the transition is applied to. The dialogs are similar to the *Transition Editor (page 2–1051)* accessed from the Motion Flow rollout.

Optimize—Optimizes the selected transitions. This uses optimization algorithms similar to those in Motion Flow.

Convert To Loopable Clip—Allows you to create a loopable clip from a transition between two clips that are clones of one another. The two clips must share the same source file, meaning they must point to the same clip in the reservoir and they must have identical time warps and scales. Although the two clips don't have to be created via cloning, it is probably the best way to create them.

The loopable clip will replace the two original clips and their transition and will begin exactly one frame before the transition starts, so at that frame, the animation will appear the same as it did

before you created the loopable clip. If you tile the loopable clip after you create it, it should loop perfectly.

Note: This dialog is available only for biped object mixes.

Motion Mixer Dialogs

Mixer Clip Source Options Dialog

Select a biped. > Motion panel > Biped Apps rollout > Mixer > Select clip(s). > Motion Mixer menu bar > Clips > Load Source > From File

Select non-biped clip(s). > Motion Mixer menu bar > Clips > Load Source > From File

The Mixer Clip Source Options dialog opens when you select one or more clips on a track and choose Load Source > From File *Clips menu (page 2–632)*. The new clip you choose replaces the selected clips. You also have the option of replacing *instances (page 3–958)* and *adaptations (page 3–905)* of selected clips.

Interface



Load Animation Into:—

- **The selected clips only—**Replaces only the selected clip with the newly specified clip.
- **All instances of the selected clip—**Replaces instances of the clip (any occurrence of the selected clip within tracks for the same biped, or other bipeds of the same size).
- **All instances and adaptations of the selected clip—**Replaces adaptations (any occurrence of the selected clip on all bipeds' tracks).

Method to Fit Animation to Clip:—

- **Scale animation to fit clip—**Scales the length of the loaded clip to match the length of the clip it replaces in the track.
- **Trim animation to fit clip—**Trims the length of the loaded clip to match the length of the clip it replaces in the track.
- **Set clip length to animation length, shifting remaining clips in time—**For selected clips, changes the animation length to the loaded clip length. All clips on the track after the selected clip are shifted depending on the length of the newly loaded clip.

Remove Weights—Removes any existing weight curves from clips selected to be replaced.

Remove Warps—Removes any existing time warps from clips selected to be replaced.

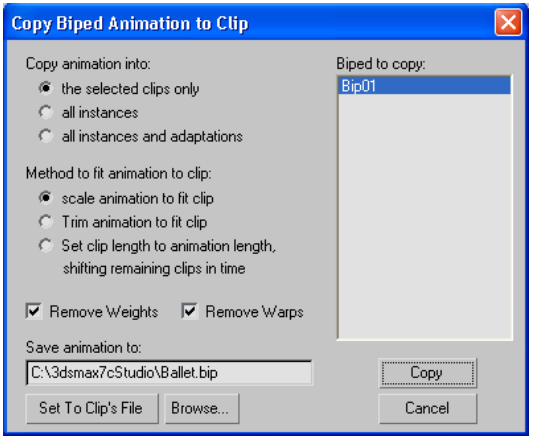
Copy Biped Animation to Clip Dialog

Select a biped. > Motion panel > Biped Apps rollout > Mixer > Motion Mixer menu bar > Clips > Load Source > From Biped

The Copy Biped Animation To Clip dialog opens when you select one or more clips on a track and choose Load Source > From Biped on the *Clips menu (page 2–632)* The animation from the biped replaces the selected clips.

Note: This dialog is available only for biped object mixes.

Interface



Copy Animation Into:—

- **The selected clips only**—Only the selected clip in the track is replaced by a new clip chosen from the reservoir.
- **All instances**—Replaces *instances (page 3–958)* of the clip (any occurrence of the selected clip within tracks for the same biped, or other bipeds of the same size) with the clip selected from the Reservoir.
- **All instances and adaptations**—Replaces any occurrence (*adaptation (page 3–905)*) of the selected clip on all bipeds' tracks in the Mixer, with the clip selected from the Reservoir.

Method to Fit Animation to Clip:—

- **Scale animation to fit clip**—Scales the length of the loaded clip to match the length of the clip it replaces in the track.
- **Trim animation to fit clip**—Trims the length of the loaded clip to match the length of the clip it replaces in the track.
- **Set clip length to animation length, shifting remaining clips in time**—For selected clips, changes the animation length to the loaded clip length. All clips on the track after the selected clip are shifted depending on the length of the newly loaded clip.

Remove Weights—Removes any existing weight curves from clips selected to be replaced.

Remove Warps—Removes any existing *time warps (page 2–620)* from clips selected to be replaced.

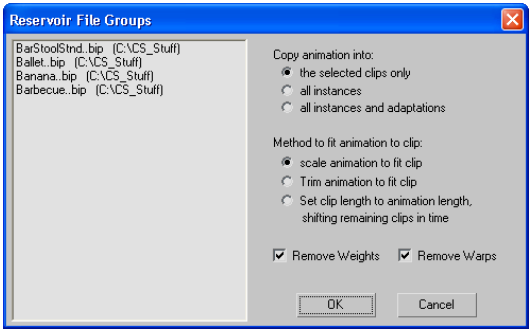
Reservoir File Groups Dialog

Select a biped. > Motion panel > Biped Apps rollout > Mixer > Motion Mixer menu bar > Clips > Load Source > From Reservoir...

Graph Editors menu > Motion Mixer... > Select a non-biped object track > Motion Mixer menu bar > Clips > Load Source > From Reservoir...

The Reservoir File Groups dialog opens when you select clips and choose Load Source > From Reservoir on the *Clips menu (page 2–632)*. This action replaces selected clips with the clip you choose from the Reservoir. This dialog shows a list of all the clips in the *Reservoir (page 2–649)*.

Interface



Copy Animation Into:—

- **The selected clips only**—Only the selected clip in the track is replaced by a new clip chosen from the reservoir.
- **All instances**—Replaces *instances (page 3–958)* of the clip (any occurrence of the selected clip within tracks for the same object, or other objects of the same size) with the clip selected from the Reservoir.
- **All instances and adaptations**—Replaces any occurrence (*adaptation (page 3–905)*) of the selected clip on all objects' tracks in the Mixer, with the clip selected from the Reservoir.

Method to Fit Animation to Clip:—

- **Scale animation to fit clip**—Scales the length of the loaded clip to match the length of the clip it replaces in the track.
- **Trim animation to fit clip**—Trims the length of the loaded clip to match the length of the clip it replaces in the track.
- **Set clip length to animation length, shifting remaining clips in time**—For selected clips, changes the animation length to the loaded clip length. All clips on the track after the selected clip are shifted depending on the length of the newly loaded clip.

Remove Weights—Removes any existing weight curves from clips selected to be replaced.

Remove Warps—Removes any existing time warps from clips selected to be replaced.

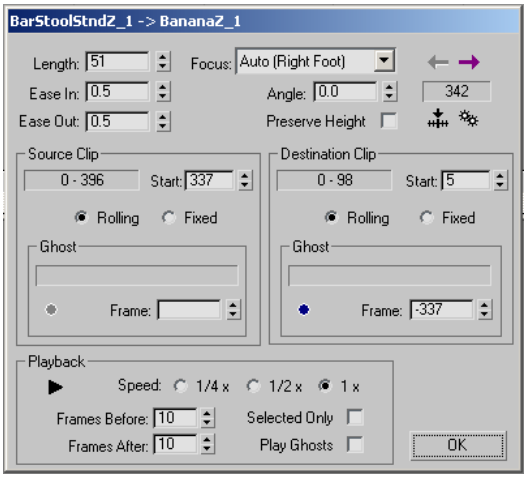
Mixer Transition Editor Dialog (Biped Object)

Select a biped. > Motion panel > Biped Apps rollout > Mixer > Motion Mixer menu bar > Transitions > Edit

The Mixer Transition Editor controls *transitions (page 3–1027)* on *transition tracks (page 3–1027)*. When it is active, you view transitionary stick figures to aid in placing your transitions. You can use this dialog to change the start and end times of transitions, change the transition focus, and perform other functions.

This dialog is very similar to the *Motion Flow Transition Editor (page 2–1051)*.

Interface



Length—Sets the number of frames for the duration of the transition.

Transitions are calculated by matching velocities in both clips. Smooth out abrupt velocity changes using longer transitions.

Ease In—Ease-in value for the source clip.

Ease Out—Ease-out value for the destination clip.

Transition Focus—Lets you specify a focus point on the biped where the transition takes place. The Mixer will attempt to match movement based on this selection. For example, if Left Foot is selected, the transition will use the left foot as a focal point during the transition, aligning the motion of the left foot in both clips as much as possible during the transition. Default=Auto.

- **Auto—**The transition focus is calculated by averaging the overall position of the biped as it transitions from one clip to the next.
- **Center Of Mass—**The transition focus is based on the center of mass position of the biped as it transitions from one clip to the next.
- **Left Foot—**The transition focus is based on the left foot position of the biped as it transitions from one clip to the next.
- **Right Foot—**The transition focus is based on the right foot position of the biped as it transitions from one clip to the next.
- **Both Feet—**The transition focus is based on an averaged foot position of both of the biped's feet as it transitions from one clip to the next.

Note: The best way to see the differences between the transition foci is by watching the yellow and red ghosts.

Preserve Height—When turned on, Preserve Height will allow gradual vertical motion to accumulate from the source to the destination clip. For example, climbing stairs will loop upward. If it is turned off, the next clip is always set so that its lowest point is at z=0. This default insures that motion does not gradually float up (or down) with each additional clip.

Angle—Sets the direction of the destination clip.

The angle of the destination clip is automatically set for best body fit between the two clips when the Start Frame values change. Use Angle to change the direction of the destination clip.



Previous Transition—Go to the previous transition in the transition track.

Displays the previous transition in the Transition Editor, moves the time slider to the start frame of the previous transition and highlights the previous clip in the transition track.

Next Transition—Go to the next transition in the transition track.

Displays the next transition in the Transition Editor, moves the time slider to the start frame of the next transition and highlights the next clip in the Scripts list.

Optimize Transition—Displays the *Transition Optimization dialog (page 2–641)*.

Options in the Transition Optimization dialog allow you to search for the location for the transition.

Go To Start Frame—Moves the time slider to the first frame of the transition.

The number field shows the start frame number.

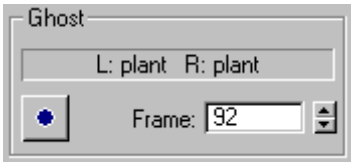


Start Frame—Set the transition start frame for the source and destination clips in their respective fields. Duration for the source and destination clips displays above the Start Frame fields.

Rolling—Keeps the clip in motion during the transition.

Fixed—For the source clip, this option freezes the biped at the Start Frame position during the

transition. For the destination clip, this option freezes the biped at the End Frame position. If Fixes is chosen for both the source and destination clips, the transition is a gradual interpolation from one frozen pose to another.



Set Start Frame—Copies the value in the Frame field of the Ghost area to the Start Frame field in the Clip area. The position of the destination clip changes to match the biped body in the destination clip to the biped body in the source clip.

Locate an appropriate start frame for the source and destination clips by using the Frame spinner and viewing the positions of both stick figures, then click Set Start Frame.

The destination clip is rotated and positioned to match both bipeds. Use the Angle spinner to reorient the destination clip.

Frame—Use the Frame spinner to scrub a stick figure back and forth, which allows you to determine a start frame for the source and destination clips. Visual feedback of the stick figures is a good way to judge which start frames are needed for the source and destination clips.

Mixer Transition Editor Dialog (Non-Biped Object)

Graph Editors menu > Motion Mixer... > Select a non-biped object clip transition. > Motion Mixer menu bar > Transitions > Edit

The Mixer Transition Editor controls *transitions (page 3–1027)* on *transition tracks (page 3–1027)*. You can use this dialog to change the start and end times of transitions, change the transition focus, and perform other functions. Mix transition blend objects are changeable from clip to clip, and any new transition that results from cloning or adding clips initially uses the blend options from the last clip in time.

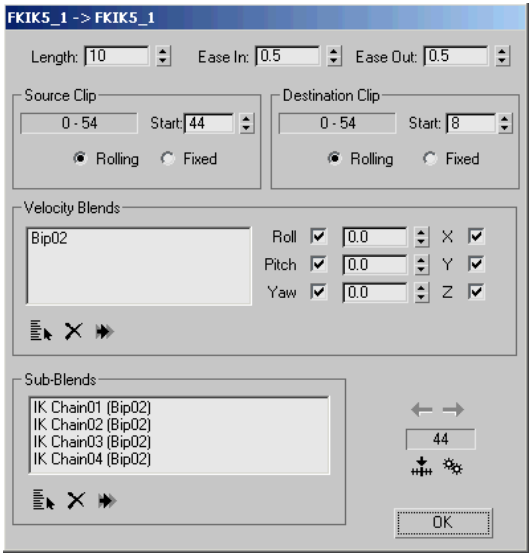
The options on this dialog are much the same as for Bipeds, except that:

- Mix transitions for 3ds Max objects do not have ghosting or playback controls.
- You can explicitly pick blend objects from any object being mixed in the track.
- The sub-blend capability provides further tuning of the transition.

Blending and Sub-Blending

Consider a character walking around a corner. The trajectory of the character's root describes an arc, while the IK feet cycle along the ground. This clip contains a few footsteps, and is used to make a mixer cycle that should take the character around in a circle. If you set the root and the feet all to be velocity blends, the transition will result in the root continuing around the corner while the feet continue in whatever direction they were moving at transition time! That is, if the right foot was just about to hit the ground with its heel, thus pointing up a bit, it would transition upward into the next clip instead of blending in the direction that the root's arc is taking. If you set the root to be a velocity blend and then set the feet to be sub-blends of the root, the transition results in the root continuing on its arc around the corner but with the feet blending in relation to it. That is, the upward-pointing right foot motion will blend with the foot motion in the new clip but in the direction of the arc, not in its own upward direction.

Interface



Length—Sets the number of frames for the duration of the transition.

Transitions are calculated by matching velocities in both clips. Smooth out abrupt velocity changes using longer transitions.

Ease In—Ease-in value for the source clip.

Ease Out—Ease-out value for the destination clip.

Source Clip and Destination Clip groups

The following options let you set different parameters proper to both source and destination clips involved in the transition.

Start Frame—Set the transition start frame for the source and destination clips in their respective fields. Duration for the source and destination clips displays above the Start Frame fields.

- **Rolling—**Keeps the clip in motion during the transition.
- **Fixed—**For the source clip, this option freezes motion at the Start Frame position during the transition. For the destination clip, this option freezes motion at the End Frame position.

  If Fixed is chosen for both the source and destination clips, the transition is a gradual interpolation from one frozen pose to another.

Velocity Blends group

A Velocity Blend object is one that has its positions blended based on velocity (like a biped's COM). Objects used for Velocity Blending are most commonly the roots of the mix hierarchy. When you pick an object for velocity blending, its animation will accumulate based on its trajectory in the scene; for example, choosing the root of a character as the blend object will ensure that the character's animation is transitioned relatively from clip to clip, so the character keeps moving through world space instead of repeating each clip as absolute. All children of the root will follow in the transition.

In some cases, you might want to use velocity blending for a child object.

By default, the Transition Editor uses no velocity blend objects.

Important: If you use velocity blending, pick your blend objects *before* you optimize the transition.

Velocity Nodes list—Lists the objects you have chosen for velocity blending.

 **Select Velocity Nodes—**Click to display a *Pick Nodes dialog (page 2–641)* and choose the objects to velocity blend.

 **Delete Velocity Nodes—**Click to delete the highlighted objects from the list.

 **Set Other Transition Velocity Nodes—**Copies the set of velocity blend objects to all other transitions in the current track.

Roll—When on, accumulates the roll (X) angle during the transition. When off, the angle is simply blended with the next clip. Default=on.

The spinner value lets you specify an additional angle to add to the roll.

Pitch— When on, accumulates the pitch (Y) angle during the transition. When off, the angle is simply blended with the next clip. Default=on.

The spinner value lets you specify an additional angle to add to the pitch.

Yaw—When on, accumulates the yaw (Z) angle during the transition. When off, the angle is simply blended with the next clip. Default=on.

The spinner value lets you specify an additional angle to add to the yaw.

X/Y/Z—When on, accumulates position in the specified axis during the transition. When off, the position is simply blended with the next clip. Default=on.

Sub-Blends group

In order for other world-space objects to transition within the space of the velocity blend or a child of the velocity blend, you need to specify a sub-blend. For example, think of a rig that has a root moving through world space, and also has world-space IK foot controls. In this case, you would pick the root as the velocity blend, then select the root in the list and add the foot controls as its sub-blends. The result of this will be the that root makes a smooth transition and the feet make as smooth a transition as possible in the direction that the legs are going (within the root's space) at transition time, not in the direction that the feet themselves are going at transition time.

Sub-Blend Nodes list—Displays the objects chosen for sub-blending.

 **Select Velocity Nodes—**Click to display a *Pick Nodes dialog (page 2–641)* and choose the objects to sub-blend.

Before you pick sub-blend nodes, you must first choose a velocity blend node in the upper list. Typically this is the parent of the nodes that will be sub-blends.

 **Delete Velocity Nodes—**Click to delete the highlighted objects from the list.

 **Set Other Transition Velocity Nodes—**Copies the set of velocity sub-blend objects to all other transitions in the current track.

 **Previous Transition—**Goes to the previous transition in the transition track.

Displays the previous transition in the Transition Editor, moves the time slider to the start frame of the previous transition and highlights the previous clip in the transition track.

 **Next Transition—**Goes to the next transition in the transition track.

Displays the next transition in the Transition Editor, moves the time slider to the start frame of the next transition and highlights the next clip in the Scripts list.

Current Frame—Displays the current frame number.

 **Go To Start Frame—**Moves the time slider to the first frame of the transition.

 **Optimize Transition—**Displays the *Transition Optimization dialog (page 2–641)*.

Options in the Transition Optimization dialog allow you to search for the location for the transition.

Important: If you use velocity blending, pick your blend objects *before* you optimize the transition.

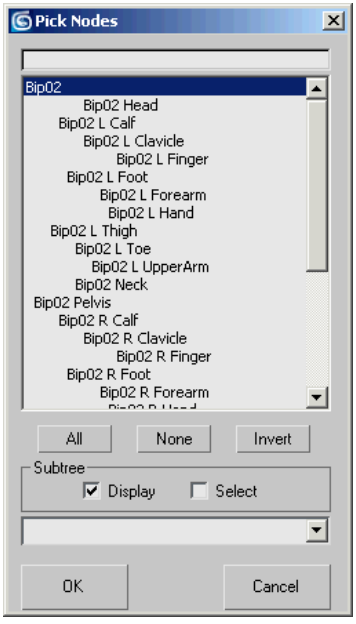
OK—Confirms the current settings and closes the dialog.

Pick Nodes Dialog (Motion Mixer)

Graph Editors menu > Motion Mixer... > Select a non-biped object clip transition. > Motion Mixer menu bar > Transitions > Edit > Mixer Transition Editor > Click one of the Select Velocity Nodes buttons.

This dialog appears when you pick nodes for velocity blending or sub-blending. If you click the Select Velocity Nodes button associated with the Velocity Blends list, the list displays only the objects being mixed in the current transition track. If you click the Select Velocity Nodes button associated with the Sub-Blends list, the list displays only the objects being mixed in the current transition track, *except* those objects that are already velocity blends.

Interface



Search field—To find an object in the list, enter the object's name and then press Enter.

Node list—Lists all the available objects.

All—Selects all objects in the list.

None—Selects no objects in the list.

Invert—Inverts the current list selection.

Subtree group

Display—When on, the object hierarchy is shown by indenting. When off, all objects appear at the same level. Default=on.

Select—When on, selecting an object selects that object's children. When off, only one object at a time is selected. Default=off.

Named selection sets drop-down list—This list shows the scene's *named selection sets (page 1–83)*. Choosing a set from this list selects its members in the Node list.

Transition Optimization Dialog

Motion Mixer > Create a transition track. > Place clips on the track. > Right-click the transition. > Optimize > Transition Optimization dialog

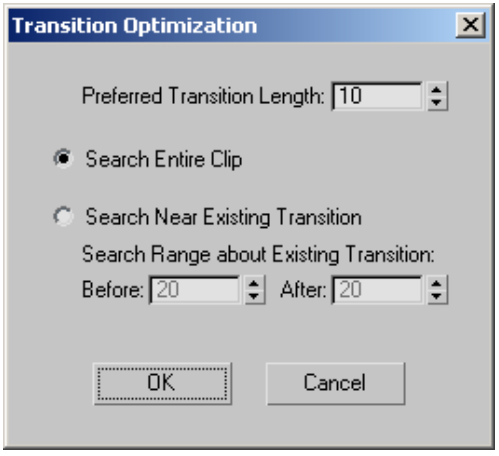
Motion Mixer > Create a transition track. > Place clips on the track. > Select the transition. > Transitions menu > Optimize > Transition Optimization dialog

Options in the Transition Optimization dialog allow you to select the range over which the optimize algorithm will search for a transition. It can search either the whole clip, or it can search near the existing transition. You must specify the preferred length of the optimized transition. The software will try to get as close to that length as possible, still opting to give you the best length.

If it searches about the existing transition, you must specify the number of frames about which it will search before and after the existing transition.

Optimized transitions compute for minimum foot sliding over the range of the transition. This method of determining transitions yields very high quality results.

Interface



Preferred Transition Length—Specify the preferred length of the optimized transition.

Search Entire Clip—Search the entire clip for an optimized transition start frame.

Search Near Existing Transition—Allows animators to find good transitions that "are in the ballpark" of their existing transitions. Transitions can be fine tuned without drastically changing the animation's timing. In general, there may be several "good transitions", so this feature narrows the range of search to the ones that are near to the current settings.

Search Range about Existing Transition—Allows animators to set a "before" and "after" range of frames to perform a search for other transitions that are close to the selected transition. Both ranges default to 20 frames around the selected transition.

Before—Set a frame value to search before the existing transition.

After—Set a frame value to search after the existing transition.

Motion Mixer Toolbar

Select a biped. > Motion panel > Biped Apps rollout > Mixer

Graph Editors menu > Motion Mixer...

Like an audio mixer where you add music tracks and use fades and blends to transition between tracks, the Motion Mixer allows you to mix motion clips. The Motion Mixer toolbar offers commands that you use to add and delete biped and non-biped objects to and from the mixer, modify clips, and change the way tracks in the mixer are displayed.

Interface



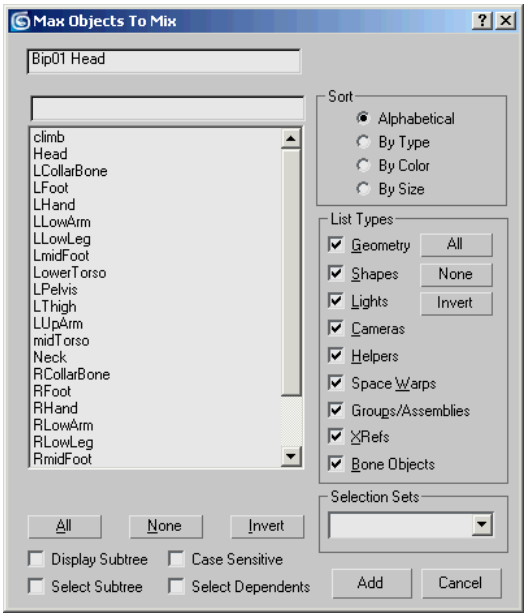


Add Max Objects—Opens the Max Objects To Mix dialog, which lists all non-biped objects in the scene. This dialog lets you add objects to be driven by your mixed animation, as well as name the mix in which they are contained.

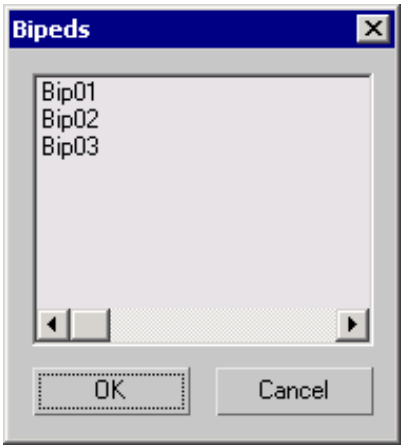
This dialog is comparable to the *Select Objects dialog (page 1–78)*.

Note: The default mix name is the root node with the most children in the list.

Note: Objects can be added to only one mix at a time. Once an object is part of a mix, it is removed from the list.



Add Bipeds—Opens the Bipeds dialog listing all the bipeds in the scene. Highlight one or several bipeds to load them into the Motion Mixer window.



Tip: You can click-drag over all, or several adjacent bipeds in the list if you want to add multiple bipeds to the Motion Mixer. If the bipeds you want to add are not next to one another, hold down the Ctrl key while choosing the bipeds you want to add.

Delete—Deletes a selected biped from the Motion Mixer window.

Select—Lets you select bipeds, trackgroups, tracks, and clips.

Move Clips—Allows selection and movement (including Shift+Clone) of bipeds and trackgroups, keeping clips and transitions intact. It also allows independent horizontal movement of clips on the same track or vertical movement of clips from one biped's track to another biped's track.

Slide Clips—Allows horizontal movement of clips on the same track or vertical movement of clips from one biped's track to another biped's track. Sliding a clip to the right will move unselected clips which are situated to the right of the selected clip; sliding to the left will move unselected clips which sit to the left of the clip.

Offset—Records how many frames have been displaced. This value appears in the status area next to Slide Clips only during interactive translations, such as moving a clip or editing a time warp. The Offset value resets to zero when the translation is complete.

Frame—Displays the frame number under the cursor during interactive translations, such as stretching a clip or editing a time warp. When stretching a clip, you can use both the Frame and Offset fields to tell you how many more or fewer frames the clip will have.

Trim Clips—Clips can be trimmed from their original lengths interactively. When using this mode, you can grab the edge of a clip and drag it to a new frame within the clip's original length. This mode will display all trimmed portions of clips in the Motion Mixer as gray. See *Adjusting Clip Timing (page 2–615)*.

 **Editable Time Warps—**Clips can be interactively time warped, given that they have had a time warp applied. When in this mode, you can place seams in a clip and then drag them within the clip's original length. See *Adding Time Warps (page 2–620)*.

 **Draggable Tracks—**Enables vertical movement of tracks with clips. Holding down the Shift key while moving a track creates a clone of the original track.

 **Lock Transitions—**When Lock Transitions is turned off (the default), moving a clip or changing its length on a transition track affect the lengths of transitions around the clip. When Lock Transitions is turned on, transition lengths do not change when you move a clip or drag clip edges not touching the transition. You can still change a transition's length by dragging the edges of the transition itself.

 **Set Range—**Matches the active segment's start and end times to the extents of the mix.

 **Pan—**Pans the Motion Mixer display horizontally and vertically.

 **Zoom Extents—**Stretches or shrinks the mix to fit in the current Motion Mixer display.

 **Zoom—**Allows horizontal stretching and shrinking of the display of the mix.

 **Zoom Region—**Zoom into a portion of a track by dragging a region selection around that part of a track.

 **Snap Frames—**An on/off toggle that sets the Motion Mixer to snap all adjustments to single frames.

 **Snap Clips—**An on/off toggle that causes clips residing on the same track to snap together. The end frame of one clip will match the start frame of the next clip. If the first clip ends at frame 100, the start frame of the next clip will snap to frame 100.

When Snap Clips is on, transition edges snap to clip edges when dragged.

 **Preferences—**Opens the *Mixer Preferences dialog (page 2–651)*, where you can change display settings for clips, transitions, and frames.

 **Reservoir—**Opens the *Reservoir dialog (page 2–649)*.

 **Horizontal—**This button is available when the Weighting button on the motion mixer toolbar or on a trackgroup is active. It locks all movement of weight nodes to the horizontal.

 **Vertical—**This button is available when the Weighting button on the motion mixer toolbar or on a trackgroup is active. It locks all movement of weight nodes to the vertical.

 **Horizontal Vertical—**This button is available when the Weighting button on the motion mixer toolbar or on a trackgroup is active. It allows free horizontal and vertical movements of weight nodes.

 **Weight Mode—**An on/off toggle that allows you to work with weight curves. Weight curves are useful for making smooth blends between tracks. The Weight Mode button on the Motion Mixer toolbar acts as a global switch that toggles the Weight Mode buttons on every track in the Mixer. The Weight Mode button at the right end of each track toggles weighting only for that track. For more information on how weight curves work, see *Motion Mixer Editor (page 2–646)*. To find out how to use weight curves, see *Adjusting Track Weight (page 2–619)*.

See also

Adding Tracks to the Mixer (page 2–607)

Using the Reservoir (page 2–626)

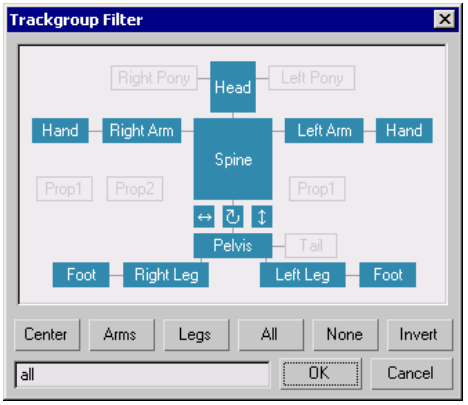
Trackgroup Filter Dialog (Biped Object)

Select a biped. > Motion panel > Biped Apps rollout > Mixer > Trackgroups menu > Filter

By default, the Motion Mixer expects that a mix you create will affect the entire biped. If you like, you can *filter (page 2–612)* a trackgroup so clips in the mix affect only selected body parts. This functionality makes it easy for you to assign one set of BIP files to control the motion of a biped's upper body, while another set controls the lower body.

For more information and procedures on the use of trackgroup filters, see *Filtering Mixer Tracks (page 2–612)*.

Interface



The Trackgroup Filter dialog consists of a rough diagram of a biped. Selectable components display in the color specified for the trackgroup by the Track Color command on the *Mix menu (page 2–629)*. If the biped does not include parts like ponytails or props, they are unselectable. Body parts can be individually activated or deactivated by clicking them, or you can use the buttons across the bottom of the Trackgroup Filter diagram to expedite selections.

Center—Activates all central body parts, comprised of the spine, pelvis and horizontal movement, rotation and vertical movement icons.

Arms—Activates the biped's arms, comprised of the right arm, left arm, and both hand icons.

Legs—Activates the biped's legs, comprised of the right leg, left leg, and both foot icons.

All—Activates all selectable body parts.

None—Deactivates all body parts.

Invert—Inverts the current body part selection. If everything is active, this deactivates everything, as if you clicked the None button.

Trackgroup Label Field—Use this field to specify a more descriptive label for a trackgroup. When you first create a trackgroup, it is labelled All. The new label entered here will appear in place of All to the left of the trackgroup. Use a unique, descriptive name to help you keep track of what body parts were filtered for the trackgroup. For example, if you selected only the pelvis and legs, you could enter **Lower Body** as the trackgroup label.

Trackgroup Filter Dialog (Non-Biped Object)

Graph Editors menu > Motion Mixer... > Trackgroups menu > Filter

By default, the Motion Mixer expects that a mix you create will affect the entire non-biped object. If you like, you can *filter (page 2–612)* a trackgroup so clips in the mix affect only selected body parts. This functionality makes it easy for you to assign one set of XAF files to control the motion of an

object's upper body, while another set controls the lower body.

For more information and procedures on the use of trackgroup filters, see *Filtering Mixer Tracks (page 2–612)*.

Interface



Search field—To find an object in the list, enter the object's name and then press Enter.

Object list—The Trackgroup Filter dialog displays a list of objects, with the hierarchy indicated by indenting (if Subtree > Display is on, which it is by default). Objects can be individually activated or deactivated by clicking them, or you can use the buttons across the bottom of the Trackgroup Filter diagram to expedite selections.

All—Activates all selectable

None—Deactivates all body parts.

Invert—Inverts the current body part selection. If everything is active, this deactivates everything, as if you clicked the None button.

Subtree group

Display—When on, all body parts are listed following a hierarchy structure. Otherwise, body parts are listed on the same level. Default=on.

Select—When on, highlighting a body part automatically select all its children.

Named selection sets drop-down list—This list shows the scene's *named selection sets (page 1–83)*. Choosing a set from this list selects its members in the Object list.

OK—Confirms selection and closes the dialog.

Cancel—Disregards current selection and closes the dialog.

Motion Mixer Editor

Select a biped. > Motion panel > Biped Apps rollout > Mixer

The Motion Mixer window contains an editor where you manage everything that is added to the mix: bipeds, trackgroups, tracks, clips, and transitions.

Interface

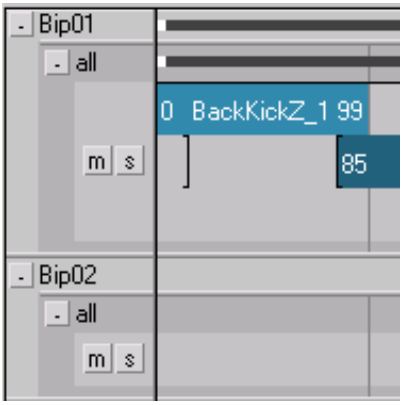


The editor contains three sections:

- *Biped/Trackgroup/Track Controls (page 2–647)*
- *Interactive Clip Controls (page 2–647)*
- *Weighting Controls (page 2–648)*

Biped/Trackgroup/Track Controls

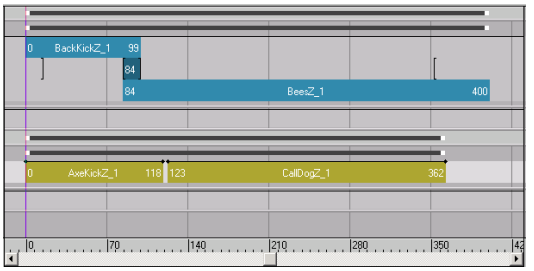
The left-most section of the editor is the Biped/Trackgroup/Track Controls section. The Biped/Trackgroup/Track Controls set the number, order, display, and characteristics of those features in the mix. Bipeds and trackgroups can be selected, added, removed and repositioned.



The "-" and "+" buttons collapse and expand the display of bipeds and trackgroups. The "m" button mutes a track, taking its influence out of the mix. The "s" button sets a track to solo mode, making it the only active track in the mix.

Interactive Clip Controls

The center section of the editor, Interactive Clip Controls, lets you manipulate the scale, time, timing, and trimming of clips and transitions.

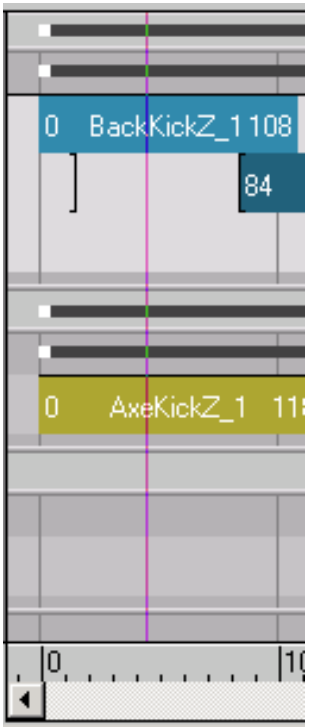


The overall length of all clips and transitions in a track can be stretched or squeezed by dragging the white ends of the gray range bars along the top of each trackgroup. Drag from the middle of the range bar if the entire mix of a trackgroup needs to be moved. Likewise, clips can be moved by dragging from the middle of a clip. Clips and transitions can be shortened or lengthened by dragging either end.

The range bar for each biped can also be moved or scaled to move or scale the entire mix for the biped.

Transition brackets appear in Transition tracks. Where a transition exists, the brackets indicate the start and end point of the transition. If there is no transition at the end of the clip, the bracket indicates where the transition will be placed when a new clip is loaded into the track.

The current time in the scene can be changed by dragging the vertical, purple timeline, a change which is reflected on the 3ds Max time slider.



The vertical timeline.

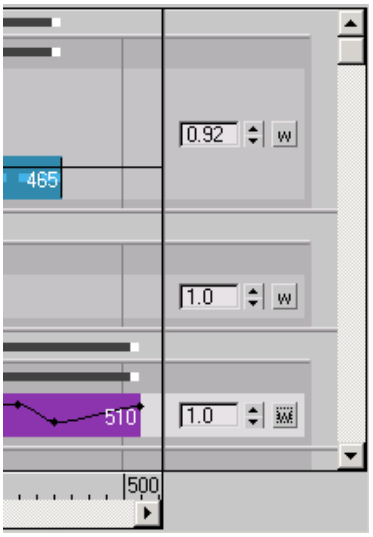
If a clip has a time warp applied, time can be distorted interactively in Time Warp mode: first create seams in the clip, then drag the seams to stretch or squeeze time.



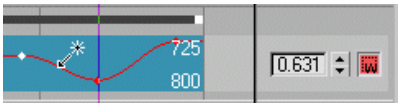
Note the corresponding changes in the frequency of the time ticks down the center of the clip. Time warps are used to speed up or slow down a motion.

Weighting Controls

The weighting controls at the rightmost end of the editor let you set the weight with which a clip or track will be blended with other tracks within the same trackgroup.



When Weight Mode is turned on, weights can be adjusted with a red weight curve on a clip or track.



On a Layer track, each clip has its own weight curve. Each Transition track has one weight curve for the entire track. By default, each weight curve has a node at each end that can be moved to change the weight at that point. Click on the weight curve to add more nodes, and drag a node to move it.

Weights can range from 0.0 (no weight) to 1.0 (full weight). The weight for the currently selected node appears to the left of the Weight Mode button. You can change the weight by moving the node or by changing the spinner value.

When you have mulitple weight nodes selected, and you use the spinner to set a new value, weight node values are changed relative to their original values. In this case, you cannot spin the spinner lower than 0 or higher than 1 on any one spin. If you have more than one weight node selected and

you type in a weight, the weights of all nodes are changed to the new value.

Weighting is evaluated across tracks in a single trackgroup. If a trackgroup has two or more tracks, the topmost track's weight is evaluated at each frame. If the weight at any frame is less than 1.0, the motion on the track is only partially used, and next track down is evaluated for its weight. If the total weight is still less than 1.0, the next track down is evaluated, and so on. In this way, motion from multiple tracks can be mixed at the same frame for the same set of body parts.

The Weight Mode button at the right end of the main toolbar turns on this mode for all tracks in the Motion Mixer. For more information, see *Motion Mixer Toolbar (page 2–642)*.

b On a balance track, the Balance Weight Mode button appears rather than the Weight Mode button. When in Balance Weight mode, you can adjust the weight curve for the *balance track (page 3–914)*. To find out how to use Balance Weight mode, see *Adjusting Biped Balance in the Mixer (page 2–622)*.

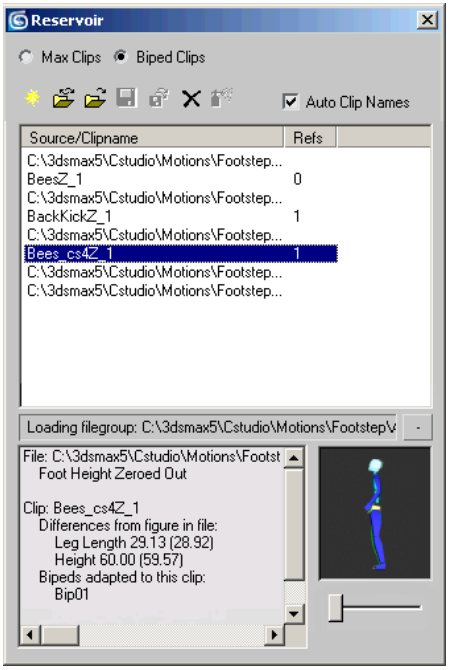
Reservoir

Select a biped. > Motion panel > Biped Apps rollout > Mixer > Reservoir

Graph Editors menu > Motion Mixer... > Reservoir

The Motion Mixer *Reservoir (page 3–1003)* is similar to the *Motion Flow Graph (page 2–1045)*, in that it acts as a holding area where clips can be listed, giving you a palette of motions to choose from. As you add *clips (page 3–975)* to objects in the Motion Mixer window, you see them appear on a *track (page 3–1023)* and they are automatically added to the Reservoir. If you add new clips directly to the Reservoir, they are available and held for future use.

Interface



The Reservoir comprises four parts:

- The *Reservoir toolbar (page 2–650)*, where all the commands are located. You can choose between listing biped clips and non-biped clips.
- The *Source/Clipname list (page 2–650)*, where motion clips (BIP and XAF files) are listed. *Instances (page 3–958)* are grouped together in one listing, while *adaptations (page 3–905)* are listed separately.
- The *Clip Status field (page 2–650)*, where you can gather information about the clips used in the Motion Mixer.
- The *Motion Preview (page 2–651)*, where you can see what a motion looks like before you use it.

  Note: The Motion Preview is available only when you choose the Biped Clips option.

Reservoir toolbar

The commands on the main Reservoir toolbar allow you to manage the clips you have in the Motion Mixer and also build up a collection of clips that you may want to use in the future.



Max Clips/Biped Clips—Choose to switch between a list of biped and non-biped motion assets. The Max Clips option displays XAF files while Biped Clips displays BIP files.


New Entry—Opens a dialog where you can choose BIP or XAF files to load into the Reservoir. You can highlight multiple files to load by holding down the Ctrl key while you click.


Reload Files—Reloads selected BIP or XAF files listed in the Reservoir.


Load File—Opens a dialog where you can choose a new BIP or XAF file to replace the selected entry.


Save File—Allows you to save a highlighted clip from the Reservoir. The Save As dialog is displayed and you can specify a name and location for the clip you are saving.


Batch Save—Opens the Save Reservoir Items dialog, where you can choose the clip or clips to save.


Delete—Removes a highlighted clip from the Reservoir. All instances of the clip are also removed from the Motion Mixer.


Clean—Similar to Delete, but only removes clips from the Reservoir that are *not* used in the Motion Mixer.



Auto Clip Names—Names the clip based on the name of the motion file. Turn off to name a clip yourself.

Source / Clipname list

The Source / Clipname List displays all the clips you have in the Motion Mixer Reservoir. Four pieces of important information can be gathered from this list:

- Drive location where the clips are stored.
- Which clips are currently being used in a mix.
- How many times each instance of a clip appears in the Mixer (*Refs*).
- Which clips are simply being held in the reservoir, as opposed to being used in the Mixer.

As you add clips to a track, from the *Tracks menu (page 2–632)* or by right-clicking a track, they are automatically added to the Reservoir. If the clip is a new *adaptation (page 3–905)*, the name of the clip is assigned an incremental suffix, such as *1* or *2*.

The Refs (Reference) column tells you how many times a clip *instance (page 3–958)* has been used in the Motion Mixer, or if it has not been used at all. If an instance shows a Refs value of 1 or more, it means the instance has been used that many times. If an instance shows a Refs value of zero, it means the clip was once used by one or more bipeds, but is not currently used on any tracks.

If a clip has been added to the reservoir, but has never been placed in the Motion Mixer, only the source path will be listed and no Refs value is displayed.

Clip Status field

The Clip Status field appears at the lower left of the Reservoir and gives you information about how a particular clip has been used. You can get file information, such as the source location of the BIP

or XAF file. Clip information is also displayed, like the name of the clip reference as it appears in the Motion Mixer. Information about the biped structure is also listed along with the bipeds to which the clip has been adapted.

Motion Preview

The Motion Preview at the lower right side of the Reservoir shows a thumbnail of the selected clip. Moving the slider across the bottom gives you an idea of what the motion looks like before you add it to a track.

Note: In order for a clip to show in the Motion Preview, it must be a BIP file that has been saved for this version of 3ds Max. If you choose a BIP file that is from an older version of 3ds Max, you will see a warning that the file is obsolete and should be resaved. Until the BIP file is resaved, it will not display in the Motion Preview.

Note: The Motion Preview is available only when you choose the Biped Clips option.

See also

Using the Reservoir (page 2–626)

Mixer Preferences Dialog

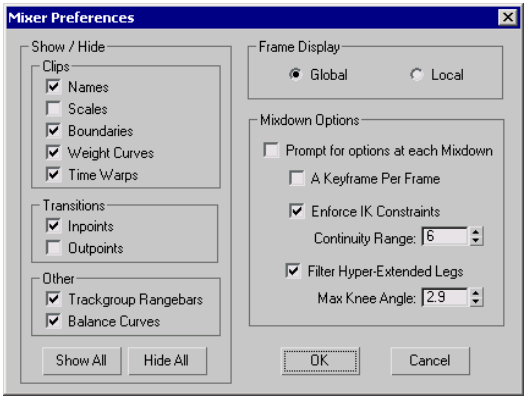
Select a biped. > Motion panel > Biped Apps rollout > Mixer > Preferences

Graph Editors menu > Motion Mixer... > Preferences

The Motion Mixer Preferences dialog lets you make changes to display settings for clips, transitions and frames. It is similar to the way you can Hide By Category on the Display panel. You can also specify mixdown options from this dialog.

The settings on the Motion Mixer Preference dialog are saved in the *biped.ini* file, as well as with the MAX file.

Interface



Show / Hide group

These settings affect how clips appear in the Motion Mixer tracks.

- **Names—**When turned off, clip names do not appear on the colored clip bars in the Motion Mixer. Default=On.
- **Scales—**Displays the clip scale. Since clips can be resized by using tools like Move Clips, displaying the scales quickly lets you know if a clip is two times its original length or half its length. Default=Off.
- **Boundaries—**Turns on and off the frame numbers at the start and end of the colored clip bars. Default=On.
- **Weight Curves—**Toggles the display of the weighting line. It doesn't matter if the Weight button is active or not. Default=On.
- **Time Warps—**If a time warp has been added to a clip, this switch toggles the display of the warp on the colored clip bar. Default=On.

Show / Hide Transitions group

The two settings affect how transitions appear in the Motion Mixer tracks.

- **Inpoints—**Toggles the beginning transition frame on the transition clip bar. Default=On.

- **Outpoints—**Displays the end transition frame on the transition clip bar. Default=Off.

Show / Hide Other group

These settings indicate if range bars and balance curves appear in the Motion Mixer tracks.

- **Trackgroup Rangebars—**Hides the gray range bar that appears along the top of each trackgroup. Default=On.
- **Balance Curves—**Displays the Balance Curves track. Default=On.

Frame Display group

These settings affect how clip boundaries and transition in/outpoints appear in the Motion Mixer.

Global—Displays clip start and end frames as frame numbers. Inpoints and outpoints of transitions display the frame at which transitions start and end. Default=On.

Local—Clip boundary values are displayed as actual lengths in frames. Along with the start and end frame number of the transition, additional values tell you by how much the transition overlaps each motion clip.

Mixdown Options group

These settings affect how transitions between biped object clips appear in the Motion Mixer tracks.

- **Prompt For Options At Each Mixdown—**When turned on, the Mixdown Options dialog is displayed when you choose Compute Mixdown from the Mix menu in the Motion Mixer. Default=On.
- **A Keyframe Per Frame—**When turned on, a keyframe is generated for each frame of animation when the mixdown is computed. Default=Off.
- **Enforce IK Constraints—**When a transition occurs between two clips where the same foot or feet are planted with footsteps or planted keys, this option forces the foot/feet to stay planted during the transition. Default=On.
  - **Continuity Range—**Sets an additional transition time after the actual transition, giving the foot time to get from its planted position to its keyframed location in the next clip. Range=0 to 100; Default=6.
- **Filter Hyper-Extended Legs—**Prevents a leg from straightening during a transition. This can occur when the COM moves in such a way that the foot can't reach its planted position unless the leg goes completely straight. This option corrects the problem by raising the biped's heel off the ground so the knee can bend. Default=On.
  - **Max Knee Angle—**Sets the maximum angle that can be reached between the thigh and calf before the heel comes off the ground. An angle of 180 indicates a perfectly straight leg. A value of 160 or 170 indicates a slightly bent leg.

Mixer Rollout

Select a biped. > Motion panel > Biped rollout > Mixer Mode

On this rollout, you can load and save *MIX files (Motion Mixer files) (page 3–973)* . A saved MIX file can be loaded onto any biped in the scene, and motions in the mix will be automatically adapted to the biped's size.


The Mixer rollout appears when Mixer Mode is turned on from the Biped rollout.

Interface



Load File—Loads a Motion Mixer file (*.mix*). These files include the following, which display in the Motion Mixer window when a MIX file is loaded:

- **Trackgroups—**Groups of *tracks (page 3–1023)* for selected parts of the biped.
- **Tracks—***Layer tracks (page 3–961)* and *transition tracks (page 3–1027)* where clips and transitions reside.
- **Clips—**References to BIP animation files used in the mix.
- **Transitions—**Connections between clips on transition tracks.

Note: Loading a MIX file does not automatically open the Motion Mixer. Click Mixer on the Biped Apps rollout to open the Motion Mixer and see the loaded mix.

Save File—Saves the currently selected biped's mix in the Motion Mixer to a MIX file.

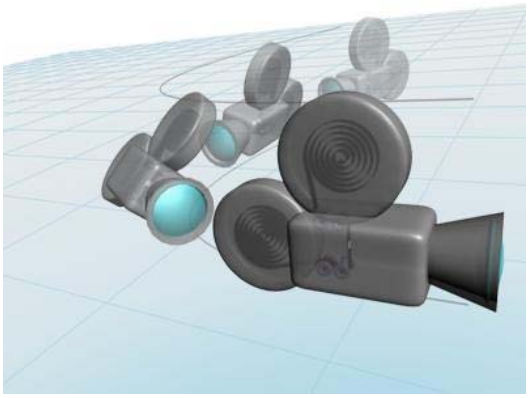
Animation Utilities

3ds Max provides a number of utilities, available from the *Utilities panel (page 3–778)*, to assist in animating scenes:

- *Follow/Bank Utility (page 2–653)* is an alternative to the *Path constraint (page 2–398)*.
- *MACUtilities Utility (page 2–665)* lets you convert an *imported TRC motion (page 3–577)* into CSM format so you can map it onto a biped.
- *Camera Tracker Utility (page 2–667)* synchronizes a background by matching the movement of a 3ds Max camera to a real camera.
- *Camera Tracker Utility (page 2–667)* synchronizes a background by matching the movement of a 3ds Max camera to a real camera.
- *Dynamics Utility (page 2–686)* generates animation by simulating real-world physics.
- *Skin Utilities (page 2–700)* provide a method for copying skin data (envelopes and vertex weights) from one model to another.

Follow/Bank Utility

Utilities panel > Utilities rollout > More button > Utilities dialog > Follow/Bank



The Follow/Bank utility applies rotation keys to an object that already has a motion trajectory. The settings are similar to those found in the Path constraint. Use this when you have an object for which you've already assigned motion and want it to follow or bank along its trajectory, but don't want to use a Path constraint.

The utility aligns the object's X axis to the direction of motion. To change the orientation, change the

pivot orientation in the Hierarchy > Pivot panel before using this utility.

Although many of the settings are the same as in the Path constraint, this utility is not dynamic; you have to click the Apply Follow button each time you change an object's animation or the utility settings. When adjusting banking settings, move to the point at which the object is at the sharpest turn, then alternately click Apply Follow and adjust the banking settings to get the right affect at the extreme corner.

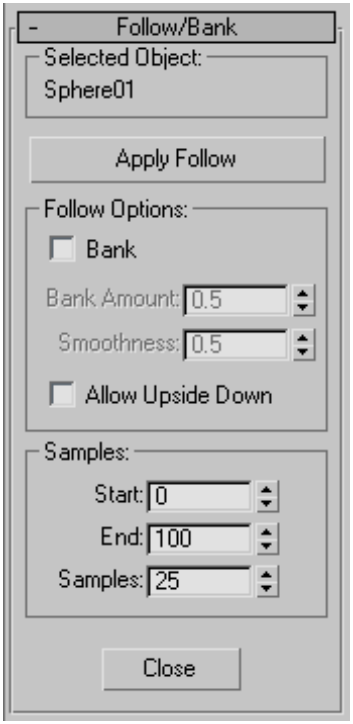
Tip: You can also use an Orientation Constraint to manually control the rotation of the object as it moves through space. While this is not automatic, it allows more fine control on individual frames than other methods.

Procedure

1. On the Utilities panel, click the More button, and choose Follow/Bank from the list.
2. Adjust the options in the Follow/Bank Parameters rollout.
3. Select the objects you want to follow/bank.
4. Click the Apply Follow button.

   The objects orient themselves at the start frame, and follow or bank along their trajectories.

Interface



Selected Object group

Displays the name of the current selection. If more than one object is selected, "Multiple Selected" displays.

Apply Follow button

Apply Follow—Applies the settings and generates rotation keys. Each time you click this button, rotation keys in the specified range of frames (set in the Samples group) are deleted and regenerated.

Follow Options group

Contains the parameters that specify whether or not to use banking, and the banking settings.

Bank—Turn on to cause the objects to bank as well as follow the trajectory.

Bank Amount—Adjusts the amount of the banking to one side or the other, depending on whether the value is positive or negative.

Smoothness—Controls how rapidly the roll angle changes as an object moves through bends in the trajectory. Smaller values make an object more responsive to subtle changes in the curve while larger values minimize jerkiness.

Allow Upside Down—Avoids the situation where an object flips when going around a vertically oriented path.

Samples group

Specifies the range of frames over which rotation keys are generated, and the number of keys to be generated.

Start—Specifies the first frame rotation keys are generated.

End—Specifies the last frame rotation keys are generated.

Samples—Specifies the number of rotation keys to be generated in the specified range.

Close button

Close—Exits the Follow/Bank utility.

Motion Capture Utility

Utilities panel > Utilities rollout > Motion Capture button

The Motion Capture utility drives your animation using peripheral devices, such as MIDI keyboards, joysticks, and the mouse. While driving the animation, you can record it in real time.

The motion capture system in 3ds Max is handled in the following way:

1. In Track View, assign motion capture animation controllers to the specific tracks you want controlled by external devices.
2. After assigning the Motion Capture controller, open the Properties dialog for the track and bind the type of peripheral device(s). As an example, the Rotation Motion Capture controller has three rotational axes to which you can bind one device each.
3. After binding devices, adjust their settings and parameters in the lower portion of the Track Properties dialog. These controls vary depending on the type of device.
4. On the Utilities panel, open the Motion Capture utility. You can test and record your motion for any combination of tracks over any range of frames.

The following motion-capture devices are included with 3ds Max:

Mouse Input Device rollout (page 2–661)

Keyboard Input Device rollout (page 2–661)

Joystick Input Device rollout (page 2–662)

MIDI Device rollout (page 2–663)

Procedures

To assign a Motion Capture controller to tracks:

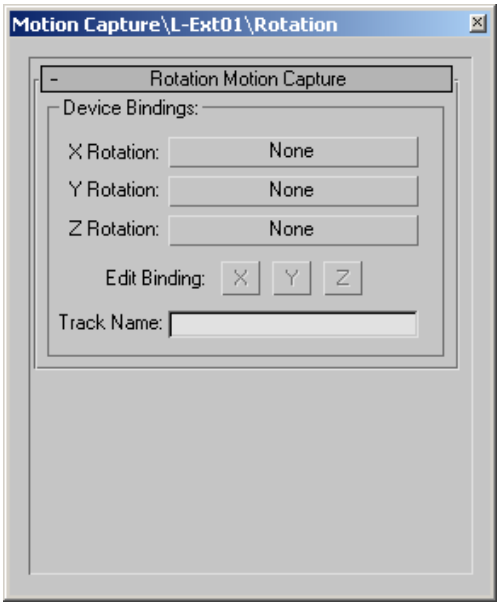
The first step in using motion capture is to assign motion-capture controllers to the tracks you want.

1. In the Track View hierarchy, highlight the Transform track you want.
2. Choose Controller > Assign.

   The Assign Controller dialog is displayed.
3. Choose Motion Capture (might start with "Position," "Float," "Scale," or other), and then click OK.

A Motion Capture dialog is displayed. Its title bar includes the name of the object and the track to which the controller is assigned.

If a Data dialog opens instead, close it, right-click the track name, and choose Properties.



When you assign a Motion Capture controller, the previously assigned controller is maintained as a child of the Motion Capture controller. This lets you continue to adjust the object using standard transform controls, while still making motion-capture control available.

To bind and adjust devices:

1. After assigning a Motion Capture controller, determine the type of device that will drive the motion.

   Depending on the type of controller, you might be able to bind one or more devices. For example, a Rotation Motion Capture controller can have three devices, one for each axis of rotation. On the other hand, a controller for the radius of a cylinder would have only one device to control the radius value.

2. Bind the devices in the Properties dialog for the track.

To specify a device:

1. Click the Properties button while the Rotation track is selected.

   The Properties dialog for rotation motion capture includes three binding buttons (one for each axis.)

2. Click the Y Rotation button. In the resulting dialog, choose Mouse Input Device, and click OK.

   The lower half of the Properties dialog display controls specific to the new device. Assuming this is a segment that isn't attached to another object, you want the horizontal mouse movement.

3. Under Mouse Input, choose Horizontal.

   The Scale spinner lets you adjust the relative effect of the mouse movement to the rotation. Flip reverses the rotational direction with respect to the mouse movement.

Example: To set up the Joystick controller to move a Free camera:

You set up the Joystick controller to move a Free Camera around on the XY plane. The Y movement of the joystick will affect the camera's forward and backward movement, and the X movement of the joystick will rotate the camera. You'll use Increment Based On Direction to allow the Y movement to be local to the camera rather than to the world.

1. Create a Free camera in the Front viewport.
2. In Track View, assign a Rotation Motion Capture controller to the Rotation track of the camera.

3. Right-click in the Rotation track to display the Motion Capture properties dialog.
4. Click the button beside Z Rotation, and assign a Joystick Input Device.
5. Under Joystick Axis, choose the X option (if it's not already chosen).
6. Check Accumulate so that the rotation won't return to zero each time you release the joystick.

Example: To set up the Position controllers of a camera for moving with a joystick:

The rotation of the camera about its world Z axis will be controlled by the X motion of the joystick. You'll set up the Position controllers of the camera to move the camera forward and back with the Y motion of the joystick.

1. Assign a Position Motion Capture controller to the camera's Position track, and access its Properties dialog.
2. Assign a Joystick Input Device to both the X Position and Y Position buttons.
3. Click the Edit Binding X button, and then choose the Y option under Joystick Axis.
4. Check Accumulate.
5. Click the Edit Binding Y button, and also choose Y under Joystick Axis.
6. Check Accumulate.

   The Y motion of the joystick will now affect both the X and Y positions of the camera.

Example continued: To add the increment based on direction options:

If you tested the camera motion at this point, moving the joystick forward or backward along the Y would produce a diagonal motion because the Position tracks would receive equal amounts of X and Y values. Adding the Increment Based On Direction options will change this.

1. Click the X Edit Binding button.
2. Under Increment Based On Direction, click the button beside Controller.

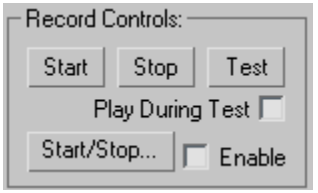
   A dialog appears with a list of controllers.
3. Find and select the Rotation Motion Capture controller for the camera.
4. Set the Direction option to Z (the local axis along which the direction points).
5. Set the Component option to X (matching the Edit Binding button).
6. Click the Y Edit Binding button.
7. Click the button beside Controller, and assign the Rotation Motion Capture controller.
8. Set the Direction option to Z, then set the Component option to Y.
9. Open the Motion Capture utility, turn on both tracks in the list, click the Test button, and move the camera while observing it in the Top viewport.

   As you move the joystick in the Y direction, the camera moves forward or backward. Moving the joystick in the X direction rotates the camera, but then, as you continue moving the joystick in the Y direction, the camera moves forward and back along its local axis.

Interface



Record Controls group



Tests and records your animation. The first three buttons let you control the recording directly, while the Start/Stop button lets you use a MIDI device to control the recording.

Start—Starts a recording using the values set under Record Range below the Track list.

Stop—Stops the recording before the Out frame is reached. You can also stop a recording by pressing Esc, or by pressing the right mouse button.

Test—Tests your motion. No recording takes place. Exit by clicking Test again, pressing Esc, or clicking the right mouse button.

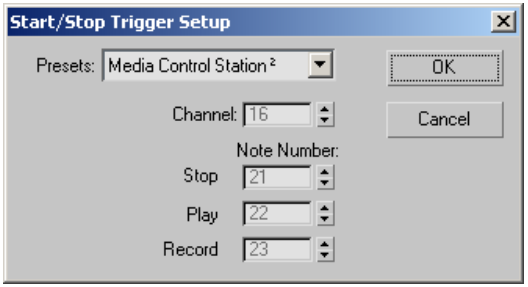
Play During Test—When turned on, and you click Test, the animation in the scene plays in a loop while you test your motion. Note that tracks in the Motion Capture list that are selected (marked red) won't play back because they're waiting for input from the assigned peripheral device.

Start/Stop—Displays the *Start/Stop Trigger Setup dialog (page 2–659)*. You can choose the type of MIDI device that will control the recording.

Enable—Uses the assigned MIDI device for recording instead of the Start, Stop, and Test buttons.

Tip: In the Time Configuration dialog, you can reduce the viewport playback speed, test or record at the lower speed, and then reset the speed to normal to view the results.

Start/Stop Trigger Setup dialog



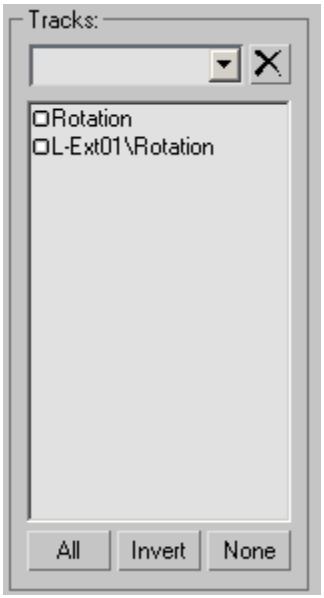
Presets—Choose the type of MIDI device. If you choose Media Control Station 2, you can use its buttons to Stop, Play, and Record. (The Media Control Station is a MIDI device containing standard VCR-style playback buttons along with a jog wheel.) By choosing Custom, you enable the remaining spinners in the dialog, where you can set specific channels and note numbers. You can also use Custom to customize the controls used by the Media Control Station.

Channel—Specifies the channel to which your MIDI device is assigned.

Note Number—The Stop, Play, and Record spinners that follow this label let you specify which note event triggers which function.

- **Stop—**This note event stops the playback or the recording. On the Media Control Station, this would be the square button. This is the equivalent of clicking the Stop button under the Record Controls group.
- **Play—**Plays the animation. This is the arrow button on the Media Control Station.
- **Record—**When this note is sent by itself, it's the equivalent of clicking the Test button under Record Controls. To record the motion capture, you must press both Play and Record at the same time (the equivalent of clicking the Start button under Record Controls). You can press the Play button before or after the Record button and release the Play button before or after the Record button. The recording starts when the button is released.

Tracks group



The Tracks area displays all tracks that have been assigned Motion Capture controllers. You select the tracks that will be affected by the three buttons in the Record Controls group. Click a track to toggle its selection box on or off. Only the selected tracks displaying the red box are affected by the Record Controls group.

You can select the tracks either by clicking them, or by using the All, Invert, None buttons. In addition, you can create named selection sets of tracks. To create a named selection, click the tracks you want, and then enter a name in the Edit field above the track list. To select a previously named set, open the Edit window, and select it from the list. To delete a named set, select it from the list, and then click the Delete button at the right.

Note: A specific track can be in only one named selection at a time.

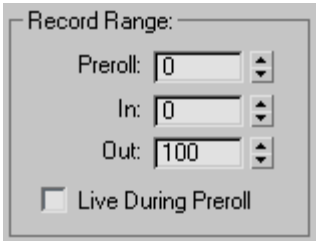
All—Assigns all tracks to the Record Controls group.

Invert—When tracks are selected, assigns the unselected tracks to the Record Controls Area.

None—Assigns none of the tracks to the Record Controls group.

By default, each track is provided a name consisting of the name of the object followed by a backslash and the name of the parameter. For example, Sphere01\Angle is an Angle parameter applied to one of the modifiers assigned to a sphere. You can rename the tracks in the Properties dialog, which you can access in Track View, or by double-clicking the track name in the Motion Capture utility. The upper rollout of each Properties dialog includes a Track Name field. Any name entered in the Track Name field appears in the Track list in the command panel, in the title bar of the Properties dialog, and beside the Data track in the Hierarchy list in Track View.

Record Range group



Specifies the frame range affected by the recording. You can also set a *preroll* range of frames (see below).

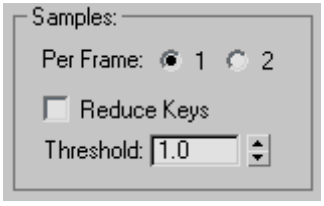
Preroll—Specifies the frame number where the animation begins playing when you press the Start button. If this number is the same or greater than the In number, the Preroll number is ignored. For example, if In: 0 and Out: 200, then a Preroll of -100 would give you 100 frames before the recording started, while a Preroll of 50 would be meaningless. If the range defined by Preroll and Out is larger than the current segment, the segment is temporarily enlarged during the recording.

In—Specifies the frame number where the recording begins after you click Start.

Out—Specifies the frame number where the recording ends after you click Start. You can stop the recording before this by clicking Stop.

Live During Preroll—When this is active, the motion capture is live through the preroll frames. The motion occurs, but isn't recorded. This can cause sudden motion shifts at the In frame because your live motion might not match the animation when you reach the In frame. When this option is inactive, the motion is disabled or frozen until you reach the In frame.

Samples group



When the motion data is captured, it's sampled at a rate relative to the current frame rate, which defaults to 30 frames per second. Increasing the frame rate increases the sample rate. The two radio buttons in this section allow you to choose one or two samples per frame. This is a shortcut that lets you sample at a field rate. If you're rendering to fields, you don't need to capture at 60 frames per second. Controllers will interpolate between the samples.

Reduce Keys—Reduces the keys generated by the motion capture.

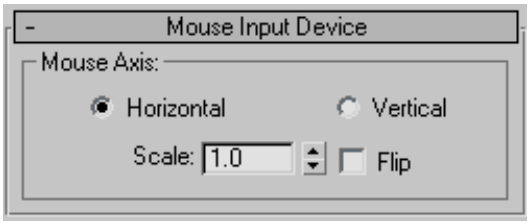
Threshold—The value you enter in the Threshold field constrains how 3ds Max changes the selected track. The difference between the new animation and the original animation, at any frame, will be less than the threshold value.

Low threshold values closely match the original animation but may not greatly reduce the number of keys.

High threshold values produce the fewest keys but may not match the original animation with much accuracy.

Note: When Reduce Keys is active, there is additional calculation time after each completed recording.

Mouse Input Device rollout



Controls animation using the horizontal or vertical motion of the mouse. The available settings include the following:

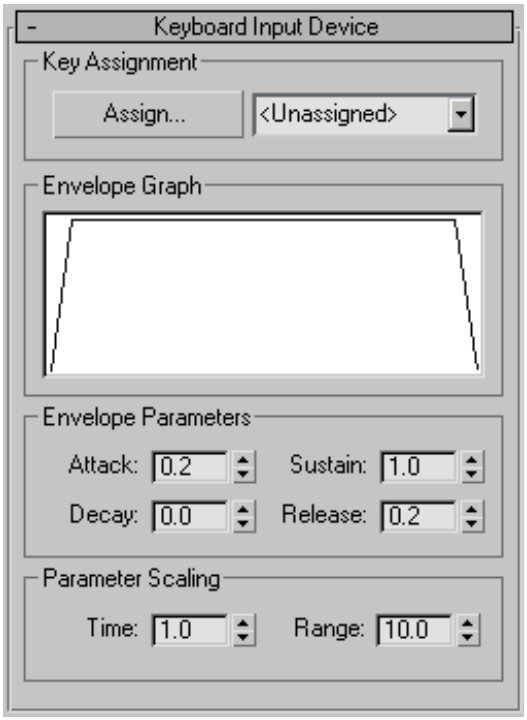
Horizontal/Vertical—Specifies which mouse motion drives the animation.

Scale—Scales the relative effect of the mouse movement to the animation response (Spinner Value: float, 0 to 999,999)

Flip—Flips the direction of the response relative to the mouse movement. For example, if moving the mouse horizontally to the right produces a clockwise effect on a Rotation controller. Activating Flip will reverse the rotation to counterclockwise.

Keyboard Input Device rollout

Lets you assign most keyboard keys to drive the animation.



Key Assignment group

Click the Assign button, and then press any key. The assigned key appears in the list at right. Alternatively, you can open the list and select a key from it.

Assign—Click the Assign button, and then press any key. The assigned key appears in the lower list at right.

Drop-Down List—You can open the list and select a key from it.

Envelope Graph group

Displays a representation of the amplitude curve over time.

Envelope Parameters group

Specifies the time over which the envelope of the action takes effect, relative to the key pressing and release.

Attack —Displays the time it takes after pressing the key for the value to reach its maximum level.

Decay—Displays the time it takes after having reached maximum for the value to fall to that specified by the Sustain spinner.

Sustain—After the Attack and Decay, the value specified here is sustained until you release the key.

Release—After releasing the key, this is the time it takes for the value to fade out to zero.

Parameter Scaling Area—Sets the scale of the envelope and the range of the output value.

Time—Specifies the scale of the Attack, Decay, and Release parameters. The value represents the number of seconds contained in 1 unit. For example, if this value is 1.0, then an Attack value of 1.0 equals 1 second.

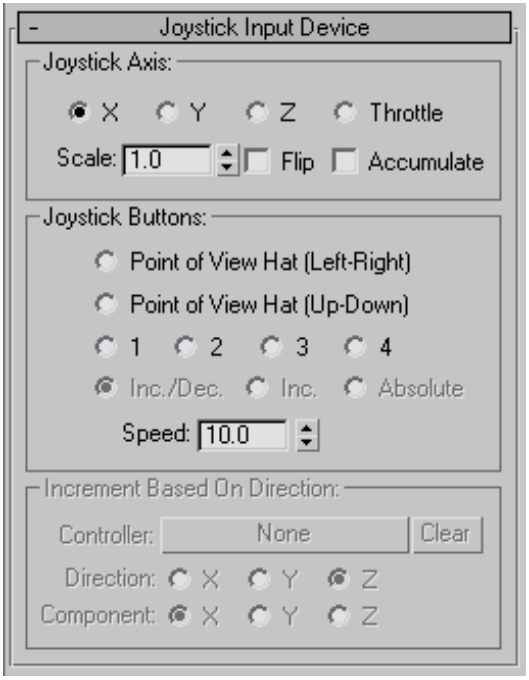
Range—Sets the maximum output value of the controller.

Note: This controller ignores the state of the Ctrl, Alt, and Shift keys.

Tip: Because a single keyboard key can only generate either a positive or negative value, if you want to control both directions of a Bend Angle, for example, use a List controller. First, make sure there's a standard controller (such as a Bezier controller) assigned first in the list to maintain the center position. Then add two Keyboard Motion Capture controllers to the list (one assigned to one key, and the other assigned to a different key.) Set the range of one to the positive extent and the range of the other to the negative extent.

Joystick Input Device rollout

The Joystick Input Device driver was designed for the Microsoft Sidewinder joystick, which contains more controls than the standard joystick. You can use this device driver for standard joysticks as well.



Joystick Axis group

X, Y, Z—Specifies which joystick direction drives the animation. (Standard joysticks provide X and Y axes only. The Sidewinder provides the Z axis when you twist the joystick.)

Throttle—On the Sidewinder, this is a slider next to the stick.

Scale—Scales the relative effect of the joystick action to the animation response (Spinner Value: float, 0 to 999,999)

Flip—Flips the direction of the response.

Accumulate—When turned off, the joystick position represents an absolute position, and

you're limited to the "rectangle" defined by the limits of the joystick. When the joystick returns to its rest position, the value generated returns to zero. When this is turned on, the joystick represents a change in the current position. Moving the joystick forward, for example, can cause an object to start moving, and it will continue to move until you return the joystick to its rest position.

Joystick Buttons group

Point-of-View Hat (Left-Right, Up-Down)—The Sidewinder includes a mini-joystick on the tip of the main joystick. Specify the direction this joystick controls in the animation.

1, 2, 3, 4—Specifies one of four buttons in the Sidewinder joystick. They work similarly to the Point-of-View Hat, except that each button increases a direction value only while pressed. When you release the button, the value returns to zero (centered).

Inc./Dec.—One of three options available only when you choose one of the numbered joystick button options. This option (Increment/Decrement) means that the value is incremented when the button is down, and decremented when the button is up.

Inc.—When turned on, the value increments when the button is down, and stays at that value when you release the button.

Absolute—When turned on, the assigned parameters jumps to the value set in the Speed spinner when the button is down, and then jumps back to zero when you release the button. Use this for on/off animations, such as blinking lights.

Speed—Controls the rate at which the value changes when using either the Point-of-View Hat or the four buttons. When using a button option and the Absolute option, this specifies the value output when you press the button (Spinner Value: float, -999,999 to 999,999).

Increment Based On Direction group

Provides controls that let you derive the direction of movement from a Rotation controller. These options are used primarily when you're animating a first-person flythrough (such as when controlling a camera).

Note: The items in this area are only available when you select Accumulate in Joystick Axis.

Controller—Assigns a Rotation controller from which the direction will be derived.

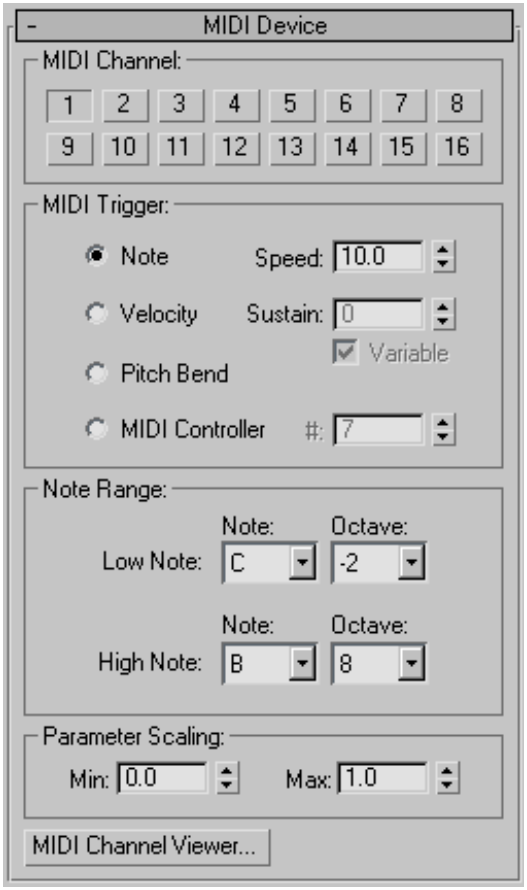
Clear—Removes the assigned controller.

Direction X/Y/Z—Specifies the local axis that will be used as the direction. For a Free Camera, for example, this would be Z because the camera points in the Z direction. However, if you had a car that pointed along its Y axis, you'd use Y.

Component X/Y/Z—Specifies the edit binding to use. Match this to the Edit Binding button under Device Bindings. For example, if the Y Edit Binding button is selected, choose the Y Component option.

MIDI Device rollout

Controls animation using a MIDI device.



MIDI Channel group

Contains 16 buttons. Click to specify the channel to which your MIDI device is assigned.

MIDI Trigger group

Defines the type of MIDI event (message) that will drive the motion. There are four options.

Note: When you choose this option, the note number or pitch defines the output value. The value is derived from where the note falls within the Note Range, specified in the Note Range group. When the note is at the bottom of the range, the value takes on the Min value specified under Parameter Scaling. When the note is at the top of the range, the value takes on the Max value from the same area. Anything in between is interpolated between the Min and Max values. (Note that Min doesn't have to be less than Max.) The generated value will slide around as different keys are pressed. The harder a key is pressed, the faster the value changes. The Speed setting defines how fast, overall, the value changes as keys are pressed.

Velocity—When you choose this option, the velocity of the note pressed determines the output value. The notes set under Note Range merely specify which notes are valid to press. The value takes on the Min value until a key within the Note Range is pressed. When the key is pressed, the value approaches the Max value based on how hard the key was pressed. (The value actually travels along a parabola toward the Max value.) The Sustain setting defines how long it takes the value to move through the parabola. When you choose Variable, the sustain length is also scaled by how hard the key was pressed.

Pitch Bend—The MIDI instrument's pitch bend knob defines the value. The Note Range doesn't apply in this case and is disabled.

MIDI Controller—Specifies a note event when you're hooked up to a different type of MIDI controller than the typical keyboard. For example, if you're using a MIDI slider box, you would select the MIDI Controller option, and then use the # spinner to specify the note event for the specify slider.

Note Range group

Low Note/High Note—Set these to specify the low and high ends of the note range when the Note option under MIDI Trigger is turned on.

Parameter Scaling group

Contains the Min and Max spinners, which specify the range of generated values. See Note and Velocity.

MIDI Channel Viewer

Clicking MIDI Channel Viewer at the bottom of the MIDI Device rollout displays a dialog that lets you test your MIDI device to see which MIDI channel is receiving events, and which notes are being triggered.

MIDI Channel group

Provides a column of 16 buttons and progress bars representing the 16 MIDI channels. Select the channel from which you want to view note activity. The channel progress bars light up when any channel has an event.

MIDI Note group

The 11 Octave buttons let you select which octave you want to view. When a note is played in that octave, a corresponding progress bar lights up in the Note column.

MIDI Controller #—When using a different type of MIDI controller, such as a slider box, you can specify a note event, and then watch the corresponding progress bar light up when you activate that event. You can find the correct note number by activating the event while watching the Note Number field at the bottom of the dialog.

Channel—This is one of four text fields that display all of the values being generated by the MIDI device as you activate an event. The Channel field displays the currently selected channel.

Event—Displays the type of MIDI event being sent. This can be:

Note On: 9

Note Off: 8

Pitch Bend: 14

MIDI Controller: 11

Velocity—Displays the velocity, which has a different meaning, depending on the event. For the most common event, a note being pressed, this value represents the velocity at which the key was struck. Other events, however, might generate a continuous value. For example, a pitch bend event transmits the position of the pitch bend.

Note Number —Displays the corresponding note number for the event. When you're using a non-keyboard MIDI device, such as a slider box, you can use this to identify the note number of a specific slider, for example.

MACUtilities Utility

Utilities panel > More button > MACUtilities

You can use the Motion Analysis Corporation utility to convert motion data originally recorded in TRC format into CSM. This allows you to easily map the motion onto a biped.

A MNM (Marker Name) file is needed to determine how to convert the marker names from the incoming data into the CSM naming convention. You can either load an existing MNM file or create a new one using the MNM Creator/Marker Selection dialog.

Procedures

To convert a TRC file into CSM:

1. Click the Browse... button next to Create.
2. Highlight a MNM file in the Select Mapping File... dialog. Click Open to confirm your selection.

   The MNM file serves as a naming convention reference for the TRC conversion.
3. Click the Browse... button over Convert TRC To CSM.
4. Highlight a folder in the Choose CSM Files Destination Folder... dialog. Click Use Path to confirm your selection.

After the conversion, the resulting CSM file is saved in this folder.

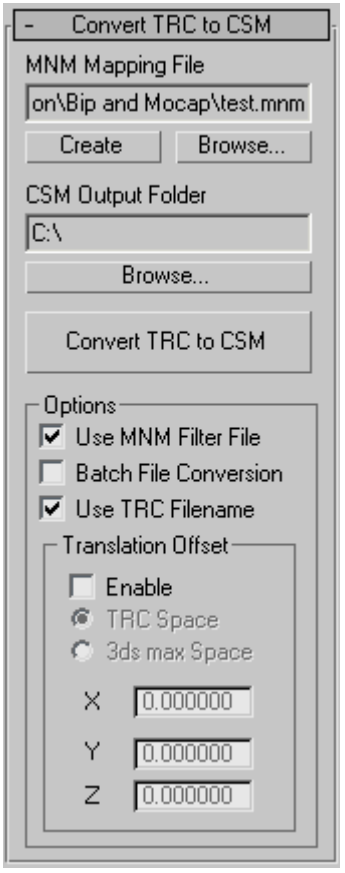
5. Turn on Use MNM Filter File.
6. Turn off Use TRC Filename.
7. Click Convert TRC to CSM.

   A Select Source TRC Files... dialog opens.
8. Highlight a TRC file and click Open to confirm your selection.

   A Selected Target CSM File... dialog opens.
9. Either highlight an existing CSM file to overwrite or enter a new filename. Click OK to confirm your selection.

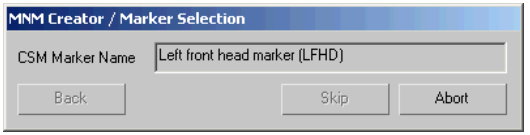
   Your file is converted into CSM format. You can now load it onto a biped.

Interface



MNM Mapping File—Displays the path and name of the chosen MNM file.

Create—Launches the MNM Creator/Marker Selection dialog, which creates a MNM file to use in the conversion process.



Browse...—Lets you pick a custom mapping file to use in the conversion process.

CSM Output Folder—Displays the folder path where the CSM file is generated.

Browse...—Lets you pick a folder in which 3ds Max saves the generated CSM file.

Convert TRC To CSM—Lets you first pick the TRC file to convert. The resulting CSM file is named after the original TRC file.

Note: If Use TRC Filename is turned off, a second dialog opens so you can pick an existing CSM file to overwrite or create a new one.

Options group

Use MNM Filter File—When on, the MNM file is used in the conversion process. Default=on.

Note: Turn off if the TRC marker names already follow the CSM standard naming convention.

Batch File Conversion—Enables you to pick multiple TRC files to convert.

Tip: Use this option to save time when you need to convert multiple TRC files.

Use TRC Filename—When off, you can customize the resulting CSM filename. When on, the CSM file is named after the TRC file. Available only when Batch File Conversion is off. Default=on.

Translation Offset group

Enable—Turn on to add an offset to the resulting data's position.

TRC Space/3ds max Space—Choose between an offset space relative to the TRC data's root or to your scene.

XYZ—Sets the offset value for each position axis.

Camera Tracker Utility

Utilities panel > Utilities rollout > More button > Utilities dialog > Camera Tracker

The Camera Tracker utility animates the movement of a camera inside 3ds Max to match the movement of a real camera that was used to shoot a background movie.

For detailed Reference information about each Camera Tracker Utility rollout see:

Camera Tracker: Movie Rollout (page 2–670)

Camera Tracker: Movie Window (page 2–671)

Camera Tracker: Motion Trackers Rollout (page 2–673)

Camera Tracker: Movie Stepper Rollout (page 2–676)

Camera Tracker: Batch Track Rollout (page 2–678)

Camera Tracker: Error Thresholds Rollout (page 2–677)

Camera Tracker: Position Data Rollout (page 2–679)

Camera Tracker: Match Move Rollout (page 2–680)

Camera Tracker: Move Smoothing Rollout (page 2–682)

Camera Tracker: Object Pinning Rollout (page 2–683)

For additional information on using the Camera Tracker Utility see:

Requirements for Camera Tracking (page 2–669)

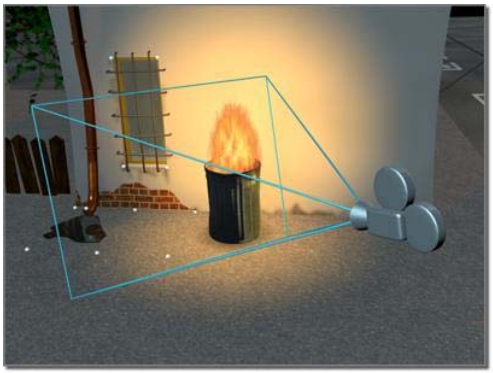
Camera Tracker: Troubleshooting (page 2–685)

Procedure

To generate a camera match-move:

1. Open the working scene in 3ds Max.

The scene should have a Free camera to be match-moved, as the tracker does not create one. The scene also should have a set of Point or CamPoint helper objects positioned in 3D to correspond to the tracking features. Optionally install the movie as an environment map-based background image in the match camera viewport. This is needed if you want to render a composite using the match-move or checking match accuracy. This doesn't automatically display the background image in the viewport. Use Views > Background Image to select the background movie and display it in the viewport.



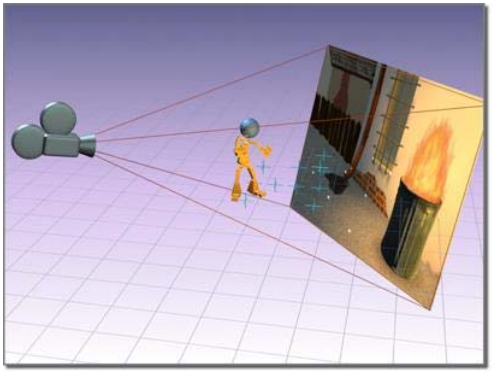
Real-world camera films a scene (the white dots will be tracking points).

2. Open the movie file in the Camera Tracker utility and create a set of feature tracking gizmos for each of the tracking features in the scene.



Resulting footage to be tracked and used as a background.

3. Position the Feature Selection box and Motion Search box for each gizmo so that they're centered on the features and have motion search bounds large enough to accommodate the biggest frame-to-frame move of the features throughout the frames that will be tracked.
4. Associate each tracker with its corresponding scene point object.



The associated scene point objects are set up in 3D space based on real-world dimensions.

5. (Optional) Use the Movie Stepper rollout to set the start and stop frames for each tracker if it's out-of-view for any of the frames that will be matched. These specify the frame range during which the tracker is visible in the scene and will be tracked as part of the matching process. This mechanism allows you to match a

move in which the view passes through a field of features with only some of them (at least six) visible at any time.

6. (Optional) Set up manual keyframes for each tracker at frames in which the feature radically changes motion or shape or is briefly occluded and so might be difficult for the computer to track. If you want, you can do this after a tracking attempt indicates where tracking errors occurred.
7. Perform the feature track using the Batch Track rollout. This is often an iterative process: correcting for tracking errors by tuning start and stop frames or manually repositioning the gizmo and motion search boxes at error frames. When you reposition a gizmo at some frame, you establish a new target feature image and subsequent frames up to the next keyframe will be retracked. You can use the error detecting features in the tracker to step through possible tracking errors. When complete, this process builds a table of 2D motion positions for each feature. You can save this to disk using the Save button on the Movie rollout.
8. Choose the camera that will be matched in the Match-Move rollout, select which camera parameters you want to estimate, set the movie and scene animation frame ranges and perform the match. This generates a keyframed animation of the selected camera parameters.



After camera tracking, scene geometry (the character) matches the filmed background.

9. Check the match for obvious errors and review the tracker gizmo positions at these frames. You can manually adjust gizmos at these frames and the matcher interactively recomputes the camera position.
10. (Optional) Apply smoothing to selected camera parameters and recompute a compensating match for the other parameters.

The Camera Tracker also has the ability to animate 3ds Max geometry to follow or match the video by following the movement of a tracker in a 2D plane. For this type of animation scene measurements aren't required. Use the Object Pinning rollout of the Camera Tracker to create this type of tracking.

Requirements for Camera Tracking

To use the camera tracker, you need the following:

- Movie footage in one of the file formats that 3ds Max supports, such as AVI or MOV. If you have sequential still images, you can use an IFL (Image File List). The file selector can automatically generate an IFL file by selecting the file name and turning on the sequence button. Or use the *IFL Manager Utility (page 3–619)* to create an IFL file.

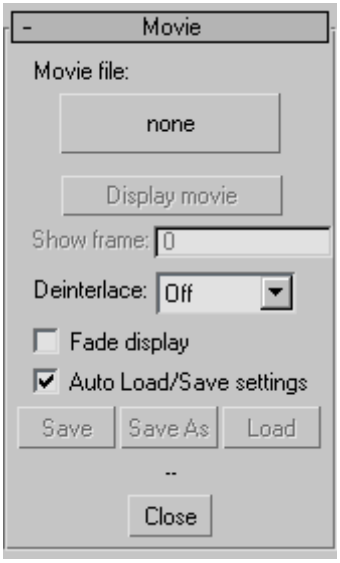
- A set of trackable features in the movie, such as wall markers, balls, corner points and so on, for which position measurements are known. The tracker will track most small features with good contrast to their surroundings, have a well-defined position in space, and are unique within their frame-to-frame areas of movement. There must be six or more tracking features in the scene at any one time, at least two of which must be non-coplanar with the others, for example four on the ground and two elevated or three on one wall and three on another. For maximum match accuracy, you should use as many features as you can and select the features that are distributed as widely as possible over the scene. The features don't have to be in any particular planes.
- A 3ds Max scene file in which a set of point objects have been created and positioned in 3D corresponding to all the tracking features in the scene. The objects should be point or CamPoint helper objects. The tracker will match-move a selected Free camera in the scene, so you need to have at least one Free camera. You can place and orient it anywhere in the scene.

Camera Tracker: Movie Rollout

Utilities panel > Utilities rollout > More button > Utilities dialog > Camera Tracker > Movie rollout

Opens the movie for tracking. Also provides controls for the display of the movie, including fade display to improve the visibility of the trackers. Provides controls for loading and saving MOT files in which tracker information is stored.

Interface



Movie file—Selects and opens the movie that will be tracked. You can open any image file format that 3ds Max supports. If you use sequential still-image files, you will use an *.ifl* (image file list) file. You can create the *.ifl* file using the IFL Manager or using any bitmap selector dialog by choosing Sequence. When you open the movie file, it's displayed in a Movie window.

Display Movie—Reopens a closed or minimized Movie window. You use the Movie window to setup and adjust feature tracking gizmos and to browse through the movie. For more details on working with the Movie window, see *Motion Tracker rollout (page 2–673)*.

Show frame—Allows you to step through the movie displayed in the Movie window. Alternatively, use the Movie Stepper rollout, which provides additional browsing controls.

Deinterlace—Causes the tracker to apply a video deinterlacing filter to the movie frames. If the movie comes from a video source and the features you're tracking show marked interlacing, you should deinterlace the video to improve tracking

accuracy. If the movie was digitized from film, you should not perform deinterlacing because the results will be less accurate. The deinterlacing is performed temporarily on loaded frames and doesn't affect the original movie file.

Off—Disables deinterlacing. When no filtering is required click this.

Odd—Interpolates using the odd lines.

Even—Interpolates using the even lines.

Fade Display—Fades the movie in a Movie window by 50 percent. Use this to see the tracker gizmos more clearly.

Auto Load/ Save—Saves the state of a tracking setup and any position data you've produced to a special file associated with each movie. Selecting Auto Load/Save Settings causes the tracker to keep this settings file up-to-date automatically as you work in the tracker. You can force a save at any time by clicking the Save button. Whenever you open a movie file in the Camera Tracker with Auto Load/Save on, the tracker reloads the state of the tracking session from this settings file.

With Auto Load/Save Settings on, the tracker constructs the setting file automatically and places it in the same folder as the movie, named with the same prefix as the movie but with file type MOT. If you move the movie file to a new folder, move the MOT file with it so the tracker can find it. You can delete this file if you want to clear the settings or if the file becomes corrupted for any reason.

When AutoLoad/Save is turned off, you can open a different movie file either with an empty tracker state or into the existing tracker state or save the current tracker state into a new file. The name of the current tracker file is displayed at the bottom of the Movie rollout.

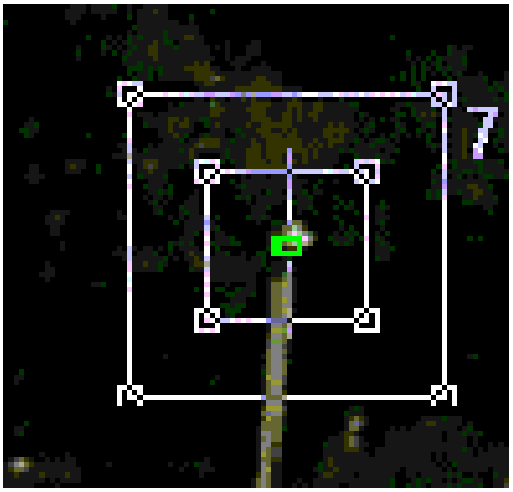
Save—Saves the state of the tracking setup, and any position data produced to the current MOT file.

Save As—Saves the current setup into a new MOT file. This file becomes the destination for subsequent Save operations.

Load—Load the tracker setup and position data from another file.

Camera Tracker: Movie Window

The tracker gizmo displayed in the Movie window consists of two enclosed boxes, a central pair of cross hairs, and a tracker number. The inner box surrounds the feature that will be tracked and is called the *feature bounds box*. The outer box defines the frame-to-frame search and is called the *motion search bounds box*. You should place the central cross-hair over the feature as close as possible to the point corresponding to the placement of its scene point object. It's this center coordinate that's used in the camera match-move correlation between image feature and scene Point coordinates.



You should make the feature bounds box large enough to enclose the feature and some of the surrounding images enough to give the feature a contrasting background of several pixels.

The motion search bounds defines the are in which the feature will be searched for from frame-to-frame. This bounds box moves with the feature box, so the area is relative to the current feature at each frame. It's important to estimate this search area well. If it's too large, the matching process will be unnecessarily slow and there's more chance of other features in the search area making the search ambiguous; if it's too small, tracking errors will occur. It might be useful to review the movie that will tracked on a real-time playback device like a video deck and estimate feature moves beforehand. If you set the Max Move/Frame spinner to the maximum move before creating the tracker, its motion search bounds box will be set to accommodate this move for you. You can also adjust this box at any time in the Movie window.

Note: It's possible to change the search bounds box for different frame ranges in the movie, so you can optimize your search capabilities. See *Keyframes and Position Data (page 2–679)*.

Working with the Movie Window

Displays the current frame of the movie being tracked and any active tracker gizmos for that frame. You can open and close the window using the Display Movie check box in the Movie rollout or minimize it using the standard window title bar controls.

When you first open a movie, 3ds Max sizes the window so it will fit on the screen and this might result in an automatic zoom out if the image is too big.

You can resize the window by dragging its sides or corners. You can zoom around in and out of the window using the tracker gizmo controls.

Working with the Tracker Gizmos

There are several ways to work with gizmos in the Movie window.

- **Selecting:** You need to select a gizmo to work on it. You do this by clicking anywhere inside its bounds in the Movie window or by clicking its entry in the tracker list box in the Motion trackers rollout. The selected gizmo displays positioning handles at the corners of both its bounds boxes. You can also select a gizmo by typing its number on the keyboard. A selected gizmo has click-selection priority in the Movie window. If several gizmos overlap, they make it hard to select and adjust the gizmo you want. You can select it first in the Motion Trackers rollout list; it remains selected for adjustment when you click it in the Movie window.
- **Zooming:** You can zoom in and out on the selected gizmo by pressing the I and O keys, for "in" and "out," respectively. Each time you press I or O, the window zooms in or out by a factor of 2 and centers the selected gizmo in the window. Pressing R resets the zoom factor to 100 percent whether you have a gizmo selected or not. The Movie window has to be the active window for any keyboard commands to work. You can make it active by clicking it.
- **Tabbing:** You can move among successively numbered gizmos by pressing the Tab key. If the gizmo you're searching for is out of view, because you're zoomed in, the image will be repositioned in the window to bring the newly selected gizmo into view. Tabbing while you're zoomed in provides a quick way to step through the gizmos to check their fine positioning.
- **Positioning:** You can reposition a gizmo by dragging anywhere within its bounding boxes, providing you don't click one of the eight box-corner handles. You can nudge a gizmo in single increments by pressing one of the arrow keys on the keyboard. The minimum increment that a gizmo will move depends on the zoom factor and the subpixel-tracking level. The increment is the smaller of either a single screen pixel in the current window or the subpixel

increment selected. So, if you're zoomed out two times and the subpixel increment is 1/16 of a pixel, the nudge increment is four image pixels, due to the zoom. However, if you're zoomed in five times, the nudge increment is 1/16 of an image pixel, due to the subpixel level.

- **Bounds:** You adjust the bounding boxes by pressing and dragging the handles of the box you want to adjust in the selected gizmo. Dragging the handles of the inner feature bounds box adjusts the opposite corner symmetrically to keep the gizmo center fixed.

Typically when setting up tracker gizmos you should create and position them all roughly in a zoomed-out view, then zoom in on one of them, fine-tune its position and bounds, then tab to the next gizmo and repeat.

Movie Window Keyboard Shortcuts

You can use the following keyboard shortcuts in the Movie window:

- I, i: Zooms in on a selected gizmo.
- O, o: Zooms out on a selected gizmo.
- R, r: Resets zoom level to 100 percent.
- Tab: Selects next numbered gizmo, cycles back to first gizmo after the last.
- **[digit]:** Selects the numbered gizmo, only works for gizmos 1 to 10 (0=10).
- **Arrow Keys:** Nudges the gizmo in the direction of the arrow.

Camera Tracker: Motion Trackers Rollout

Utilities panel > Utilities rollout > More button > Utilities dialog > Camera Tracker > Motion Trackers rollout

The Camera Tracker utility's Motion Trackers rollout sets up and controls the image feature trackers. All the trackers created are listed here: showing enable check mark, tracker number, associated scene point object name, and active frame range. To select a tracker, click in the list. Provides controls for subpixel tracking as well.

Interface



On—When turned on, the feature will be tracked in the image when the movie is stepped through in the Movie Stepper, or when a full track is performed. An X appears in the tracker list when turned on.

Off—When turned off, the tracker will sit at the closest known position prior to the current frame. Because feature tracking is computationally intensive, sometimes it's useful to disable the trackers you're not working on to speed up movie stepping.

All—Enables all the trackers.

None—Disables all the trackers.

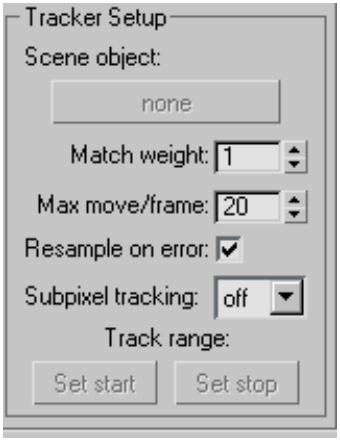
New Tracker—Creates a new movie feature tracker. This places a feature tracker gizmo in the center of the Movie window and a new entry in the tracker list. The new tracker takes on the parameters currently displayed under Tracker Setup. You can change these parameters after you create the tracker, but it's often useful to set the Max Move Per Frame parameter prior to creating the new

tracker as this automatically establishes an initial search bound for the tracker gizmo.

Delete Tracker—Deletes the selected tracker from the list and the Movie window.

Tracker Setup group

Sets up the parameters for the currently selected tracker. Any tracker you create will take its initial parameters from the current values in this group.



Scene object—Associates a tracker with a scene point object that corresponds to the image feature in the Movie window. Pressing this button puts 3ds Max into standard object picking mode. When you've selected the object, the entry in the tracker list and the text in the button change to the name of the selected object. You can change the connected object at any time by pressing this button.

It's recommended that you use point or CamPoint helper objects, although you can use any 3ds Max object. The tracker uses the pivot point of the selected object as the correlation reference point during the match computation.

When you set up a tracking project, you can make the tracker save the setup in a *.mot* settings file. The association between tracker and scene point object that you set up here is saved as a scene object name. If you reload the movie in the tracker, the object association is reestablished using the object's name. This means you must be careful to name the point objects uniquely, otherwise you might establish incorrect associations. Further, you must load the working scene in 3ds Max before opening the movie in the camera tracker utility, so the named objects can be found. If you happen to open the movie first, the object names in the Tracker list will all read "<none>". Simply reopen the movie in the tracker and it will establish the connections correctly.

Match weight—Sets the camera matching priority separately for each tracker. Use this to improve the match accuracy for features close to where a computer graphics generated element will be inserted. It helps to minimize sliding and jitter in that part of the scene.

The match-move algorithm works by adjusting the camera to minimize the error between the feature objects seen through the camera and their associated feature points in the movie plate. With all weights set to 1, the algorithm randomly distributes the error between all the object projections and their tracking features. If you bump the Match Weight of a tracker, it redistributes the error in favor of that tracker, reducing the distance between that tracker's feature and object projection, while possibly increasing the error in other trackers. If you bump weights on two or more features, the locking may progressively lessen, since the algorithm distributes the improvements between the high-weight trackers. Use weights of two to six to get a good lock.

Max Move/Frame—Sets the maximum number of pixels that the tracked feature moves from frame-to-frame over the entire tracking range and establishes an initial motion search bounds box for newly created trackers. You can also set this dynamically by adjusting the motion search bounds box directly in the Movie window.

Resample on error—You can make the tracker identify possible tracking errors and you can use this during a manual tracking review and during actual tracking to correct errors. Typical sources of error include frame-to-frame jumps greater than the search bounds, image contrast changes, or a feature changing shape over a sequence of frames, such as a corner might if the camera moves by it. The error detector is controlled by measurement thresholds set in the Error Thresholds rollout.

When you place a gizmo on some frame, a keyframe is created and a new target feature image is sampled at that frame for searching in subsequent frames. Turning on Resample On Error causes the tracker to detect errors during tracking. If one is found, it backs up to the last good frame and places a new keyframe there causing a new target feature image to be sampled at that frame and continues tracking from that frame.

Subpixel Tracking—Usually, an image feature won't move an integral number of pixels from frame-to-frame. With pixel-level tracking, however, the nearest estimate for a feature center is always at a pixel boundary. This can lead to camera match errors. With subpixel tracking you can zoom in on a feature and position its gizmo within a pixel. This makes the tracker attempt to track at that resolution by up-sampling the target image and search bounds, thereby improving match accuracy.

You can turn on subpixel tracking selectively for each tracker by choosing a level in the Subpixel Tracking drop-down menu on the Motion Trackers rollout. This defines the resolution at which feature tracking is performed in the tracker up to 1/32 of a pixel.

Two things happen when you select a subpixel level:

- The grid over which you can position a tracker gizmo in the image window is adjusted to the new subpixel level. You can position the gizmo can be positioned inside pixels in zoomed views. You can zoom in on a tracker gizmo by selecting it and pressing the I or O key to zoom in or out, respectively.
- Feature searching is performed on internally scaled-up versions of the target image feature and portions of the search window. These versions are scaled up in inverse proportion to the subpixel level using bicubic interpolation, which allows the feature tracker to search at this scaled pixel level.

When manually positioning gizmos at keyframes, it's crucial that you zoom in enough so you can place the gizmo in the center of the feature to the degree of accuracy allowed by the selected subpixel level. If you don't do this, the tracker won't be able to track the correct center at the selected resolution.

Tip: Even though the feature is first tracked to the nearest pixel at whole pixel increments and then tracked at the subpixel level within that one pixel range, the tracking time can increase significantly in proportion to the inverse square of the subpixel level. For this reason it's suggested that you set the subpixel level to as low as possible. Usually 1/8 of a pixel is adequate for features well distributed around the scene. Use greater levels if the features are fairly close together, or move very slowly within the frame.

You can also reduce feature tracking times by making the inner feature bounds box in the gizmo as tight as you can while still retaining a couple of pixels worth of surrounding contrast. Search times are proportional to the feature bounds box size.

Track Range

Use these when a feature moves off-screen or is not visible in the Movie window.

Set Start—Sets the start of the active frame range for a tracker.

Set Stop—Sets the end of the active frame range for a tracker.

You use the Set Start and Set Stop buttons in the Setup box together with the Movie Stepper rollout to set the active range. Pressing the buttons will set the start or stop point at the frame currently shown in the Movie Stepper.

Tracker Keyframes and Position Data

Any time you manually position a tracker gizmo, such as at the first frame or at some tracking error, you create a special keyframe position. The tracker never attempts to reposition these keyframes, but uses them as start points for the frames that follow them up to the next keyframe. At keyframes, you can adjust the position and the feature or search bounds.

During tracking, the selected feature at a keyframe is sampled and that feature becomes the search target for subsequent frames. Using any adjusted search bounds up until the next keyframe, provides the start point, target feature, and search bounds box for the frames that follow.

The position data for each frame in each gizmo can be one of three possibilities:

- A keyframe defining a new start position, target feature image, and search bounds box.
- A tracked position.
- Unknown.

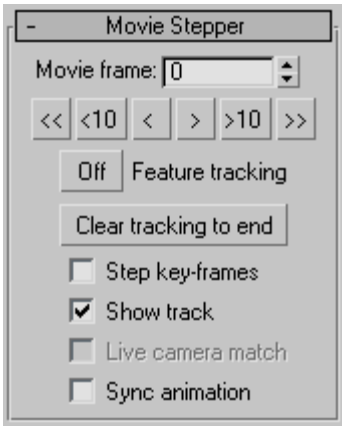
All frame positions except the first are initially unknown and will be named as tracking proceeds or gizmos are positioned. When you manually set a keyframe, any previously tracked positions following that keyframe up until the next keyframe are deleted and so become unknown. The Complete Tracking button in the Batch Track rollout searches for unknown positions in the currently enabled trackers and tracks those frames. The Check Status button in the same rollout checks for any currently unknown positions as well as tracking errors. See *Batch Track Rollout (page 2–678)*.

Camera Tracker: Movie Stepper Rollout

Utilities panel > Utilities rollout > More button > Utilities dialog > Camera Tracker > Movie Stepper rollout

The Camera Tracker utility's Movie Stepper rollout browses through the movie during the tracking phase either to set start and stop frames, to review tracking results or to manually perform a track over a sequence of frames. The Movie Stepper rollout consists of a frame counter, a set of stepper buttons, and some option check boxes.

Interface



Movie Frame Counter—Advances the Movie window to the designated frame.

VCR controls—Navigates through the movie in the Movie window:

<<—Steps to the start frame of the selected tracker.

<10—Steps back 10 frames.

<—Steps back one frame or keyframe if Step Keyframes is turned on.

>—Steps forward one frame or keyframe if Step Keyframes is turned on.

>10—Steps forward 10 frames, either directly or one frame at a time if the Master Track Enable check box is turned on.

>>—Steps forward to the end frame of the selected tracker, either directly or one frame at a time if the Master Track Enable check box is turned on.

Feature Tracking—If turned on, any forward steps will cause frame-by-frame feature tracking to be performed for all enabled trackers with unknown position data at the frames you previously stepped through. You can use this to perform a tracking check while testing bounds settings or to review tracked positions one frame at a time. The tracker also moves the gizmos to previously tracked positions, so you can review tracking results a few frames at a time. If you want, you can force a retrack by nudging a keyframe gizmo back-and-forth which causes the gizmos to drop all subsequent tracked positions up to the next keyframe.

Note: Make sure this button is turned off if you just want to step through the movie without tracking, such as when browsing through the movie to find and place Start and Stop frames for trackers that come on screen or go off screen during the tracking range.

Clear Tracking to End—Removes all tracked positions and keyframes from the current stepper frame onward for the currently selected tracker. This is often helpful if you have too many errors and error correcting keyframes at the end of a track and you want to clear them and track to the end again.

Step Keyframes—Modifies the operations of the single-step buttons in the Movie Stepper. If turned on, the single step buttons step immediately to the next or previous keyframes, allowing you to browse through them quickly. In this mode, feature tracking is disabled.

Show Track—Displays a visible tracking line in the Movie window for the currently selected tracker. This button shows where the tracker has currently tracked. A red line is displayed for tracked segments, a white dashed line for segments still to be tracked. White dots are centered on each frame position with a green box around each frame position at which you or the tracker create keyframes.

Live Camera Match—Enabled after you set up the match parameters and select a match camera in the Match rollout. If turned on, the match camera position is recomputed and adjusted as you move gizmos in the Movie window. This can be useful for making corrections to bad match frames by adjusting the trackers or to check the effect of tracker placement on the match computations. This is best done with the nudge arrows. Turn on the 3ds Max Auto Key button to record adjusted match camera positions.

Sync Animation—Controls whether the scene time slider is advanced in step with the movie frames during stepping. This is sometimes useful if you have the movie set up as an animated background image in the match camera window. Check camera positioning at the same time you check feature positioning. This can slow stepping substantially, so it's turned off by default.

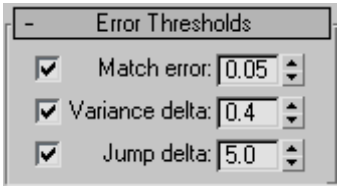
Camera Tracker: Error Thresholds Rollout

Utilities panel > Utilities rollout > More button > Utilities dialog > Camera Tracker > Error Thresholds rollout

The Camera Tracker utility's Error Thresholds rollout attempts to identify feature tracking errors by taking measurements and comparing them to the thresholds that you set up in this rollout. You

can use this detection can be used both during a manual tracking review (see *Batch Track Rollout (page 2–678)*) or during actual tracking to correct errors using the Resample On Error control in the *Motion Trackers setup rollout (page 2–673)*.

Interface



Match Error—A sum-of-differences-squared measure in RGB space of the best match in the current frame to the feature target from the previous keyframe. The error shown is a percentage of the maximum possible image difference. Good matches are usually below 0.05 percent.

Variance Delta—The difference in RGB color variance between the feature target image and the current best match. This measures whether the target and match have similar color ranges and can compensate for mistakes made by the sum-of-differences measure.

Jump Delta—The tracker computes a moving average for the preceding five frames and measures the difference between that average and the jump in the current frame. If this difference exceeds the current Jump delta threshold, a potential error is flagged. If the threshold is set to five, any jump in the current frame that's five pixels more than the average jump in the last five frames is flagged. This picks up sudden accelerations rather than just large jumps.

You can adjust the thresholds as needed and enable or disabled individual measures in the detector using the check boxes to the left of each measure.

Camera Tracker: Batch Track Rollout

Utilities panel > Utilities rollout > More button > Utilities dialog > Camera Tracker > Batch Track rollout

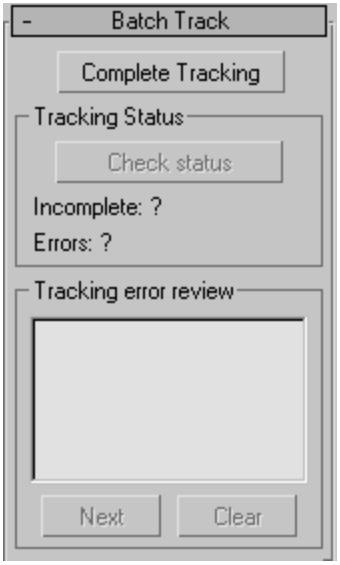
The Camera Tracker utility's Batch Track rollout performs feature hands-off tracking once the trackers have been set up. You can set up a set of trackers, start a batch tracking of them all, and then use the other tools in this panel to check for and review detected tracking errors. This panel also checks for tracking completion.

Procedure

To review and correct tracking errors:

1. Select an error in the list.
2. Browse through the errors in sequence by repeatedly clicking Next, under the list.
3. When you've selected an error, the associated tracker gizmo is selected and brought into view and the movie is set to the frame at which the potential error was detected.
4. If there is indeed an error, adjust the gizmo in the Movie window to correct it, and place a new keyframe.
5. If inspection shows that there isn't an error, the thresholds may be set too finely, so try adjusting the associated threshold and running Check Status again. This is often a good way to calibrate the thresholds for your scene, but be careful not to increase them too much and miss a real error.
6. To clear an entry out of the list, click the Clear button. This is basically a cosmetic action that removes items from the list that you've either fixed or determined are not errors while working through the potential errors. If you perform a Check Status again, some of the errors may show up as well.

Interface



Complete Tracking—Searches for unknown positions in enabled trackers and undertakes tracking for those frames. It performs this optimally, and only tracks the frames for which enabled trackers have no position data.

Tracking Status group

Check Status—Scans for tracking errors and untracked frames. It displays two lines of information underneath it in the Status box:

Incomplete—Lists the tracker numbers with untracked frames in their active frame ranges.

Errors—Lists the tracker numbers containing one or more tracking errors as determined by the current error threshold settings.

Tracking Error Review group

Displays all the potential errors in a list. For each error the list shows a tracker number, error frame and details about the error. The details include a code identifying the thresholds and the error measure.

Error Codes

Me—Match error threshold exceeded.

Vd—Color variance delta threshold exceeded.

Jd—Jump delta threshold exceeded.

In each case, the number following the code is the actual error measure. You can see how the error measure compares to the current thresholds in the Error Thresholds rollout. If you adjust the thresholds in this rollout and then rerun the Check Status again, the list is refilled with those points in the tracking that exceed the newly adjusted thresholds.

Next—Moves to the next error in the list.

Clear—Clears an entry out of the list.

Camera Tracker: Position Data Rollout

Utilities panel > Utilities rollout > More button > Utilities dialog > Camera Tracker > Position Data rollout

Interface



Apply to group

Contains three radio buttons that control which trackers will be operated on by the action buttons below it.

Selected Tracker—Operates on the currently selected tracker.

Enabled Trackers—Operates on all the enabled trackers in the Motion Trackers rollout.

All—Operates on all the trackers.

Clear to End—Clears the position data in the trackers specified in the Apply To group from the current Movie Stepper frame onward. This is useful for clearing out and redoing a portion of the tracking.

Clear All—Clears all the position data in the trackers specified in the Apply To group.

Show Data—Opens a text window that displays the position data for the trackers specified in the Apply To box. The display includes some title lines showing the tracker number and associated tracking object followed by comma separated columns of numbers for the position data one row per frame. All rows contain at least three columns: frame number, X position, Y position. Keyframe frames also contain another eight columns: feature bounds left, top, right, and bottom, and Search bounds left, top, right, and bottom. The coordinates in decimal pixels show any subpixel tracking as decimal fractions.

The text window is a standard MAXScript source editor and you can copy, edit or save the contents to file.

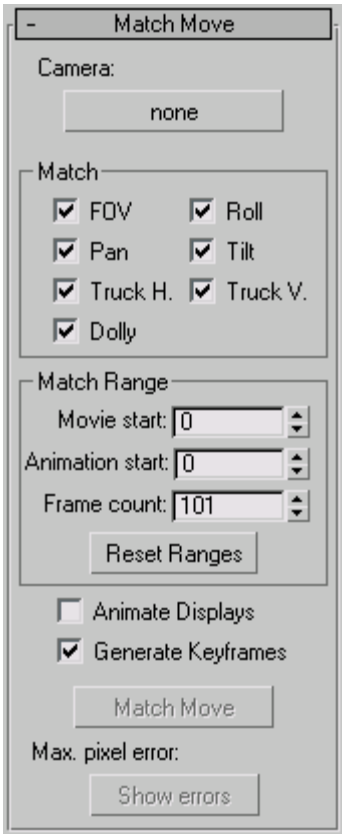
Export—Creates a text file containing position data for the trackers specified in the Apply To group. Pressing this button displays a standard Save File dialog asking for the location and name of the file you're creating. The file type defaults to *.csv*, the comma-separated value format readable by Excel and other data processing applications. The format of the data is exactly as for the Show Data command described previously.

Camera Tracker: Match Move Rollout

Utilities panel > Utilities rollout > More button > Utilities dialog > Camera Tracker > Match Move rollout

The Camera Tracker utility's Match Move rollout automatically generates a camera animation by correlating the 2D position data of the movie features gathered during the tracking phase with the associated 3D feature points in the 3ds Max scene.

Interface



Camera—Enters an object-picking mode for you to select the camera that you will match-move. You can either select the camera directly in the scene or

press H to open the *Select Objects dialog (page 1–78)*. Once selected, the camera name is placed in the picker button, and other buttons and controls are enabled in the utility. The selected camera is remembered in the MOT settings file and restored on reopening the movie in the camera tracker. As with the tracker point scene objects, the camera name is stored and used when reloading, so make sure it's unique. The camera selected must be a Free camera. It can be at any position and orientation in the scene.

Match group

Lets you choose which camera parameters will be estimated during the match-move. Those parameters turned off will be derived from the current Camera animation or position, so it's possible to pre-position or pre-animate the match camera with known moves or parameter settings and have the matcher estimate just the missing parameters. The accuracy of the estimation process increases substantially as the number of parameters that need estimation decreases.

FOV—Camera FOV (field-of-view).

Pan—Rotation about the local camera Z axis.

Tilt—Rotation about the local camera X axis.

Roll—Rotation about the local camera Y axis.

Dolly—Movement along the local camera Y axis.

Truck-H—Movement along the local camera X axis.

Truck-V—Movement along the local camera Z axis.

Match Range group

Let's you set the range of frames that will be matched to set up the relationship between the movie frames and the 3ds Max scene animation frames. By default, the start frames are both set to zero and the frame count is set to the smaller of the number of frames in the movie or the scene animation. Setting a different Movie Start or Animation Start frame lets you position the movie frame range relative to the animation frame range. The position data for a match will be taken from the tracker frames starting at the Movie Start spinner value and the camera will be matched and animated starting from those frames.

Movie Start—Sets the start of the range to be matched in the movie.

Animation Start—Sets the start of the range in the animation. Use this if the animation range is different from the movie range.

Frame Count—Sets the number of frames to be matched.

Reset Ranges—Reloads the maximum frame number and count ranges into the match range spinners from the current movie and 3ds Max scene settings. This is useful if you change the animation frame range in the scene while the tracker is open.

Animate Displays—Controls whether the 3ds Max viewport appears and if the Movie window updates frame-by-frame to show the match results. This display update can be very time-consuming, so it's often best to leave this turned off and review the match results by looking at the Camera function curves in Track View or by stepping through the 3ds Max animation a frame at a time after the match.

Generate Keyframes—Generates keyframes for the estimated parameters on each frame.

Match-Move—Starts the match process.

Max. Pixel Error—Displays the running maximum tracking error and frame. The error is shown in pixel fractions, and gives the largest distance between any feature and the viewed position of its associated tracking object over all the frames matched. You can use this to give a general estimate of how good the match is and the frame to investigate if errors are problematic. The pixel

error display also interactively updates as you nudge tracker gizmos in Live Camera Match mode, showing the current maximum match distance for the camera; see *Movie Stepper Rollout (page 2–676)*.

Note that this error might be exaggerated when you employ match point weighting, as this tends to distribute the errors toward the nonweight points.

Show Errors—Opens a text window containing a list of the maximum errors for each of the matched frames, so you can look over the whole set of match frames and identify which frames to check.

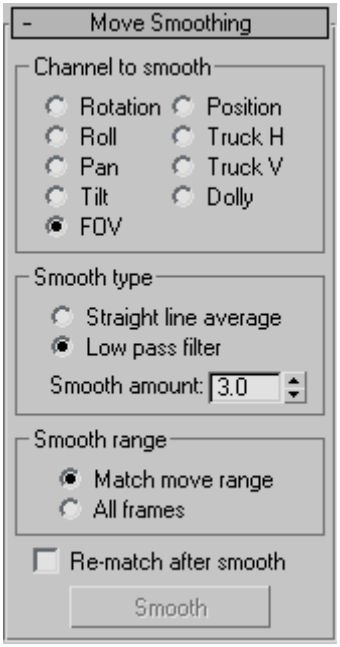
Camera Tracker: Move Smoothing Rollout

Utilities panel > Utilities rollout > More button > Utilities dialog > Camera Tracker > Move Smoothing rollout

Because the match-move algorithm is not always exact and is highly susceptible to field measurement error or feature tracking errors, and because the match for each frame is computed independently, the results can be "noisy", particularly if field of view is estimated. In match-move situations where the real camera performs smooth moves, it might be useful to filter the match results to eliminate this noise.

You use the Move Smoothing rollout to control and perform these smoothing operations after a match has been created.

Interface



Channel To Smooth group

Lets you choose which parameter or set of parameters to smooth. You can choose to smooth the entire Rotation or all the Position tracks at once, or you can smooth individual parameters or axes.

Rotation—Pan, Tilt, and Roll.

Position—Dolly, Truck-H, and Truck-V.

FOV—Camera FOV (field-of-view).

Pan—Rotation about the local camera Z axis.

Tilt—Rotation about the local camera X axis.

Roll—Rotation about the local camera Y axis.

Dolly—Movement along the local camera Y axis.

Truck-H—Movement along the local camera X axis.

Truck-V—Movement along the local camera Z axis.

Smooth Type group

Straight line average—Causes Move smoothing to generate a straight line average for each of the selected parameters. This is useful when you know a parameter is fixed, but don't know its actual value. Examples might include field-of-view or vertical position. Use the camera matcher to initially estimate the parameter, then construct a straight line average fit and rematch the remaining parameters

Low pass filter—Causes Move smoothing to apply a low-pass filter to the selected parameters, eliminating high-frequency noise.

Smooth amount—Implements smoothing using a Gaussian low-pass filter whose width and variance are controlled by the Smooth Amount spinner. Increase this value to increase the smoothing.

Smooth Range group

Controls the range of frames over which the smoothing will be performed.

Match-Move Range—Applies the smoothing to the range of animation frames specified in the Match Move rollout. This allows you to break up a match-moving job into several frame ranges. Use this to isolate and experiment on a troublesome range of frames, without disturbing other satisfactory frames.

All Frames—Applies the smoothing to all the keyframes in the selected parameters.

Rematch After Smooth—Controls whether a compensating match-move is performed after the filtering. If this is turned on, the parameters that were smoothed are disabled in the Match rollout and another match-move is performed. Performing this rematch again is advised because it adjusts the other camera parameters to compensate for the smoothing and often results in a smoother estimation for the unsmoothed parameters.

Smooth—Performs the move smoothing based on the smooth type, range, and amount settings. If this button is disabled, no current match camera is assigned.

Camera Tracker: Object Pinning Rollout

Utilities panel > Utilities rollout > More button > Utilities dialog > Camera Tracker > Object Pinning rollout

Object Pinning lets you use the 2D image feature tracking data to directly animate objects in the 3ds Max scene. This is useful for replacing moving features in the scene with 3D objects when the movement is approximately planar. After the tracking data is created you can use this function to do such things as synchronizing the movement of 3D objects with feature movements in the background movie. If the 3D objects directly cover the background features, they will appear to replace them. For example, you could take the video of a juggler juggling balls and replace the balls with 3D balls of fire or synchronize the movement of an animal's head with a 3D jaw to create talking animal animations.

Procedure

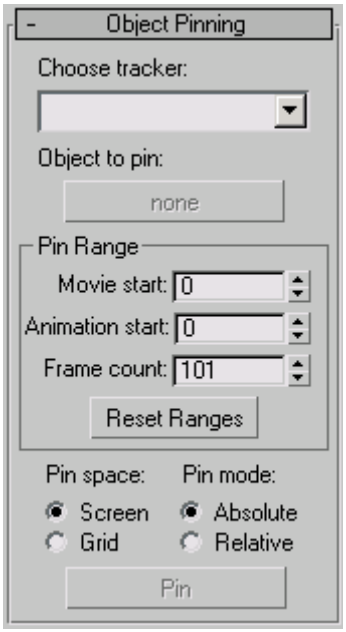
Example: To create a juggler juggling 3D balls of fire:

This assumes you already have a video of a juggler. Otherwise, find a juggler and a video camera.

1. Load a video of a juggler.
2. Create tracker gizmos to associate with the juggling balls.
3. Use the Movie Stepper to track the gizmos.
4. Go to Object Pinning and choose the tracker.
5. Under Object To Pin, associate an object with a tracker. In this case, create a 3D ball of fire using fire or an animated map to generate the fire.

6. Adjust the Pin Range, and then click the Pin button to move the object to follow the tracker.

Interface



Choose Tracker—Shows all the current trackers as set up in the Motion Trackers rollout. Use this to select the tracker whose feature movements you will use to control the pinned 3ds Max object. The feature must have been tracked already over the frames you want to use, with the Track and/or Movie Stepper rollouts. Note that the accuracy of the pinning is determined by the accuracy of the match, so subpixel tracking is highly recommended.

Object To Pin—Lets you choose the object in 3ds Max scene for pinning and animating. By default the tracking object associated with the selected tracker is preloaded into this button. You can use this picker if you wish to pin a different object.

Pin Range group

Movie Start—The movie frame from which to start using tracked feature positions.

Animation Start—The frame in the 3ds Max animation at which to start keyframing the move of the pinned object.

Frame Count—The number of frames to animate. There must be enough tracked positions in the selected tracker to cover this number of frames.

Reset Ranges—Reloads the maximum frame number and count ranges into the above spinners from the current movie and 3ds Max scene settings.

Pin Space group

Lets you select the plane of motion in which the keyframed pinning will occur.

Screen—Moves and animates the pinned object in the plane of the screen at its current depth in the scene. This is equivalent to dragging an object around in the Screen reference coordinate system in 3ds Max.

Grid—Moves and animates the pinned object in the plane of the currently active grid. This lets you set up an arbitrary plane of motion for the pinning by constructing a grid helper object in the desired place and making it the active grid. If you don't have a grid helper set up and active, the pinning will occur on the 3ds Max Home Grid.

Pin Mode Group

Lets you choose between two different pinning modes: absolute and relative.

Absolute—Places and moves the pinned object exactly over the tracked feature. If combined with grid space, the object is also placed and moved directly on the active grid plane.

Relative—Leaves the pinned object at its current position in space and moves it so that its projected position relative to the tracked feature remains the same throughout the pinning frame range. This makes it possible to animate an object that you don't want placed exactly over a feature, or to animate an object using several features in successive frame ranges. (For example, a computer generated jaw can track an ear for some frames, then a nose for others in a talking animal shot). If combined with grid space, the object is *not* moved onto the grid plane, but keeps its height above the grid and is moved parallel to it.

Pin—Performs the actual pinning, moves the object to follow the selected feature and places a keyframe at each frame in the pin frame range. As with the other keyframing operations in the tracker, you can undo a pin using the 3ds Max Undo function. You might also find it useful to perform keyframe reduction afterwards in Track View.

Camera Tracker: Troubleshooting

Feature Tracking Halts Immediately

If the feature tracker attempts to match the first couple of frames but halts immediately, check these steps:

1. One of the enabled error thresholds might be set too low. Perform a Check Status in the Batch Track rollout and look at the reports in the Tracking Error Review list. Also try disabling the thresholds in the Error Thresholds rollout one at a time and see if this fixes the problem. Adjust the problematic threshold as needed. You can also disable error threshold checking altogether by turning off the Resample On Error check box in each of the trackers.
2. Make sure the search bounds box for each tracker is big enough to accommodate the frame-to-frame motion of the feature. This needs to be at least big enough to enclose the moved feature and the surrounding contrast area you've selected in the feature bounds box.
3. Make sure the scene objects are associated correctly with their trackers. You can verify this by checking the entries in the list at the top of the Motion Trackers rollout. If there is an object associated, you will see its name in the list entry, if not you will see "<no object>". If you open a movie file for which you have previously set up trackers in the camera tracker before you open the associated 3ds Max scene, the objects won't be associated with their trackers. You should make sure the correct 3ds Max scene is open then close and reopen the movie file in the Camera Tracker so that it can now find the scene objects.
4. If you have inadvertently swapped the scene objects for one or more pairs of trackers, the match-move can't find a solution.

Repeated Match-Move Errors

If you get repeated camera match-move errors, here's a list of things to check:

1. Make sure there are at least six features actively being tracked at the error frame and that at least two of them are a good distance out the plane of the others.
2. Check the tracker gizmos in the Movie window around the frame on which the error occurs to ensure that no obvious feature tracking errors exist.
3. If you've disabled any parameters in the Match section of the Match-Move rollout prior to a match-move, the camera being matched must be already set correctly in the scene for those parameters. For example, if you disable roll and FOV, you must already have set the camera to the correct FOV and roll orientation, either directly in 3ds Max, or as a result of a previous match-move. The matching algorithm uses all

the "fixed" camera parameters to compute the ones being estimated and simply won't work if they're fixed at incorrect values.

4. You may know that some of the camera parameters don't vary during a match sequence, but you aren't sure of their values. One approach to this situation is to enable all parameters for an initial match and then apply a straight-line average filter to them in the Move Smoothing rollout. This automatically disables them in the Match section and with luck sets them at a good estimate of the fixed position.

Dynamics

Utilities panel > Utilities rollout > More button > Utilities dialog > Choose Dynamics.

The term *dynamics* refers to a system of controls that generate keys to produce animation that simulates real-world physics. For example, using standard keyframing techniques to animate a bouncing ball, you create keys that move the ball down to the floor, squash the ball, move the ball back up, and so on.

Using a dynamics system, you assign physical properties to the ball and the floor (such as friction or amount of bounce), specify which object will collide against another object (for example, the ball with the floor), place an effect (such as gravity), in the scene, and then calculate a solution over a range of frames. The result is a number of keys that produce an animation in which the ball drops due to the gravity effect, collides with the floor and reacts naturally based on its surface characteristics.

Note: As of 3ds Max 5, the preferred tool for dynamics simulation is **reactor**, available from the Utilities panel. The **reactor** plug-in lets you control and simulate complex physical scenes with ease. **reactor** supports fully integrated rigid and soft-body dynamics, cloth simulation, and fluid simulation.

Objects and Space Warps Used to Create Dynamics Simulations

A very basic simulation would involve creating a sphere and a gravity space warp, adding the sphere to a new simulation, assigning gravity as an effect on the sphere and solving the simulation. The result is that the sphere falls under the force of gravity.

You can use *dynamics objects (page 1–395)*, dynamics space warp deflectors (*PDynaFlect (page 2–81)*, *SDynaFlect (page 2–85)*, and *UDynaFlect (page 2–86)*), and space warp forces such as *Gravity (page 2–73)* and *Wind (page 2–75)* to add complexity to a simulation. You can use a particle stream from a particle emitter as a force on an object: the particles can collide with and be deflected by the object, as well as move it. You can attach special dynamics objects to objects, such as *Spring (page 1–400)* to simulate the effects of a spring.

You can combine effects: Wind, gravity, a spring object, collision, particle deflection and collision, as well as surface properties such as friction, can all work on an object in a simulation. For a full understanding of what dynamics can do, explore this topic and areas of 3ds Max that pertain to dynamics.

Dynamics-Specific Areas in 3ds Max

The Dynamics Utility (this topic)

The Dynamics Utility is the main control center for dynamics simulations. You specify which objects are used in the simulation, what their interactivity is with each other and with the effects in the scene. The simulation is then "solved," generating the keyframes.

The effects of collisions between objects depend on the velocity of the objects and their properties. For collision to work between two objects, each object must have the other object assigned for collision. For example, when bouncing a ball, both the floor and the ball are assigned collisions.

The Material Editor: The *Dynamics Properties rollout (page 2–1479)* in the Material Editor lets you assign the dynamics properties to the surface of an object, like friction and bounce. Using a multi/sub-object material, you can have different surface properties at the face level of any object.

Note: You can override material surface dynamic properties with controls on the *Edit Object dialog (page 2–696)* in the Dynamics utility.

Dynamics Objects, Particles and Dynamics, Dynamics Interface: Special objects like Spring and Damper, Space Warp forces like Gravity and Wind, as well as Space Warp deflectors like PDynaFlect can all influence a dynamics simulation. You must create these objects and space warps first in other areas of 3ds Max before you can use them in a dynamics simulation. See the following topics for details on their creation and use:

- Create panel > Geometry > *Dynamics Objects (page 1–395)*

  You can use the Spring and Damper objects for dynamics effects.

- Create panel > *Space Warps (page 2–55)* > Particles and Dynamics.

  Effects (forces) are special space-warp objects that emulate natural phenomena, such as wind or gravity. In a dynamics simulation, you must place gravity in the scene if you want the objects in your simulation to fall.

- Create panel > *Space Warps (page 2–55)* > Dynamics Interface.

  The Dynamics Interface space warps can cause particles to deflect from and affect an object in a simulation.

Combine Keyframed Objects in a Simulation: Keyframe an object and have it interact with other objects in a dynamics simulation by turning on the This Object is Unyielding check box in the Edit Object dialog for the keyframed object. Objects can bounce off of a keyframed sphere for example.

Dynamics Controller: When a Dynamics simulation is solved, a new list controller is created that holds both the generated dynamics keys and the original keys. This allows you to restore your original keys, if necessary. Undo is not supported by Dynamics.

For example, if a sphere is bouncing in a Dynamics simulation, and the sphere already contains position keys from a previous animation, the following tracks display in Track View:

```
Transform
     Position
          Dynamic Position Controller
          Old Position
     Rotation
          Dynamic Rotation Controller
          Old Rotation
```

Procedures

Example: To create a basic simulation of bouncing boxes:

1. Create a thin box in the Top viewport.

   Have the box almost the same height and width as the viewport. This will act as the ground plane.

2. In the Front viewport, create six small boxes above the "ground" box.

   Position some of them to collide with each other when they fall.

3. On the Create panel > Space Warps > Particles and Dynamics > Object Type rollout, click Gravity.
4. Drag in the Top viewport to create a Gravity gizmo.
5. On the Utilities panel, click Dynamics.
6. On the Dynamics rollout, click New.

   Dynamics00 appears in the Simulation Name field.
7. Click Edit Object List.

   The Edit Object List dialog displays.
8. Select all the boxes in the dialog and click the > button, then click OK.

   All the objects are moved to the Objects in Simulation list on the right side of the dialog.
9. On the Dynamics rollout, click Edit Object.

   The Edit Object dialog displays.
10. In the list under Object, choose Box01.
11. Turn on Dynamics Controls group > This Object is Unyielding.
12. Click Ok to close the Edit Object dialog.

    The "ground" box won't move when the other objects collide with it.
13. On the Dynamics rollout, in the Effects group, turn on Global Effects.
14. Click Assign Global Effects, select Gravity in the dialog and click > (right arrow), then click Ok.

    This assigns gravity to all objects in the simulation.
15. On the Dynamics rollout, in the Collisions group, turn on Global Collisions.
16. Click Assign Global Collisions, select all the boxes in the dialog and click > (right arrow), then click OK.

    Collisions are active for all the boxes.
17. Turn on Update Display with Solve, and then click Solve.

    In the viewports the objects fall and collide with each other and the ground plane.
18. Turn on the Auto Key button, move the time slider to frame 15, and then select and move the "ground" box upward along the Z axis.

    The ground will "move" during the simulation.
19. Click Solve.

    The keyframed "ground" box moves up and collides with the boxes. The ability of a keyframed object to be part of a simulation is one of the useful features in 3ds Max. You could use this capability to strike a ball with a bat, for example.

    For further experimentation, create a spring object in Create panel > Geometry > Dynamics Objects and attach the ends of the spring object to two of the boxes and then solve the simulation. The spring will stretch and follow the bouncing boxes.

You can use the space warps in Create panel > Space Warps > Dynamics Interface, such as SDynaFlect (Spherical Dynamics Deflector) to cause a particle stream to "push" an object in a dynamics simulation.

To remove the dynamics tracks and restore the original animation tracks in Track View:

When you solve a dynamics simulation, 3ds Max creates a list controller that holds both the generated dynamics keys and the original keys. This lets you easily restore the original keys. Undo is not supported by Dynamics.

To layer simulations, reverse this method. In other words, after you've solved the first simulation, copy its controllers to the old tracks, and then set up the next level of the simulation. The new simulation will base its actions on the previous one instead of overwriting it, as it normally does. You

can repeat this as many times as you like to layer simulations *ad infinitum*.

Since Undo is not supported by Dynamics, you can also use Hold and Fetch in its place.

1. In Track View, open the Position track for an object animated by a dynamics simulation.
2. Select the Old Position track, and click Copy Controller on the Track View toolbar (the second button from the left).
3. Select the Position track (the parent), and click Paste Controller (to the right of Copy Controller).
4. Click OK in the Paste dialog.

   The two sub-tracks and the Position parent track are replaced by a single Position track containing the original keys that were in the Old Position track.

   Repeat the above steps with the Rotation track.

To set up dynamics:

1. Assign materials to the objects included in the simulation and adjust the surface characteristics in the Dynamics Properties rollout of the Material Editor. (For a bouncing ball, you'd use this to create a rubber-like surface.)
2. If you're using a linked hierarchy, set the Move and Rotate locks in the Hierarchy/Link Info panel to limit the motion and rotation of the linked objects.
3. Create space warp effects in the scene where needed. (For a bouncing ball, you'd need a Gravity space warp.)
4. Use the Dynamics utility to create a new simulation. Specify which objects are included in the simulation, which effects influence which objects, and which objects should collide with which. (For the bouncing ball, the ball and the floor are in the simulation because one collides with the other. You assign the ball the floor to collide with each other, and assign the gravity effect to the ball.)
5. Use the Dynamics utility to specify the range of frames to which keys will be generated, and to calculate the animation and generate the keys. (In the case of a bouncing ball, a number of position and rotation keys are generated for the ball.)
6. Play the animation to see if the effect is what you were looking for. If one or more objects fly off into space, or move through objects they should have bounced off, it's likely that you need to increase the Calc Intervals Per Frame value.

To reduce the number of keys generated by the Dynamics utility:

When you solve a dynamics simulation, Position and Rotation keys are generated at every frame of the specified range for every object affected in the simulation. Not only does this result in an excess of keys for later editing, but it can increase the size of the *.max* file tremendously. The following steps show how to reduce the number of keys using Track View.

1. After solving the simulation, check the animation to make sure it's what you want, and then save that version of the scene.
2. Deselect every object in the scene.
3. In the Dynamics panel > Objects in Simulation group, click the Select Objects In Sim button.

   All objects included in the simulation are selected.
4. Open a Track View window, and then set its filter to show Animated Tracks Only and Selected Objects Only.
5. Right-click the top object in the Hierarchy list (Objects), and choose Expand All.

   Track View now shows all tracks in the simulation that have keys.

6. Go to Edit Time display mode and select all of the tracks containing keys (or right-click over the hierarchy, and choose Select All).
7. Double-click any key to select all keys in all tracks.
8. Click the Reduce Keys button, set the Threshold to what you want, and then click OK.

   All selected keys are reduced.
9. Save the reduced version of the scene either under a new name, or by replacing the original file.

To use linked hierarchies in a simulation:

When linked hierarchies are included in your simulation, you must set locks for the children in the simulation to confine the dynamics results to specific axes. Do this on the Hierarchy panel > Link Info > *Locks rollout (page 2–500)*.

The Locks rollout contains three rows of check boxes affecting the XYZ axes of the three possible transforms: Move, Rotate, and Scale. The Scale transforms are ignored, and only the Move and Rotate locks are used. When a check box is turned on, that axis of the specific transform is locked.

When you manipulate a forward-kinematics hierarchy directly using the Move or Rotate tools, you might not bother with the Link Info locks, because you can specify axis constraints using the X, Y , Z, and XY buttons in the toolbar. However, when you use that same hierarchy in a dynamics simulation, where there are several forces at work (gravity, wind, collisions), the only thing that maintains the linkage between the objects is the locks you set in the Link Info > Locks rollout. As a result, no matter what combination of Move and Rotate locks you use, you'll always want at least one Move lock in place, or your objects won't really be linked.

The following lists all of the combinations of Move and Rotate locks that make sense within a dynamics simulation, and the effect on the link of such combinations. An asterisk (*) indicates those combinations that are more typically useful.

The format of this list is as follows:

X=check box on.

O=check box off.

One group of settings is made up of the three Move check boxes over the three Rotate check boxes. Here's an example:

XXO=X and Y Move check boxes on, and Z off.

OXO=Y Rotate check box on, and X and Z off.

1. **1 Move Lock:** Turn on any single Move. (This is like a long pin sliding in a loose, long slot.) The joint can transmit force in one direction only. The objects can slide with respect to each other in two directions and rotate freely.
2. **2 Move Locks:** Turn on any two Moves. (This is like a sliding ball joint; a freely rotating joint at the end of a sliding shaft, which can slide and rotate in a hole.)
3. *** 3 Move Locks:** Turn on three Moves. (This is like a ball joint, or the theoretical "pin joint" of the statics and dynamics texts, in that it transmits any force but never transmits any torque.)
4. **1 Move + 1 Rotate (unique):** Turn on any one Move and any one Rotate, but not in the same column. (This is like two long pins, parallel, sliding in a single long slot.) The joint can transmit force in one direction only and restrain rotation about the axis of the "pins." This combination is of limited application.
5. **2 Moves + 1 Rotate (matching):** Turn on two Moves, plus one Rotate turned on that's in the same column as one of the selected Moves. (This is the same as 1 Move + 1 Rotate, above, except that the pins can no longer slide vertically in their slot.) If the assembly rotates

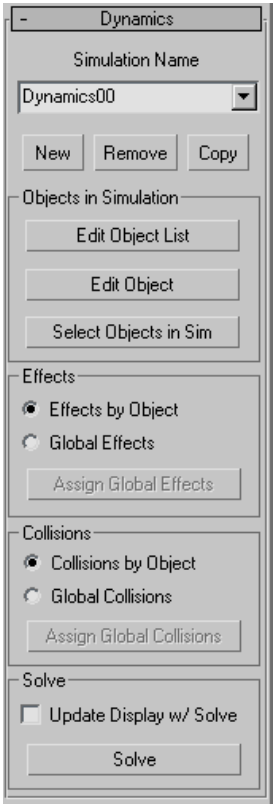
so that one pin travels further into the slot, the other must ride higher in the slot. This is of limited application. The possible check box combinations are: XXO XXO XOX XOX OXX OXX OXO XOO XOO OOX OOX OXO

6. *** 2 Moves + 1 Rotate (complementary):** Turn on two Moves, plus one Rotate that's not in the same column as any of the selected Moves. (This is a sliding universal joint like the splined output shaft between the automatic transmission of a rear-drive car and the drive shaft.) It can transmit torque and constrain translation in two directions, both orthogonal to the axis of rotation. The possible check box combinations are: XOX XXO OXX OXO OOX XOO
7. **3 Moves + 1 Rotate:** Turn on three Moves plus one Rotate. (This is a universal joint without the sliding.) It's typical of automotive applications where the rear axle is located with the trailing drive shaft. This is an uncommon application.
8. *** 1 Move (complementary) + 2 Rotates:** Turn on one Move that's complementary to two Rotates. (This is like a hockey puck on ice.) The joint can slide anywhere on a plane, but cannot fall or tip, and it cannot leave the surface of the plane. The possible check box combinations are: XOO OXO OOX OXX XOX XXO
9. **2 Moves (one complementary) + 2 Rotates:** Turn on two Moves, one of which is complementary to one of the two Rotates that are selected. (This is like a hockey puck with a nail through it, and the nail is sliding along a groove in the ice.) It's free to travel in one direction, and to rotate about an orthogonal axis. One possible check box combination is: XXO XOX
10. *** 2 Moves + 2 Rotates (matching):** Turn on two Moves and two matching Rotates. This results in a sliding axle (a shaft that can both slide in and out of a hole, and rotate with the hole). The clear Move and Rotate axis specifies the axis along which the joint can slide and rotate. The possible check box combinations are: XXO XOX OXX XXO XOX OXX
11. *** 2 Moves + 3 Rotates:** Turn on two Moves and all three Rotates. (This is a prismatic or sliding joint.) The joint transmits no torque, and force in only one direction. You can use this in conjunction with the Push effect to make a hydraulic cylinder. The one clear Move specifies the axis of movement.
12. *** 3 Moves + 2 Rotates:** Turn on all three Moves, and any two Rotates. This is an axle (the most common type of joint.) The one clear Rotate specifies the axis of rotation.
13. **All Locked:** All six check boxes are on. This is a completely rigid joint.

Interface

Dynamics rollout

Contains all of the surface dynamics controls.



Simulation Name—Displays the name of the current simulation. You can edit the name to rename any existing simulation.

You can create any number of simulations in your scene. Each must have a unique name and is stored in the *.max* file. For example, you might have a simulation named Bouncing Ball that bounces a ball down a flight of stairs, while another simulation named Paper Airplane flies a paper airplane across the room.

List—Displays the name of the current dynamics simulation, and lists all simulations in the scene. If the list contains two or more simulations, choose one from the list to make it current. All remaining panel settings are specific to the current simulation.

New—Creates a new simulation. Its name consists of the word "Dynamics" appended by a number, starting with 00. This number is incremented by one for each new simulation.

Remove—Deletes the current simulation. Dynamic simulations can use a lot of memory. Removing old or unused simulations reduces the size of your *.max* files. When you remove a simulation, all timings and other settings are deleted. However, any keys generated by the simulation remain.

Copy—Creates a duplicate of the current dynamics simulation. All of the settings are identical to the original simulation, with the addition of "01" appended to the name.

Objects in Simulation group

Lets you add and remove objects from the simulation, and edit the properties of objects in the simulation.

Edit Object List—Displays the *Edit Object List dialog (page 2–700)*, which lets you specify which scene objects are to be included in the simulation.

Edit Object—Displays the *Edit Object dialog (page 2–696)*.

The Edit Object dialog is the main interface for object dynamics attributes. Use this dialog to set collisions, effects, surface properties and mass for each object in the simulation.

Select Objects in Sim—Adds all objects in the simulation to the current selection set. One use for this function is to bring selected objects into Track View for further manipulation, keyframe reduction, and so on.

Effects group

Specifies which effects are included in the dynamics calculation.

Effects by Object—Only effects assigned to specific objects with the Edit Object dialog > Assign Object Effects button are considered in the calculation.

Global Effects—Only effects included in the Assign Global Effects dialog (accessed by clicking the button of the same name) are included in the calculation.

Assign Global Effects—Displays the Assign Global Effects dialog.

Select effects (space warps) in the list on the left and use the > button to move them to the list on the right. Effects thus chosen affect all objects in the simulation except unyielding ones.

The Assign Global Effects dialog functions similarly to the *Edit Object List dialog (page 2–700)*.

Collisions group

Specifies which collisions are included in the dynamics calculation.

Collisions by Object—Collisions assigned to specific objects through the Edit Object dialog > Assign Object Collisions button are included in the calculation.

Global Collisions—Collisions assigned in the Assign Global Collisions dialog (accessed by clicking the button of the same name) are included in the calculation.

Assign Global Collisions—Displays the Assign Global Collisions dialog.

Select the objects in the list on the left and use the > button to move them to the list on the right. All objects thus chosen collide with each other in the simulation.

The Assign Global Collisions dialog functions similarly to the *Edit Object List dialog (page 2–700)*.

Solve group

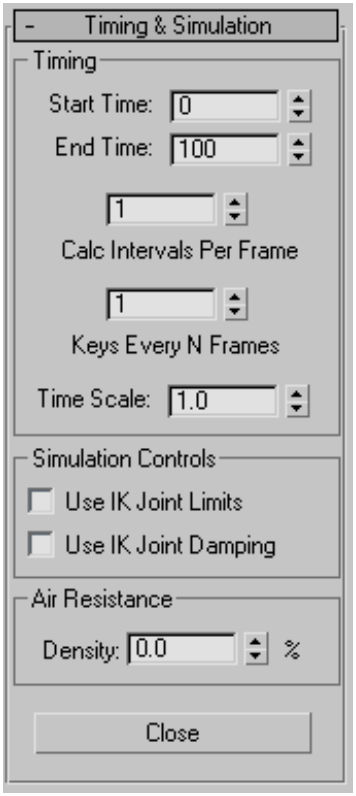
Update Display w/ Solve—Displays each frame of the solution in the wireframe viewports during the calculations. This slows down the calculation process.

Solve—Calculates the dynamics solution, generating keys over the range of frames specified in the Timing area. A progress bar appears in the status/prompt line. Press ESC to cancel the calculation.

Note: You cannot undo the generation of a dynamics simulation solution. If there's a chance you might want to restore the scene to its state prior to the solution, either *save (page 3–390)* the scene or use *Hold (page 1–95)* before solving it.

Timing and Simulation rollout

Lets you specify the range included in the calculation, how IK is included in the simulation, and what the air density is for the simulation.



Timing group

Controls how keys are generated over time.

Start Time—Specifies the first frame to generate keys, which is the first frame to be considered for the solution. Default=0.

Tip: If you set a start time that's later than a keyframe-animated object's last animation frame, you may get unexpected motion during the interim frames. For example, if you animate a box's position from frames 0 to 10, and then use the box in a dynamics simulation that starts at frame 20, the box will move between frames 10 and 20 because of the Bezier controller's default interpolation. To avoid this, before solving the simulation, set the last animation keyframe's Out tangent type to Linear or Step; see *Bezier Controllers (page 2–310)*. Alternatively, set the last key of the keyframed "input motion" *after* the start time of the simulation, or set a start time before the last key.

End Time—Specifies the last key considered for the solution. This spinner is set to the last frame of the *active segment (page 3–904)* when you create a new simulation. For example, if your active segment ends at frame 200, when you click New to create a new simulation, End Time is set to 200.

Calc Intervals Per Frame—Specifies how many calculations are performed for each frame of the simulation time range. Range=1 to 160.

Finding the right number for this spinner is a matter of experimentation. As a general rule, the faster things are moving in the simulation, the higher you should set this value.

Note: If you find that some objects aren't colliding properly with others (they're going through them), increase the Calc Intervals Per Frame value.

Keys Every N Frames—Specifies the frequency with which keys are generated, per object. If this were set to 2, keys would be generated in every other frame.

Warning: When you reduce the key count by increasing this setting, important information can be lost. For example, if a collision occurs on frame n, the Dynamics utility normally sets keys at frames n, n+1, and n-1. But if you've set Keys Every N Frames to 2, a keyframe for the impact itself might not be generated, while keys for the sudden reversal of motion would be generated on either side of the (missing) impact. Thus, the motion controller is left to interpolate motion in a region where the motion should be sharply defined. When this happens, motion can be incorrect and the remainder of the solve is affected. In the aforementioned example, a key describing the impact is lost and the motion controller interpolates motion so that objects that should collide actually intersect, ruining the simulation.

Time Scale—Slows down or speeds up the overall effect of the simulation. It applies a linear scale

factor to the outside forces affecting each object (gravity, wind, and so on).

The default value of 1 results in normal speed. You can scale down the simulation (make it slower) by using values below 1 (from 0.1 to 1), and you can scale up the simulation (make it faster) by using values greater than 1 (from 1 to 100). If you speed up your simulation and objects begin to behave incorrectly (going through objects, for example), increase the Calc Intervals Per Frame value to compensate.

Simulation Controls group

Relates to IK settings and the transfer of momentum.

Use IK Joint Limits—Uses the current IK joint limit settings as constraints for hierarchies in the simulation.

Use IK Joint Damping—Uses the IK damping settings as constraints for hierarchies in the simulation.

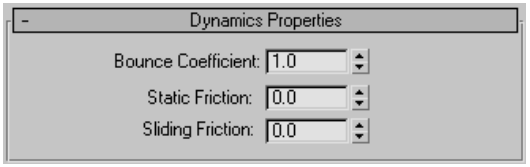
Air Resistance group

Density percent—Sets the air density in the simulation. A setting of 100 is the air at sea level. A setting of 0 is a total vacuum.

When anything moves, it hits air resistance (except in space). The faster it moves, the higher the relative air resistance with the square of the speed. Thus, air resistance imposes an upper limit on the speed of things that are falling with gravity, and also makes objects tumble due to the effect of air resistance on each face of the object.

Close—Closes the Dynamics utility.

Dynamics Properties Material Editor rollout



The three spinners in the Dynamics Properties rollout in the Material Editor let you specify surface properties that affect the animation of an object upon collision with another object. If there are no collisions in your simulation, these settings have no effect.

Because the Dynamics Properties rollout is available at the top level of any material (including sub-materials), you can specify different surface dynamic properties for each face in an object. There are also controls in the Dynamics utility that let you adjust the surface properties at the object level, but only the Material Editor lets you alter the surface properties at the sub-object level (through use of a Multi/Sub-Object material).

As a default, the values in the Dynamics Properties rollout provide a surface that's similar to Teflon-coated hardened steel. This is with values of Bounce Coefficient equal to 1; with both Static Friction and Dynamic Friction set to 0.

Bounce Coefficient—Determines how far an object bounces after striking a surface (the higher the value, the greater the bounce.) A value of 1 represents a bounce in which no kinetic energy is lost.

Static Friction—Determines how difficult it is for an object to start moving along a surface (the higher this value, the more difficult the movement). If something weighs 10 pounds and sits on Teflon (a static friction of near zero), it takes almost no force to make it move sideways. On the other hand, if it sits on sandpaper, then the static friction might be very high, around .5 to .8.

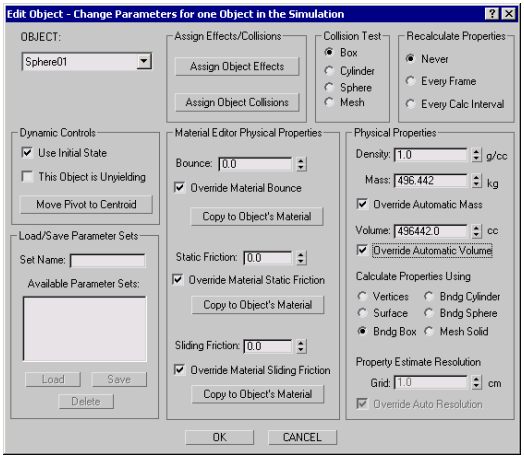
Sliding Friction—Determines how difficult it is for an object to keep moving over a surface (the higher this value, the more difficult for the object to keep moving). Once two objects begin to slide over one another, static friction disappears and sliding friction takes over. Generally, sliding friction is lower than static friction due to surface tension effects. For example, once steel starts sliding over brass (a value of static friction that might run from .05 to .2), the sliding friction drops to a significantly lower value: .01 to .1.

Edit Object Dialog

Utilities panel > Utilities rollout > More button > Dynamics > Dynamics rollout > Edit Object button > Edit Object dialog

The Edit Object dialog is the main interface for setting dynamic properties for any object in the simulation. You assign dynamics properties one object at a time. Select an object in the list under Object, and then set that object's parameters. For example, you can cause the object to be affected by gravity or wind, and specify collisions with other objects.

Interface



Object—Displays the name of the object for which you're setting the dynamic properties. All settings in the Edit Object dialog affect the object listed here. To change the object you're affecting, open the list, and choose from a list of all objects assigned to the simulation.

Dynamic Controls group

Use Initial State—When turned on, the simulation solves for the motion, energy, and momentum of the object at the simulation's starting frame (specified in the Start Time spinner in the Time and Simulation rollout > Timing group.) If Use Initial State is turned off, the object is assumed to be sitting motionless at the starting frame. The Dynamics utility overwrites keys. Either the object's keys are clocked with the simulation, or it's not moving (or both). If you don't solve for the initial state, the object's keys are interpolated at the instant the simulation starts and all subsequent motion comes from the interaction of the simulation itself.

This Object is Unyielding—Lets you use keyframe-animated objects in dynamics simulations. Objects with this option turned on are immovable relative to other objects that collide with them, but can be animated (keyframed).

You can animate objects with this option turned on by themselves or as part of a keyframed hierarchy. For collisions, objects with Use Initial State turned on cannot move objects with This Object is Unyielding turned on.

Move Pivot to Centroid—When this button is on (green) and you exit the Edit Object dialog by clicking OK, the object's pivot point is moved to its center of mass. Having the pivot aligned with the center of mass speeds up the dynamics calculation and makes manual keyframing of the objects appear more natural. However, if the pivot point is already positioned as part of a linked hierarchy,

turning on Move Pivot to Centroid will alter the positioning of your hierarchical linkage.

Load/Save Parameters Sets group

Saves the current Edit Object parameter settings, or loads previously saved parameter settings. You can use these controls to save the settings for one object, and then apply them to another object.

Set Name—Accepts a new name for a parameter set. Click Save to save the parameter set with this name.

Available Parameters Sets—Lists previously saved parameter sets.

Load—Loads the selected parameter set.

Save—Saves the named parameter set. The saved parameter sets appear in the list.

Delete—Deletes the selected parameter set.

Assign Effects/Collisions group

Let you specify which effects in the scene will affect the current object, or which objects in the scene can collide with the current object.

Assign Object Effects—Displays a subdialog that lets you specify which effects in the scene will affect the current object. The dialog consists of two list windows. All effects (typically space warps) in the scene are listed in the window on the left. Select effects in the list, and then click the > button to transfer the highlighted effects to the window at right.

The Assign Object Effects dialog functions similarly to the *Edit Object List dialog (page 2–700)*.

Assign Object Collisions—Displays a subdialog that lets you choose which objects in the scene are considered for collision with the current object. This dialog works the same as the Assign Object Effects dialog, except that it lists only objects in the simulation. Objects included for collision in this dialog can collide with the current object.

Note: For every potential collision, you should explicitly specify both colliding objects. For example, if you specify that Box01 is to collide with Box02, you should also specify that Box02 is to collide with Box01. This is primarily for keeping track of your simulation; if you specify the collision for only one object, it works for the other as well. While it is not necessary to specify that Box02 collides with Box01, this is the only way to update the Assign Object Collision list for Box02.

The Assign Object Effects dialog functions similarly to the *Edit Object List dialog (page 2–700)*.

Collision Test group

Specifies the type of boundary used for collision testing.

Box—A bounding box is used to test collision. This is the fastest method. If your object is a box or close to a box shape, this option is both fast and accurate.

Cylinder—A cylindrical form is used to test collision. The "height" axis of the cylinder is aligned with the local Z axis of the object.

Sphere—A spherical form is used to test collision.

Mesh—Uses the surface of the object to calculate the collision. This option is the most accurate, but also takes far longer to calculate than the previous three methods. Use this only when your object is too complex to work properly with the first three options.

Recalculate Properties group

Specifies when to recalculate the properties of an object that changes over the course of the animation. Specifically, the properties refer to the mass moment of the object. The mass moment of an object can be defined as the measure of how

the mass in an object is distributed with respect to the center of mass of that object. Altering the mass moment changes the way an object responds to torque. If your object is changing shape over time, and you want to take this into account for the simulation, you should recalculate the properties.

Never—No calculation is performed beyond the first calculation for the dynamics solution.

Every Frame—The object's properties are recalculated for every frame.

Every Calc Interval—The object's properties are recalculated at each "Calc Interval." You set the Calc Interval (interval of calculation) in the Time Parameters rollout; it specifies how many calculations are performed for each frame of the simulation time range.

Material Editor Physical Properties group

Three of an object's physical properties are taken by default from its material: bounce, static friction, and sliding friction; see *Dynamics Properties Rollout (page 2–1479)*. Use these settings in this group to accept or override the material's dynamics properties.

Bounce—Specifies how far the object bounces after collision. Available only if Override Material Bounce is on.

Override Material Bounce—Enables the Bounce spinner when you want to use a specific value that's different than that assigned by the object's material. When using a multi/sub-object material, you can assign different Bounce values to the sub-materials, and thus to different faces of an object. But if you use an override value, it affects the whole object only.

Copy to Object's Material—Copies the value in the Bounce spinner to the material assigned to the object.

Static Friction—Specifies how hard it is to start moving on a surface. Available only if Override Material Static Friction is on.

Override Material Static Friction—Enables the Static Friction spinner when you want to use a specific value that's different than that assigned by the object's material. When using a multi/sub-object material, you can assign different Static Friction values to the sub-materials, and thus to different faces of an object. But if you use an override value, it affects the whole object only.

Copy to Object's Material—Copies the value in the Static Friction spinner to the material assigned to the object.

Sliding Friction—Sliding friction determines how hard it's for an object to keep moving over a surface. Available only if Override Material Bounce is on.

Override Material Sliding Friction—Enables the Sliding Friction spinner when you want to use a specific value that's different than that assigned by the object's material. When using a multi/sub-object material, you can assign different Sliding Friction values to the sub-materials, and thus to different faces of an object. But if you use an override value, it affects the whole object only.

Copy to Object's Material—Copies the value in the Sliding Friction spinner to the material assigned to the object.

Physical Properties group

Provides controls that specify the physical properties of the current object.

Density—Specifies the density of the object in grams per cc. A setting of 1 is the equivalent of water, and useful for anything wooden, plastic, or organic. The more dense an object, the slower it will react to forces. There is a direct relationship between density and mass, so when you alter the

Density setting, the Mass value changes (when its automatic value is not overridden).

Mass—Technically, mass is the measure of how many subatomic particles are in an object. In 3ds Max, it's calculated as a result of density times volume. Mass can is the resistance to acceleration given a constant force. The greater the mass, the more resistance. The Mass value, when not overridden, is derived from the density and volume, and the volume, when automatically calculated, is affected by the Calculate Properties Using option.

Override Automatic Mass—Enables the Mass spinner so you can specify the object's mass. When you override the automatic value, the Mass value is no longer affected by the Density or Volume values.

Mass is affected by the Calculate Properties Using options.

Volume—The volume of the object, measured in cubic centimeters. When automatically calculated, this value depends on the option chosen under Calculate Properties Using.

Override Automatic Volume—Enables the Volume spinner so you can specify the object's volume. When you override the automatic Volume value, it's no longer affected by the Calculate Properties Using options, and it no longer affects the Mass value.

Calculate Properties Using options

Specifies the type of geometry the simulation will use to calculate the mass, the volume, and the mass moment. When the Mass and Volume values are automatically calculated (their overrides are turned off), changing these options affects the volume, and thus the mass.

Vertices—Treats an object as a "point cloud," or a collection of vertices without segments. Each vertex is given a mass of 1 gram, but the object itself has no volume.

Surface—Treats the object as a hollow shell whose thickness is 1 centimeter. The mass is derived from the surface area and the 1-centimeter thickness, but the object has no volume.

Bounding Box—A bounding box surrounding the extents of the object is used to calculate both the volume (of a solid bounding box) and the mass (based on the volume).

Bounding Cylinder—Similar to the Bounding Box option, except that a bounding cylinder is used, whose Z height axis is aligned with the local Z axis of the object.

Bounding Sphere—Similar to Bounding Box, except that a bounding sphere is used.

Mesh Solid—The geometry of the object is used to calculate both the volume and the mass, according.

Property Estimate Resolution options

Available only when the Mesh Solid option is chosen. The Property Estimate Resolution settings affect the accuracy with which the mass moment is calculated. This determines how accurately the rotation of the object responds to torque.

Grid—A sample grid is used to calculate the mass moment. The smaller the cells in the grid, the more accurate the calculation. The grid cells are measured in centimeters, regardless of the current display unit.

Override Automatic Resolution—Lets the Dynamics utility automatically set the grid size based on the complexity and size of the object. In addition, the lowest value that will be automatically calculated is plugged into the Grid spinner (which is also unavailable). If you want to specify your own grid value using the Grid spinner, but want some idea of where to start, turn off Automatic Resolution and then turn it back on again.

Edit Object List Dialog

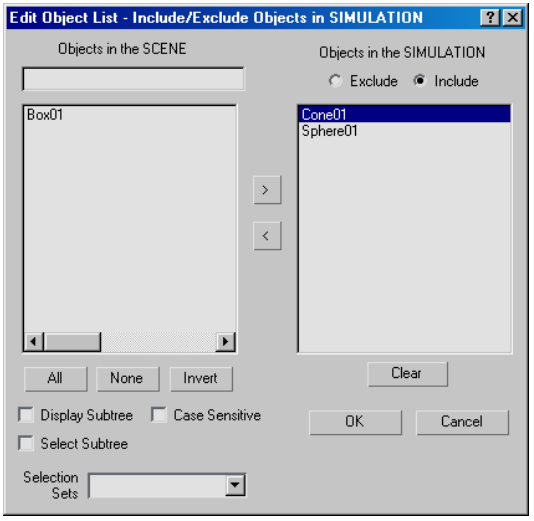
Utilities panel > Utilities rollout > More button > Dynamics > Dynamics rollout > Edit Object List button > Edit Object List dialog

The Edit Object List dialog lets you select the objects to include in a dynamics simulation. Highlight objects from the Objects in the Scene list on the left, and then click the > button to place them in the Objects in the Simulation list on the right side of the dialog.

The Exclude and Include option buttons above the right window determine whether the items listed are excluded from or included in the simulation.

Note: All the selection and inclusion dialogs in the Dynamics Utility are displayed and function similarly.

Interface



Objects in the Scene—To add scene objects to the simulation (assuming Include is chosen), select objects from the list on the left, then use the > button to move them to the list on the right.

Search Field—The edit box above the Scene Objects list lets you search for object names by entering names that use wildcards. For example, searching for "sphere*" finds all objects whose names start with "sphere".

Objects in the Simulation—To exclude objects from the simulation (assuming Include is chosen), select objects from the list on the right, then use the < button to move them to the list on the left.

Exclude/Include—Choose whether the simulation will exclude or include the objects named in the list on the right. If you choose Exclude, only objects in the list on the left are included in the simulation.

All/None/Invert—Affect the list on the left. All selects all objects, None deselects all objects, and Invert selects unselected objects and deselects selected objects.

Display Subtree—Turn on to indent the list according to the object hierarchy.

Case Sensitive—Turn on to use case sensitivity when searching for object names.

Select Subtree—Selects all objects in a hierarchy at the level you click and below.

Selection Sets—Displays a list of named selection sets. Choosing a selection set from this list selects those objects in the Scene Objects list.

Clear—Clears all entries from the Exclude/Include list on the right.

Skin Utilities

Utilities panel > More button > SkinUtilities

Skin Utilities provide a method for copying skin data (envelopes and vertex weights) from one model to another. Skin Utilities work by embedding the skin data in a copy of the source mesh, then using the object copy to map the data onto the target mesh.

To use Skin Utilities, you must have two meshes to which the *Skin modifier (page 1–791)* has been applied, both in the current scene. You must have already assigned bones to the Skin modifier for both meshes.

Note: Skin Utilities work by matching bones from one character to another, so the task of copying the skin data will be greatly simplified if you name the bones in each mesh with similar naming conventions.

Procedure

To use Skin Utilities:

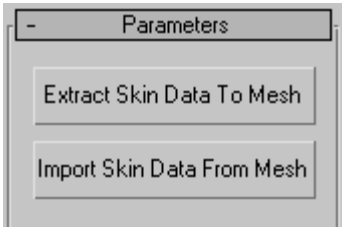
1. Load a scene that contains two skinned meshes, one from which you want to extract the skin data, the other to which you want to paste it.
2. Select the source mesh, the mesh with the correct envelopes and vertex weights.
3. Choose Utilities panel > More button > Skin Utilities.
4. Click Extract Skin Data to Mesh.

   A new object is created, identical to the selected object. This is a new mesh with all envelope-assigned and manually-assigned vertex weights "baked" into the mesh. The object is named with the prefix *SkinData_* followed by the same name as the original object.
5. Move the skin data mesh on top of the mesh to which you would like to paste the skin data, so the two meshes are coincident.
6. If necessary, adjust the skin data object at the Vertex sub-object level to make it fit the other mesh as closely as possible.
7. Select both the skin data object, and the object to which you want to paste the data.
8. Click Import Skin Data From Mesh.

   The Paste Skin Data dialog appears.
9. Match up the bones by highlighting one on each side of the dialog, and clicking the left arrow to move the target bone to its match on the left side.

   You can also highlight several bones on each side at once, and click the left arrow to move them all over at the same time. Only bone pairs on the left side of the dialog will be pasted.
10. If the meshes' vertices do not match exactly, set the Interpolation method to Match by Face, and increase the Threshold to a unit value that will allow the meshes' faces to match up.
11. Click OK to paste the vertex weights and close the dialog.
12. Delete the skin data mesh.

Interface

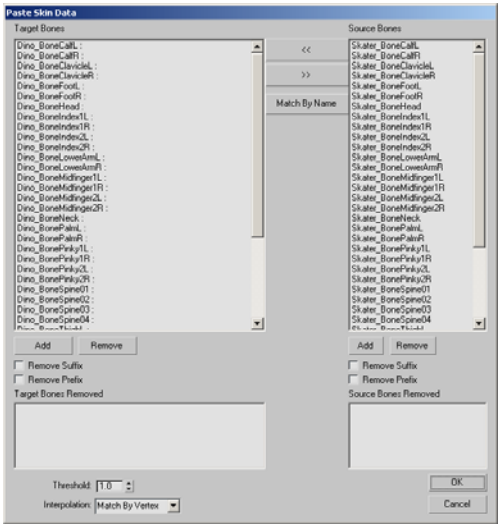


Extract Skin Data to Mesh—Extracts the skin data from the selected mesh and embeds it in a new mesh named *SkinData_* followed by the original object name.

Note: The skin data is stored in channels. To see the data, select the *SkinData_* mesh and use the *Channel Info Utility (page 2–1738)*.

Import Skin Data From Mesh—Opens the Paste Skin Data dialog. Before clicking this option, you must select the *SkinData_* mesh and the mesh to which you want to paste the data.

Paste Skin Data dialog



To use the Paste Skin Data dialog, highlight matching Target Bones and Source Bones, and click the left arrow to match them. Only matched sets of bones listed under Target Bones will be mapped when you click OK.

Target Bones—Lists target bones available for matching, and lists matches moved from the Source Bones listing with the left arrow.

Source Bones—Lists source bones available for matching.

Add—Adds removed bones back to the source or target list.

Remove—Removes highlighted bones from the source or target list.

Remove Suffix—If bone names end with an underscore followed by text, this removes the suffix from the displayed bone names from the underscore to the end of the name. Use this option to remove suffixes and make source and target bone names match, which will allow you to use the Match by Name option.

Remove Prefix—If bone names start with text followed by an underscore, this removes all characters from the displayed bone names up to and including the underscore. Use this option to remove prefixes and make source and target bone names match, which will allow you to use the Match by Name option.



Left Arrow—Matches highlighted target bones to highlighted source bones. The source bone name is moved to the left side of the dialog, where it follows the target bone name. You can also match more than one source bone to a single target bone by highlighting the target bone and the source bones and clicking the left arrow. When you click OK, only the matches listed on the left side of the dialog will be mapped.



Right Arrow—On a highlighted match on the left side of the dialog, removes the match and moves the source bone to the right side of the dialog.

Match by Name—Matches the source and target bones by name. You can use this option only if you have removed prefixes or suffixes to make the displayed source and target bone names match exactly.

Threshold—Sets the unit distance that the utility will search for vertex or face matches. Default=1.0.

Interpolation—Sets the method by which pasting will take place. Use the Match by Vertex option only if the meshes are identical. Otherwise, use Match by Face. Default=Match by Vertex.

13 reactor

reactor is a toolset that allows animators and artists to control and simulate complex physical scenes in 3ds Max. reactor supports integrated rigid and soft body dynamics, cloth simulation, and fluid simulation. It can simulate constraints and joints for articulated bodies. It can also simulate physical behaviors such as wind and motors. You can use all of these features to create rich dynamic environments.

Once you have created an object in 3ds Max, you can assign physical properties such as mass, friction, and elasticity to it with reactor. Objects can be fixed, free, attached to springs, or attached together using a variety of constraints. By assigning physical characteristics to objects like this, you can model real-world scenarios and then simulate them to produce physically accurate, keyframed animations.

After you set up your reactor scene, you can preview it quickly using the real-time simulation display window. This allows you to test and play with a scene interactively. You can alter positions of all physical objects in the scene, dramatically reducing the design time. You can then transfer the scene back into 3ds Max with a single mouse click while retaining all the properties needed for the animation.

reactor frees you from having to hand-animate time-consuming secondary effects, like exploding buildings or draping curtains. reactor also supports all standard 3ds Max functionality such as keyframes and skinning, so you can use both conventional and physical animation in the same scene. Convenient utilities, such as automatic keyframe reduction, let you tweak and alter the physically generated parts of an animation after it has been created.

The remainder of this chapter describes each of reactor's features in detail. Also, the included tutorials step you through creating some typical reactor scenes. Together, we hope these will help you to get the most from reactor.

If you would like to find out more about dynamics simulation, see *Introducing Dynamics Simulation (page 2–707)*.

Getting Started

This section shows you where to find the various reactor options in 3ds Max, as well as introducing you to reactor's helper icons. You'll see how to use each of the options in the relevant section of this guide.

Command Panel

You can use the reactor options on the Create panel to create various reactor elements. To find most reactor objects, go to the Helpers sub-panel, and then, from the drop-down list, choose reactor.



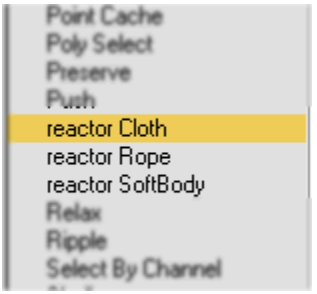
You can also find a space warp, used for *water (page 2–801)*, in Space Warps > reactor.



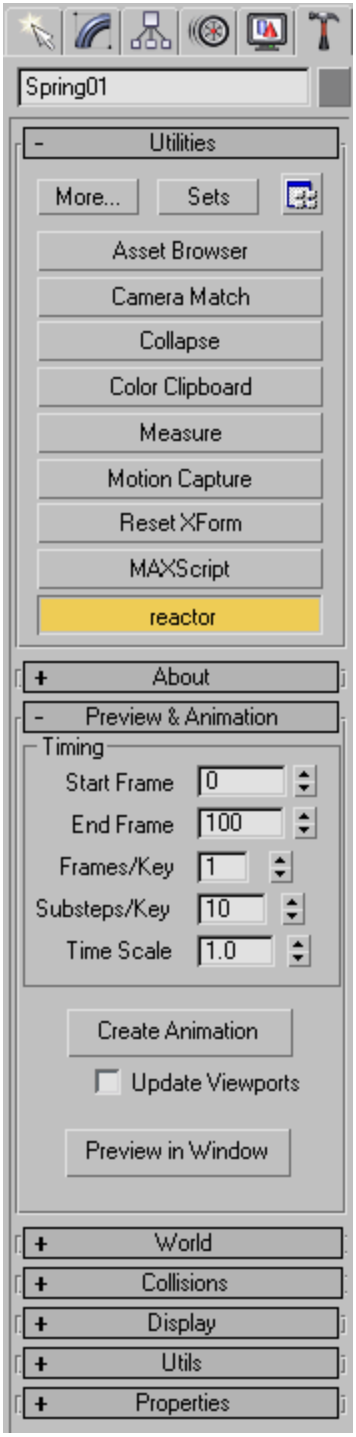
Once you've created a reactor object, selecting the object and opening the Modify panel allows you to configure its properties.



There are also three reactor modifiers, used to simulate *deformable bodies (page 2–777)*:



You'll find most of the remaining reactor functions on the Utilities panel. This provides access to functionality such as previewing the simulation, changing world and display parameters, and analyzing the convexity of objects. It also lets you see and edit the rigid body properties associated with objects in the scene.



Note: The reactor menus and toolbar provide shortcuts for many of the reactor functions provided in the command panel.

The reactor Toolbar

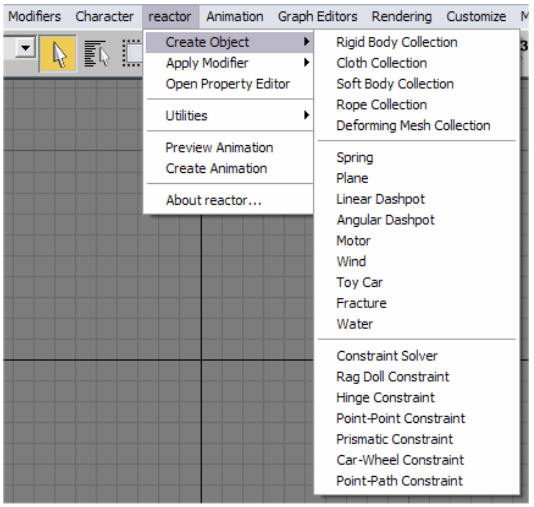


The reactor toolbar is a handy way to access much of reactor's functionality. It has buttons that let you quickly create constraints and other helpers, display physical properties, generate animations, and run the real-time preview.

To display the reactor toolbar:

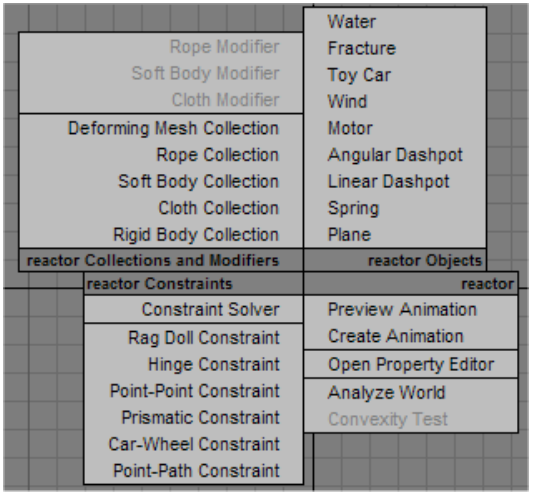
1. Right-click an empty area of the main toolbar to display a list of available toolbars.
2. Click reactor.

The reactor Menu



The main reactor menu is another way to access reactor functionality.

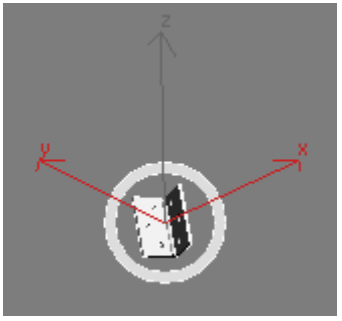
The reactor Quad Menu



A further quick way to access reactor options is the reactor quad menu. To open this menu, press Shift+Alt and right-click in the active viewport.

Helper Icons

Many reactor elements, such as constraints and the Rigid Body Collection, have their own special helper icon that appears in the viewport when you add them to the scene. For instance, the following illustration shows the Hinge constraint icon:



Although the helper icon doesn't appear in your rendered scene, the icon's appearance (and in some cases, its position and orientation) will help you to set up your reactor scene correctly.

When selected, reactor icons are white and are also larger than when not selected. When not selected, the icon for a *valid* element is blue, and for an *invalid* element is red. What constitutes validity depends on the particular reactor element. For instance, a *hinge (page 2–747)* is valid if it has the correct number of objects attached to it; a *Rigid Body Collection (page 2–723)* is valid if it's not empty. Invalid elements are excluded from the simulation, and reported as errors.

Certain icons provide additional information about how the element behaves in the simulation. For example, the display for a valid hinge indicates the hinge position, and, when selected, the hinge axis and any limits you have specified for the movement of the hinged bodies.

Introducing Dynamics Simulation

This section gives a broad overview of physical simulation, and introduces some concepts you might find useful when working with reactor.

Physical Simulation

Havok's physics technology, used in reactor, relies on a process known as *physical simulation* in order to provide a dynamic environment for the objects in a scene. So what exactly does *physical simulation* mean?

It is a process that automatically determines the motion of objects according to their physical properties. It achieves this by encapsulating some physical laws, like Newton's laws of motion, within a robust, efficient engine, which calculates the position of each object as time passes. In much the same way that a motion picture is composed of many individual images or frames, physical simulation splits time into small discrete steps and predicts the motion of each object during each step. The cumulative effect of all these steps is fluid, continuous, believable motion.

Unlike traditional keyframe-based animation, where the animator needs to specify a set of keyframed configurations, physical simulation determines how objects move based their properties. This takes the burden from the animator, who now doesn't need to manually animate every piece in an explosion, every bone in a character stunt animation, or every vertex in a piece of cloth.

In a physical simulation, physical properties such as mass and elasticity are assigned to all objects in a scene. This is then complemented with a set of external forces like gravity or wind and/or constraints, like a spring force or a ball-socket configuration. From all this information, the physics engine calculates a continuous set of states that can then be displayed in real-time, if the calculations are fast enough, or converted into keyframes to be played back later.

What Does A Physics Engine Do?

A physics engine like Havok has three basic tasks to perform:

1. Collision Detection

   Track the movement of all the objects in the scene and detect when any of them have collided.

2. Update System

   Determine an appropriate response for objects that have collided by resolving the collision according to the object properties and for all other (non-colliding) objects update them according to the forces acting on them.

3. Interface with Application

   Once reactor determines the new positions and state of all objects, it can display the objects in a 3D window or store their states in the form of keyframes.

Note: A physics engine knows nothing about how the objects it is simulating are displayed. It simulates the motion and interaction of these objects based on a physical (not graphical) description of the objects, and this information can be used to generate a display that "tracks" the simulation.

Given that we are talking about simulating a continuously evolving state (in other words, objects are moving and colliding and reacting all the time in general), we need to map this to a series of snapshots in order to generate an animation. In a computer game, for example, we typically want to display the world 60 times per second, because this is how frequently many graphics systems can redraw the screen. When creating an animation, we may want to store the state (that is, create a keyframe) for every frame. At 60 FPS, for example, this would also create 60 keyframes every second, What this really means is that the physics engine must be capable of evolving the world by 1/60th of a second knowing the state of all the objects at the start of this time interval and knowing the external forces acting on these objects. As an example we'll look at the simple case of a cannon ball and we'll assume we're interested in animation at 60Hz (Hz = cycles or frames per second).

Simulating a Cannon Ball

Let's forget about collisions for now, and consider only the simulation of a cannon ball immediately after it has been fired from the cannon. We know the ball's position (and orientation, but we'll ignore this for now), speed, acceleration, and weight, and we assume we know the state of the environment (air resistance, wind force, gravity). Armed with this knowledge we can start to make predictions using a physics engine.



This figure illustrates what we would like to achieve. Over a period of time the cannon ball's rate of ascent should slow due to gravity, and it should eventually fall to the ground having traveled through a classic parabolic arc (assuming no air resistance).

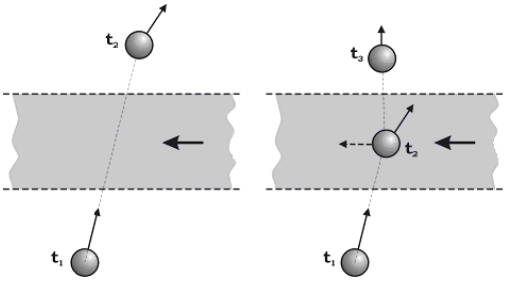
At a given point in time we can examine the state of the ball (its speed *v* and acceleration *a*) and knowing the external forces acting on it we can make a guess as to its change in position after a period of time has elapsed (call this period *h* seconds). This guess is a combination of a number of factors:

- We assume that Newton's laws of motion govern the motion of the ball
- We assume that in the time period *h* all the external forces acting on the ball are constant, so air resistance and wind and gravity do not change during this time.
- We assume that the math we use to calculate the new position is accurate

In general, the first assumption is usually valid, except at relativistic or quantum scales, which we can assume should be handled by other systems. The remaining two, however, cause problems and are closely linked to the time period *h* over which we're performing the calculations. We'll now examine the effect of the size of this time period on the accuracy of the simulation.

Time Steps

In general, the forces acting on an object are rarely truly constant; gravity is close to being constant, but most other forces like wind and air resistance are not. So, taking the cannon ball example, imagine a windy layer in the atmosphere that the cannon ball passes through, as shown in the next figure.



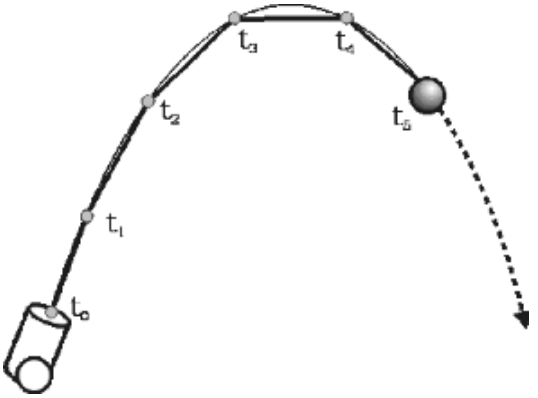
In the simulation on the left we assume we're taking steps of one second; this is actually a relatively large interval for a physics simulation, but is used here to illustrate the point. We know all the forces acting on the ball at time t1 so we use some math to predict the new position and velocity at time t2, after one second has elapsed. During this period, we assume that the wind force acting on the ball is constant. In this example, we'll calculate the new position above the region of high wind, so we'll effectively have missed the windy bit by taking too great a jump. In the second example on the right, we're using time steps of ½ second. In this case, after determining the new position at time t2 we find the ball in the middle of the windy region. This region causes a large wind force to act on the ball which is taken into account during the next time step. At that point we reevaluate the math and determine a new position for the ball at time t3. This is different from the position determined in the simulation on the left, even though the same amount of time has been simulated in each case. In other words, the wind has blown the ball to the left a bit and has reduced the velocity of the ball

In general, the smaller the time step taken, the more accurate the result at the end of the time step. Thus, to step forward in time by a large time step

t it is better to split this into n steps of a smaller time interval t/n.

This is also true of the math. As the simulation becomes more complex, the math required to calculate the new positions and velocities of objects in a simulation also becomes more complex, and as a result the guesses produced by the math give progressively less-accurate results.

So the principle is to take small time steps, evaluate all the forces acting on the objects, determine the new positions and velocities (and other parameters) of the objects at the end of the time steps, and then start over. We end up with a series of snapshots of the state of the system as it evolves, as shown in the following illustration.

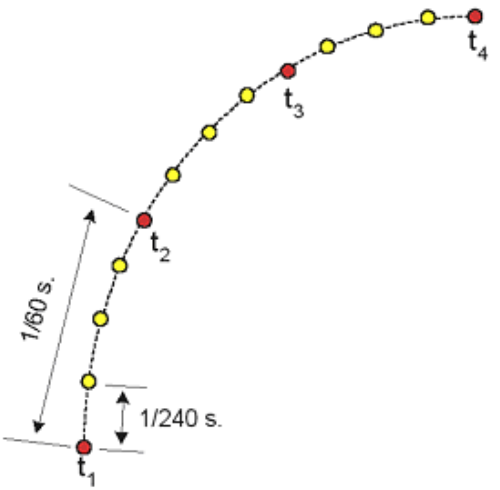


Substeps

In the previous section, we said that the physics engine passes on the necessary information to update the display after it determines the new positions for all its objects. However, what if you don't want to update the display in each simulation time step?

Let's say we absolutely need to update the application once every 1/60th of a second, either because we're playing a real-time game that refreshes the screen at 60Hz. or we're creating a keyframe every frame for a 60FPS movie. This effectively means that we want to step the physics engine at intervals of 1/60th of a second. In many cases this does not present a problem, but if, hypothetically, the accuracy of the simulation was not sufficient (remember: smaller time steps mean better accuracy), then we'd like to decrease the time step even further, say to 1/120th of a second. But this would mean we generate twice the number of images/keyframes we are interested in, which is wasteful. To get around this, the Havok engine allows you to specify the number of substeps to take per key.

The Substeps/Key parameter specifies the number of steps the physics engine takes before updating the application. This gives control over the granularity of the physics simulation independent of the display update or keyframe-creation frequency. So with Substeps/Key=1, reactor uses a single simulation step for each update to the 3D display. With Substeps/Key=2, reactor takes two physics steps each time before updating the application, and so on.



In this figure, we have specified that the physics simulation should step at intervals of 1/240th of a second, but that we create a keyframe only once every 1/60th of a second. The red spots indicate a keyframe and simulation step, the yellow spots

are simulation steps only. This was achieved by instructing the physics engine to employ four substeps per key. Thus, for each four simulation steps we create only keyframe. By setting the number of substeps we can control the accuracy of the physical simulation independent of the number of keyframes created.

Rigid Bodies and Deformable Bodies

Rigid Bodies

Havok simulates most objects in a simulation as *rigid bodies*. A rigid body is an object whose geometry doesn't change over the course of the simulation. You can simulate any real-world object that doesn't noticeably change its shape, from a pen to a boulder hurtling down a mountainside, as a rigid body. Simulating objects in this way facilitates rapid physical simulation in real time; the physics engine can make certain assumptions when detecting collisions based on the fact that the objects' shapes don't vary from simulation step to simulation step.

Both the Havok 1 and Havok 3 engines can simulate rigid bodies, but Havok 3 simulations are faster and more accurate.

Deformable Bodies

To simulate cloth, rope, or other material whose shape changes over time, you need to use a different type of body: a deformable body. With deformable objects, collision detection becomes much more difficult, given that the object can change shape dramatically between time steps and can also attempt to collide with itself. For this reason, deformable bodies are more expensive to simulate.

Note: Only the Havok 1 engine can simulate with deformable bodies.

Scale

Our scientific knowledge of physics is extensive. What we are concerned with here would more accurately be described as a mechanical simulation of the interactions of objects at real-world scales. We are dealing with Newtonian mechanics; that is, the well-understood laws of motion, popularized by Sir Isaac Newton, that describe the behavior of objects under the influences of other objects and external forces. Since then we've discovered that these laws break down at very small (i.e. subatomic) and very large (i.e. planetary) scales.

New physics systems have been devised to work with these scales (for example, relativistic and quantum), but these are beyond the scope of the reactor physics engine. The Havok physics simulation technology works at the scale of objects we interact with on a daily basis, such as chairs, cars, buildings, and footballs. By default, the engine works in units of meters and kilograms.

It's important to keep in mind the scale in which you're working. For example, a common mistake people make is to start by creating a cube 100 meters on a side, and then they wonder why it takes so long to fall. A box of this size when viewed at a distance sufficient to be able to see the entire box (say 1km away) will appear to fall at the same speed as an aircraft hangar dropped from a height and viewed from a kilometer away: slowly.

Changing Scale

The Havok physics engine does not care what units of measurement you use when specifying the size of objects or the strength of gravity; it cares only about the numbers. So you could, for instance, work in inches. However, for realistic (or at least predictable) results, it's important to be consistent. So, for instance, if you're working in meters, make sure gravity is set to an appropriate value in meters. To produce Earth-like gravity, use

9.8 m/s^2. Otherwise objects might appear to fall faster or slower than you expect.

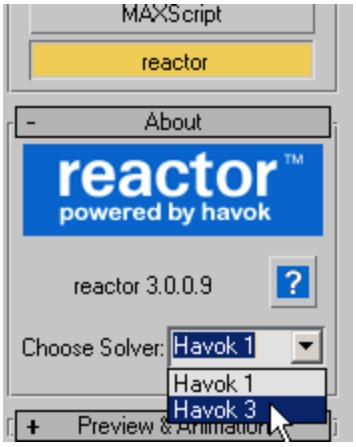
Note: Due to CPU floating-point precision, a physics engine is most accurate when dealing with numbers as close in magnitude to 1 as possible. In other words, values like 10,000,000 work poorly, as do values like 0.0000001. Therefore, for real-world scenes, when creating objects of 1*1*1 size it is most useful to be working in meters or feet, rather than centimeters/inches or kilometers/miles, in that you most often simulate objects larger than sugar cubes and smaller than football fields. It is for this reason that the default values in many physics engines are usually specified in meters.

Special Features in reactor

A number of reactor features enhance speed and usability. The software lets you create new and better animations that were virtually impossible without reactor, and you can set up those animations in an efficient, intuitive way. Here we present some of the main features in reactor.

Two Simulation Engines

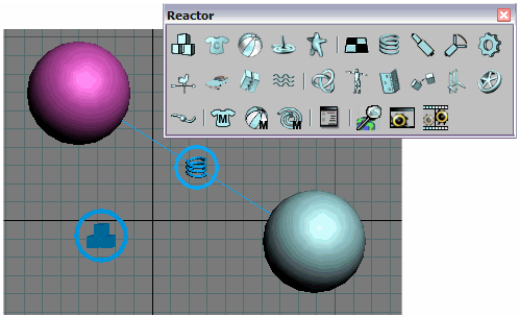
From the reactor utility interface, you can choose to run your simulation with either of two engines: Havok 1 or Havok 3.



If your reaction involves cloth, rope, or soft bodies, you'll need to stick with Havok 1, the version included with previous releases of 3ds Max. However, if you're using rigid bodies only, you can take advantage of the enhanced accuracy available in Havok 3. For more information, see *The reactor Utility (page 2–805)* and *Havok 1 World/Havok 3 World Rollout (page 2–808)*.

Integrated User Interface

The reactor interface is fully integrated in 3ds Max. You can access functionality in reactor through menus, quad menus and toolbars. All 2D and 3D icons follow a consistent look and feel. The *reactor utility (page 2–805)* parameters are arranged for easy access.

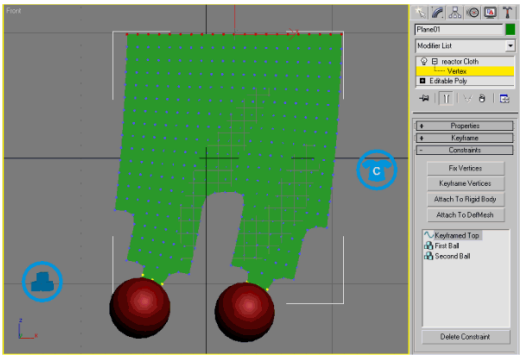


You can find more details in:

- *reactor (page 2–703)*
- *The reactor Utility (page 2–805)*

Vertex Selection Inside reactor Modifiers

The reactor modifiers for *deformable bodies (page 2–777)*: Cloth, Soft, and Rope, let you select vertices and apply constraints (fix points, keyframe points, attach to rigid body, attach to deforming mesh) to those vertices without leaving the modifier or having to apply extra modifiers. You can create and manipulate *deformable constraints (page 2–795)* inside the modifier.



Cooperative Constraints

reactor includes a number of *cooperative constraints (page 2–735)* that facilitate the simulation of articulated bodies and machinery.



The *Rag Doll constraint (page 2–737)* allows the simulation of constrained bodies with relative rotation and twist angles limits similar to those found in human and animal joints

The *Hinge constraint (page 2–747)* allows the simulation of hinges (limited or not) and hinge-like joints where movement is limited around a specified axis, like elbows and knees

The *Prismatic constraint (page 2–754)* allows the simulation of translation-only joints (limited or not) like those found in robots and machinery

The *Car-Wheel constraint (page 2–757)* is particularly suited to the simulation of wheels attached to a chassis. The wheels rotate (and can be powered) about a given axis. Limited linear motion relative to the chassis is allowed along a user-defined suspension axis.

Using *Point-Point constraint (page 2–750)* with the Limited option lets you limit the relative rotation of the attached objects by a given degree around each axis.

Using *Point-Point (page 2–750)* with the Stiff Spring option constrains both objects as though attached by a fixed-length bar (a very stiff spring).

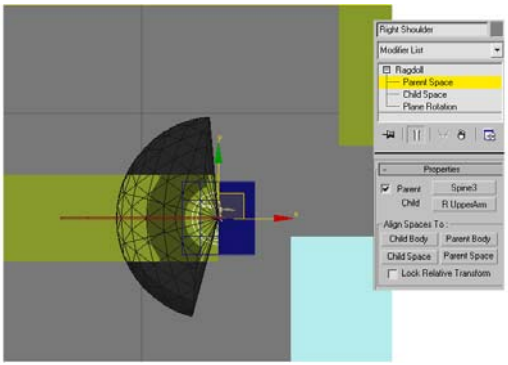
Fracture Object

The *Fracture (page 2–770)* object in reactor offers excellent usability and behavior. Objects inside Fracture are standard Rigid Bodies and, as such, you can add them to constraints, assigned initial velocities, etc. You can tell pieces to break at a specific time.



Constraint Manipulation and Constraint Spaces

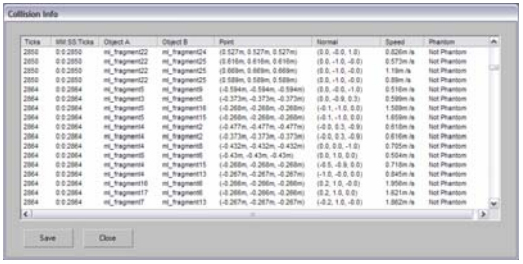
Constraints (page 2–724) in general restrict the relative movement between two bodies (parent and child) or between a body (child) and the world, and use consistent, intuitive nomenclature and setup methods. Springs, linear dashpots and angular dashpots are referred as *Simple Constraints (page 2–727)*, while the other constraints (those that are grouped with a *Constraints Solver (page 2–736)*) are referred to as *Cooperative Constraints (page 2–735)*.



In reactor, you define the effect of constraints by the specifying and manipulating two *constraint spaces (page 2–725)*. Limits, such as minimum and maximum rotation angles, are defined and displayed around those spaces, which you can modify using sub-object manipulation. reactor provides tools for automatically aligning and manipulating those spaces.

Storage and Access of Collision Information

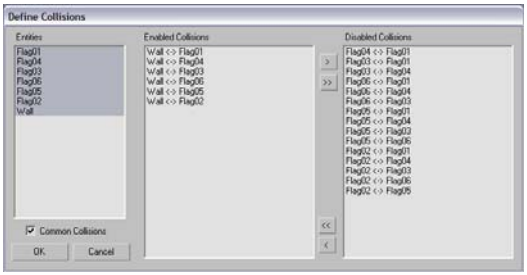
reactor can store information of all rigid body collisions occurred during the simulation. The information includes the objects involved, the point of collision, and the relative velocity during collision, and can be used by animators to generate particles or other effects, trigger sounds, etc. You can access the information via MAXScript and save it to a text file.



For more information, see *Storing and Accessing Collisions (page 2–774)*.

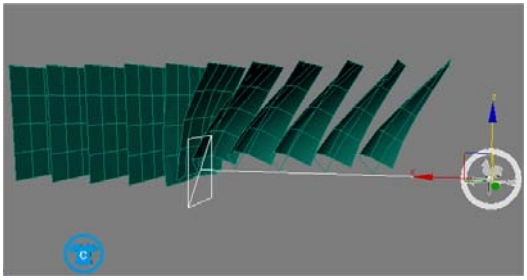
Support for Global Collisions

In reactor, you can enable or disable collisions globally instead of inside the *Rigid Body Collection (page 2–723)*. The software stores disabled collisions inside *the reactor Utility (page 2–805)*. You can disable collisions not only for rigid bodies, but also for cloth, soft and rope. And you can access disabled collisions through MAXScript.



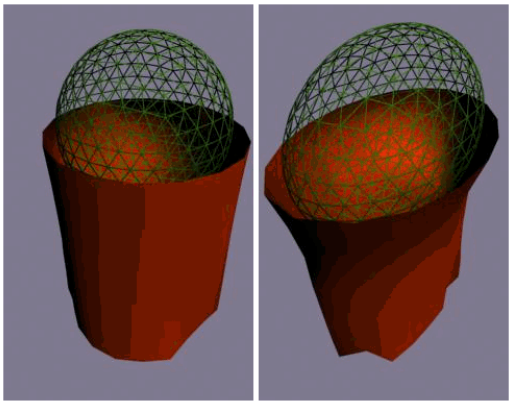
Animatable Wind

You can animate most parameters in the reactor *Wind (page 2–803)* object, including wind speed and direction. Range and Falloff parameters let you set up the range of action for the wind.



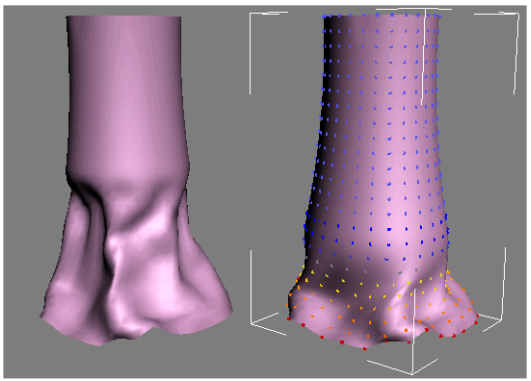
Cloth/Soft/Rope Attachments to Deforming Meshes (Skin)

The *Attach to DefMesh (page 2–799)* deformable constraint allows vertices in cloth, soft bodies, and rope to follow a non-rigid mesh such as skin.



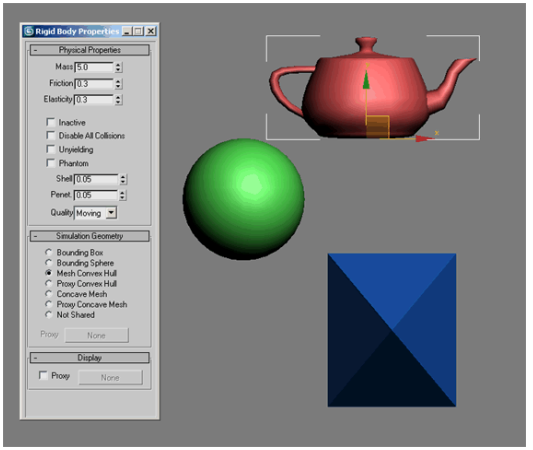
Soft Selection for Cloth/Soft/Rope

The *deformable (page 2–777)* (cloth/soft/rope) modifiers in reactor can now use the soft selection flowing to that modifier and blend the current vertex animation with the reactor animation, facilitating the transition between skin-driven and reactor-driven animation



Floating Rigid Body Property Editor

In reactor you can change the *rigid body properties (page 2–716)*, like mass, elasticity or friction without having to use the *reactor Utility (page 2–805)*. You can open a floating MAXScript window to modify rigid body properties from the reactor menus, quad menus, and toolbars at any time.



MAXScript Access

Virtually all parameters and functionality in reactor are accessible through MAXScript.

Animation Features

reactor can automatically create list controllers (or **character studio** layers) to store the animation of rigid bodies. For more information, see *Preview & Animation Rollout (page 2–806)*.

Also, the *ragdoll script (page 2–818)* is aware of being applied to **character studio** bipeds, and, in such cases, sets up the constraints so that, when you create an animation, the biped is physically animated.

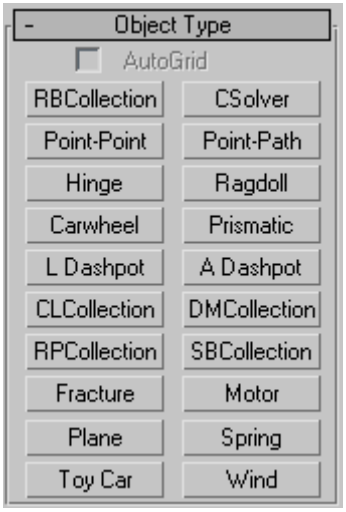
reactor Helpers

Create panel > Helpers > reactor

reactor menu > Create Object

This topic simply provides links to the various helper objects that are part of the reactor system.

Interface



Rigid Body Collection (page 2–723)

Constraint Solver (page 2–736)

Point-Point Constraint (page 2–750)

Point-Path Constraint (page 2–762)

Hinge Constraint (page 2–747)

Rag Doll Constraint (page 2–737)

Car-Wheel Constraint (page 2–757)

Prismatic Constraint (page 2–754)

Linear Dashpot (page 2–730)

Angular Dashpot (page 2–732)

Cloth Collection (page 2–781)

Deforming Mesh Collection (page 2–794)

Rope Collection (page 2–792)

Soft Body Collection (page 2–788)

Fracture (page 2–770)

Motor (page 2–765)

Plane (page 2–764)

Spring (page 2–727)

Toy Car (page 2–766)

Wind (page 2–803)

Rigid Bodies



The rigid body is the basic building block of reactor simulations. You can use a rigid body in reactor to simulate any real-world object that doesn't change its shape, from a pen to a boulder hurtling down a mountainside.

You can create a rigid body using any geometry in a 3ds Max scene. reactor then lets you assign the properties the object should have in the simulation, such as mass, friction, and whether the body can collide with other rigid bodies. You can also restrict the possible movement of your rigid bodies in the simulation using constraints such as hinges and springs.

This chapter shows you how to create and work with rigid bodies in reactor. It includes the following topics:

Rigid Body Basics (page 2–717)

Rigid Body Collection (page 2–723)

Constraints (page 2–724)

Plane (page 2–764)

Motor (page 2–765)

Toy Car (page 2–766)

Fracture (page 2–770)

Storing and Accessing Collisions (page 2–774)

To find out how to simulate objects that change their shape over time, see *Deformable Bodies (page 2–777)*.

Rigid Body Basics

You can make a rigid body from any geometry in your scene. A rigid body can be a single object, or it can comprise several objects grouped together, known as a *compound rigid body (page 2–722)*. If you designate an object whose geometry changes over time as a rigid body, the simulation uses its geometry at the start frame.

reactor lets you assign the physical properties that each body will have in the simulation, such as mass, friction, and whether the body can collide with other rigid bodies. You can specify a proxy geometry for a rigid body, which allows reactor to treat the rigid body as an easier-to-simulate shape for the purpose of the simulation. You can also specify how your rigid body should be displayed when previewing the simulation.

An object is simulated as a rigid body after you have added it to a *Rigid Body Collection (page 2–723)*; you can edit its rigid body properties before or after you do this.

This section tells you how to:

Edit rigid body *properties (page 2–717)*, including

- *physical properties (page 2–718)*
- *simulation geometry (page 2–719)*
- *display properties (page 2–722)*

Rigid Body Properties 

Utility panel > reactor > Properties

Menu bar > reactor > Open Property Editor

reactor toolbar > Open Property Editor button

You assign physical properties, simulation geometry, and display properties to a rigid body using the Rigid Body Properties rollout or dialog. The interface consists of three rollouts or sections, each of which is described in detail in this topic.

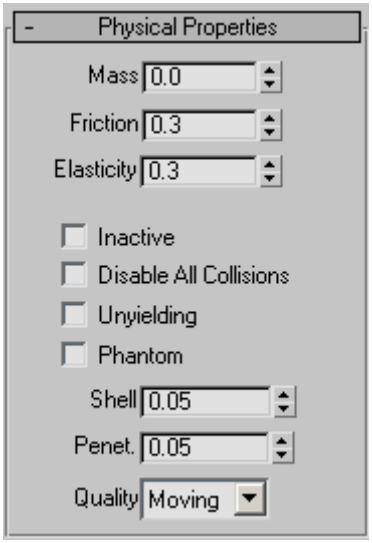
Procedures

To edit the properties of a rigid body:

1. Select the appropriate object in the scene.
2. Open the Properties rollout or dialog. If the item you have selected cannot become a rigid body (for example, if it is a helper object such as a Rigid Body Collection), or if you have no object selected, the dialog still displays when you choose this option, but its controls are unavailable until you select a valid object.
3. Use the settings to specify the required properties, as described in the following sections.

Interface

Physical Properties



Mass—A rigid body's mass governs how the object interacts with other objects. When its mass is set to 0.0 (the default value), the object remains fixed in space during the simulation, although other objects will be able to collide with it. For example, you could use a fixed rigid body to create a slope for other objects to roll down. Other values allow the object to move during the simulation, depending on other circumstances. A valid value for Mass is greater than or equal to 0.0.

Friction—The coefficient of friction for the object's surface. This affects how smoothly the rigid body moves relative to surfaces it's in contact with. The friction values for both objects combine to produce a coefficient for the interaction. To achieve realistic results, use values between 0.0 and 1.0. However, values up to 5.0 are accepted.

Elasticity—This value governs the effect collisions have on the velocities of the rigid body; in other words, how "bouncy" the object is. Like Friction, this is a pair-wise coefficient: When two objects collide, their elasticity values combine to produce a coefficient for the interaction. To achieve realistic results, use values between 0.0 and 1.0. However, values up to 5.0 are accepted.

Inactive—When on, the rigid body starts the simulation in an inactive state. This means it requires interaction with another object or system, or the mouse, before it becomes active in the simulation. For example, if you place an object in midair, give it a mass and set it to Inactive, when the simulation starts it sits in midair until something interacts with it. Inactive objects require less computation during simulation.

Disable All Collisions—When on, the object doesn't collide with other objects in the scene; it simply passes through them.

Unyielding—When on, the rigid body takes its motion from the animation it already has in 3ds Max, rather than the physical simulation. Other objects in the simulation can collide with it and react to its motion, but its motion is governed solely by the current animation in 3ds Max, and reactor will not create keyframes for it.

Phantom—A phantom object has no physical presence in the simulation. Like an object with Disable All Collisions on, it simply passes through other objects. Unlike an object with disabled collisions, however, a phantom maintains collision information about any objects that it passes through during the simulation. You can then use this collision information, for instance, to trigger sounds or other effects. You can find out how to access collision data in the *Storing and Accessing Collisions (page 2–774)*.

Shell—The radius of an extra "shell" around convex shapes, which reactor uses as the shape's surface for collision-detection purposes. The simulation tries to ensure that the distance between this shell and other objects is always more than zero; in other words, that the distance between the original

convex shape and other objects is always more than the combined radii of the objects. Default=0.05.

Important: Applies to Havok 3 only.

Adding a shell to an object can improve performance. The core convex-convex collision-detection algorithm is fast when objects are not interpenetrating, and slower when they are. Adding a shell makes it less likely that the shapes themselves will interpenetrate, thus reducing the likelihood of using the slower algorithm. Using a shell is thus faster in situations that involve the risk of shapes interpenetrating; for instance, when an object is settling or sliding on a surface, when there is a stack of objects, or when many objects are jostling together.

Penet.—(Penetration) The amount of penetration reactor permits.

To avoid trying to solve insoluble physical problems, the Havok 3 engine allows penetration between objects even if *continuous physics* is enabled. Default=0.05.

Important: Applies to Havok 3 only.

Quality—Lets you set individual settings for each object based on the desired level of interaction. Default=Moving.

Important: Applies to Havok 3 only.

The available Quality settings are:

- **Debris—**Low-importance objects used for adding visual quality.
- **Moving—**Regular objects that populate the world, such as furniture, etc.
- **Critical—**Essential objects that are never allowed to interpenetrate.
- **Bullet—**Fast-moving projectiles.

Simulation Geometry

Lets you specify the physical representation of your object that will be used in the Havok simulation.

Proxies

reactor supports substituting one object for another in two different ways: geometry proxies and display proxies.

Geometry proxies allow you to specify a different body's geometry as the simulation geometry for an object. For example, you can have a complex object displayed on-screen, but replaced for simulation purposes by a box, which is much easier and faster for reactor to simulate. The box governs the movement of the object, and dictates how it collides, its position, and orientation. However, on-screen you can see the complex object in full detail. Geometry proxies are applied per object.

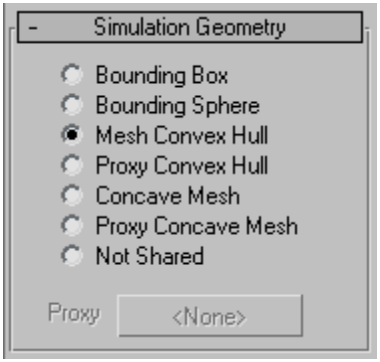
Display proxies replace a rigid body's display body with that of another object. Therefore, they affect display only during the *real-time preview (page 2–815)* and do not affect animation. They are applied per rigid body, rather than per object. This means that you can create a compound rigid body of several objects and simulate these, but display an alternative mesh for the body during the preview. You can find out how to specify a display proxy in the *Display (page 2–722)* section.

Convex and Concave Objects

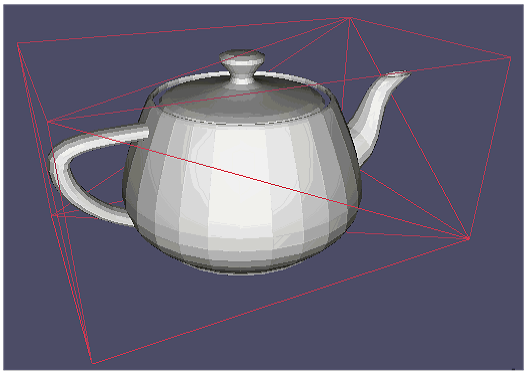
A rigid body primitive is defined as convex if, given any two points inside the object, you can always go in a straight line from one to the other without leaving the object. Convex objects include spheres, cylinders, and boxes. For example, a sphere is convex but a golf ball is concave because of the concavities (dimples) in its surface. Also, by definition, non-closed meshes (planes, hollow hemispheres) are always concave.

Convex objects are faster to simulate than concave objects. Because of this, you should aim to use convex objects as often as possible for simulation. Treating concave objects as convex for simulation purposes allows you to take advantage of their faster processing time. This is the default setting for an object's simulation geometry (Mesh Convex Hull). If you are unsure whether an object is convex or concave, you can perform a convexity test on it. With the object selected, open the reactor - Utilities menu and select Convexity Test.

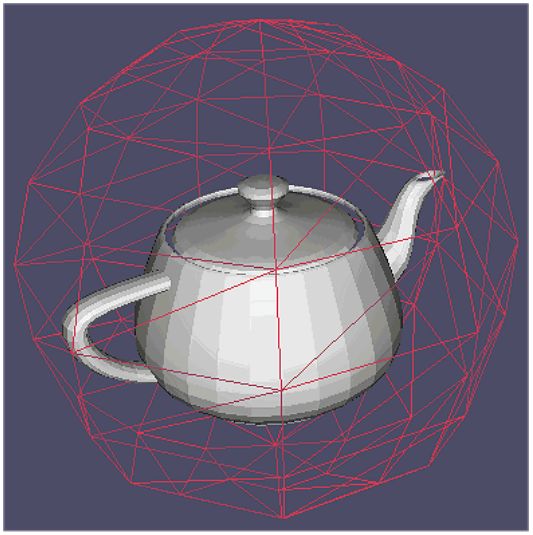
You can specify one of the following simulation geometry options to define how your object will be represented in the physical simulation. To view the simulation geometry for your objects in the preview, as in the following examples, select Sim Edges from the Display menu in the Preview Window.



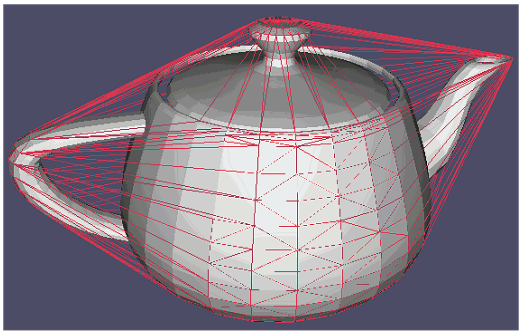
Bounding Box—The object is simulated as a box whose extents are determined by the object's dimensions.



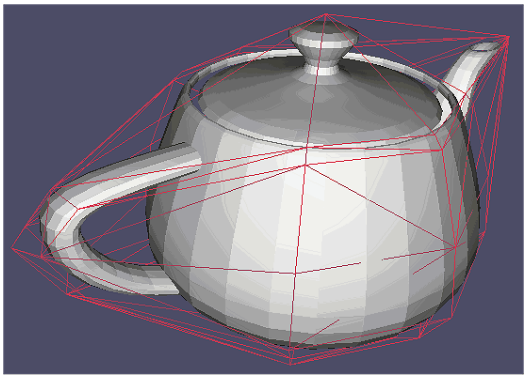
Bounding Sphere—The object is simulated as an implicit sphere. The sphere is centered on the object's pivot point and then minimally encloses the object's geometry.



Mesh Convex Hull—This is the default option. The object's geometry is passed through an algorithm that creates a convex geometry using the geometry's vertices, completely enclosing the geometry's vertices. To visualize this, imagine shrink-wrapping a teapot: The teapot is concave but its shrink-wrap forms a convex hull.

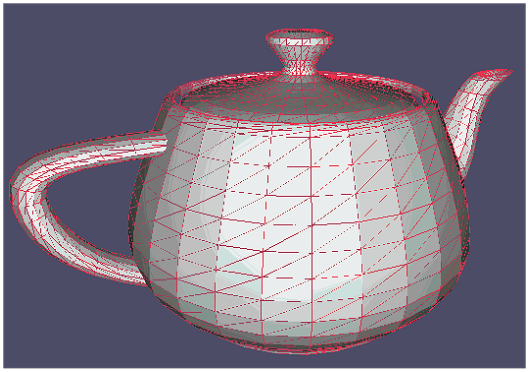


Proxy Convex Hull—The convex hull of another object is used as the physical representation of the object in the simulation. For instance, you could use the convex hull of a low-poly teapot to simulate a high-poly teapot. The proxy object's pivot point is aligned with that of the rigid body.



Concave Mesh—The actual mesh of the object is used for simulation. Although the convex hull of an object and the object's actual mesh may be exactly the same shape, using the convex hull simulates much more quickly, as reactor can make certain assumptions for convex objects. If you try to use Concave Mesh for a convex object you will get a warning. Not heeding such warnings could dramatically reduce the speed of your simulation. In some cases, though, you might want to ignore that warning. If, for example, you want to place objects *inside* a convex object, and make them collide with the internal faces of that object, you should simulate it as concave instead of convex.

Note: If you are using a standard 3ds Max plane as a rigid body (as opposed to the special reactor *Plane (page 2–764)* object), you must set its simulation geometry to Concave Mesh.



Proxy Concave Mesh—Another object's concave mesh is used as the physical representation of the object. For example, you can use a low-poly teapot to animate a highly tessellated teapot. The proxy object's pivot point is aligned with that of the rigid body.



Not Shared—This option is active only when multiple objects with different Simulation Geometry settings are selected, and cannot be chosen by the user.

Display

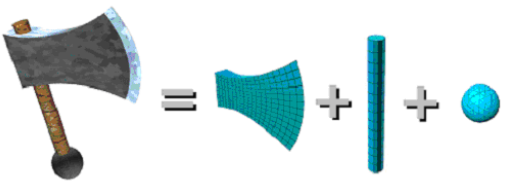


Use Display Proxy—When on and you specify a proxy object by clicking the Proxy button, reactor displays the specified proxy In the Preview Window in place of the object.

Compound Rigid Bodies

Rigid bodies can be made of one or more objects, or primitives. For instance, if you assign physical properties to an object in a scene and add it to a Rigid Body Collection, we say that the rigid body has one primitive. However, reactor can also simulate a rigid body made up of more than one primitive. To do this, you need to group objects together using the 3ds Max Group menu. You can then add the group to the collection, and the objects in the group then become the primitives that make up the rigid body.

Creating such compound rigid bodies can be useful, as it is much faster to simulate a rigid body made up of convex pieces than to simulate a more complex concave shape. In addition, the primitives that make up a compound rigid body can have different masses, so you can create rigid bodies that have nonuniform mass distribution. For example, you can create a hand axe where the head is heavier than the handle, so that if you toss the axe it spins around its head.



You can do this because the Mass property on the Rigid Body Properties dialog is assigned to a specific primitive rather than an entire rigid body, unlike most of the other properties in this dialog. For single-primitive rigid bodies, this distinction is unimportant. However, in the case of a compound rigid body, the group parent is the rigid body and cannot have a Mass value assigned: Mass can be assigned only to its children, the primitives.

Simulation geometry is also defined at a primitive level, so you can combine concave and convex hull geometries in a single rigid body.

Procedures

To create a compound rigid body:

1. Arrange the objects to use to make up your compound rigid body.
2. Select all of the objects.
3. From the Group menu, choose Group.

   Your objects are now all child objects of a new group parent. You can add this parent object to a Rigid Body Collection.

To edit compound rigid body properties:

1. With the compound rigid body selected, choose Group > Open.
2. To edit the compound rigid body's properties, select the parent (dummy) and go to the properties rollout of the utility or open the Rigid Body Properties dialog. You cannot edit the Mass or Simulation Geometry values for a group parent because they are primitive properties.

   Tip: It's easier to select a group parent using the *Select Objects dialog (page 1–78)* with Display Subtree enabled: The group parent is displayed above its members.
3. To edit one of the rigid body's primitive's properties, select the appropriate child and go to the Properties rollout of the utility or

open the Rigid Body Properties dialog. You can give this object a Mass value and, if you wish, change its Simulation Geometry setting. Specifying any other rigid body properties for a child object has no effect unless you remove the object from the group.

Rigid Body Collection 

Create panel > Helpers > reactor > RBCollection
Menu bar > reactor > Create Object > Rigid Body Collection
reactor toolbar > create Rigid Body Collection button

The Rigid Body Collection is a reactor helper object that acts as a container for rigid bodies. Once you have added a Rigid Body Collection to your scene, you can add any valid rigid bodies in the scene to the collection. You can find out more about rigid bodies in the *Rigid Body Basics (page 2–717)* section.

When you run the simulation, reactor examines the Rigid Body Collections in the scene and, provided the collections are not disabled, adds the rigid bodies they contain to the simulation.

At a lower level, a collection also allows you to choose the mathematical method to be used for solving the rigid body behaviors for its member bodies.

Procedures

To create and use a Rigid Body Collection:

1. Choose any of the above commands, and then click in any viewport to add the Rigid Body Collection.

   Note: The icon's position has no effect on the collection's behavior.

   The collection icon is added to the scene. You can add rigid bodies to the collection in two ways: by picking or by using a selection list.

2. Add objects to the collection using either (or both) of these methods:

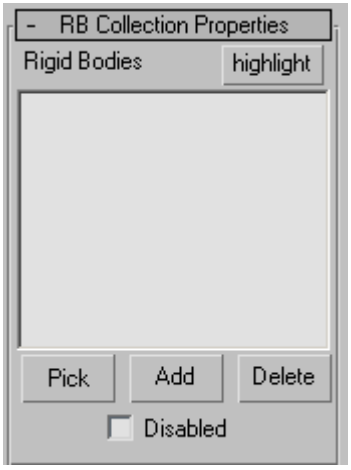
   - On the RB Collection Properties rollout, click the Pick button, and then in the viewport click an object to add.
   - On the RB Collection Properties rollout, click the Add button, and then use the Select Rigid Bodies dialog to specify one or more objects to add.

To create a collection and add bodies in a single step:

1. Create and select the object(s) to add to the collection.
2. Create the collection using the reactor toolbar, reactor menu, or quad menu. reactor instantly places the collection in the scene and adds the selected objects to the collection.

Interface

RB Collection Properties rollout



Highlight—Cause the objects in the Rigid Bodies list to display momentarily in the viewports as if selected.

Rigid Bodies—Lists the names of the objects in the Rigid Body Collection.

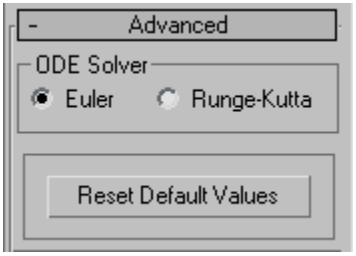
Pick—Lets you add an object to the Rigid Body Collection. Click this button, then in the viewport move the cursor over the object you would like to add to the collection. If the object can be used as a rigid body the cursor will change from an arrow to a cross and you can select the object to add it to the collection.

Add—Lets you add one or more objects from the scene to the collection. Click the button to open the Select Rigid Bodies dialog. Make a selection in the provided list, and then click the Select button to add the objects to the collection.

Delete—Lets you remove objects from the collection. In the rigid bodies list, select the bodies you would like to remove from the collection and click this button.

Disabled—When on, the collection and the bodies it contains are not added to the simulation.

Advanced rollout



ODE Solver—Lets you choose the method by which reactor simulates the collection:

- **Euler—**The collection calculates the behavior for its rigid bodies using an Euler ODE (Ordinary Differential Equation) solver. Euler is a fast method that provides good results in most cases.
- **Runge-Kutta—**This method is more accurate in some cases but requires more computation. Use Runge-Kutta if you have many object connected using *simple constraints (page 2–727)* like springs or dashpots, because those kind of systems can easily become unstable.

Reset Default Values—Sets values for the collection to their defaults. In this instance, it sets the ODE Solver back to its default value.

Constraints

With reactor, you can easily create a simple physical simulation by simply assigning rigid body properties to objects and adding them to a Rigid Body Collection. When you run the simulation, objects can fall from the sky, slide across each other, bounce off each other, and so on. However, let's say you want to simulate a real-world scene such as a person pushing open a door. How, for instance, do you make sure that the door rigid body doesn't just fall onto the ground, or that it swings out in the proper direction when pushed?

To accomplish this you use *constraints*. These let you restrict the possible movement of objects in the physical simulation. Depending on the type of constraint you use, you can hinge objects together, attach them together with springs that will snap back if the objects are pulled apart, or even simulate the movement of a human body joint. You can constrain objects to each other or to points in space.

This section shows you how to work with constraints in reactor. It includes the following sections:

Constraint Concepts (page 2–725)

Simple Constraints (page 2–727), which includes the following topics:

- *Spring (page 2–727)*
- *Linear Dashpot (page 2–730)*

- *Angular Dashpot (page 2–732)*

Cooperative Constraints (page 2–735), which includes the following topics:

- *Constraint Solver (page 2–736)*
- *Rag Doll Constraint (page 2–737)*
- *Hinge (page 2–747)*
- *Point to Point (page 2–750)*
- *Prismatic (page 2–754)*
- *Car-Wheel (page 2–757)*
- *Point to Path (page 2–762)*

Constraint Concepts

This section introduces some common concepts that you'll need to understand to work with any of the reactor constraint types.

Constraints, Simple Constraints and Cooperative Constraints

You use constraints in reactor to specify limitations in the movements of objects. Without constraints, the movement of objects could be limited only by collisions and deformations. This is the case in the real world: For example, a door's movement is limited by a set of pieces that form a hinge; the collisions between those objects limit the movement of the door. The motion of two objects attached by a spring is limited by the deformation of the spring. A train can move only along the path defined by a rail due to the collisions of its wheels with the rail. In many cases, though, it is preferable to specify explicitly the effect of those objects (hinge, spring, rail) rather than model them and simulate them. This is what constraints are for.

A constraint lets you limit the way an object can move. Once you specify a constraint, reactor tries to enforce it during the simulation. For example, you can use a *Hinge (page 2–747)* constraint to simulate the effect of an actual hinge on an object: No translation is allowed, and rotation is allowed around only one axis. Similarly, you can use a *Spring (page 2–727)* constraint to simulate the effect of a spring (translation is limited to a certain length); or a *Point-Path (page 2–762)* constraint to simulate the effect of a rail (translation and orientation are limited to follow a path).

Sometimes you will have a system of many objects constrained together. For example, if you want to simulate a character falling down the stairs, you might constrain the different bones of the character using many constraints (like *Rag Doll (page 2–737)* or *Hinge (page 2–747)*). Because all the bodies are connected, maintaining one constraint may affect the other constraints, so it is better if they are simulated together, so they are aware of each other. Thus, some constraints require you group them so they can be solved as a system. Those constraints are called *Cooperative Constraints (page 2–735)* and are usually more stable, although they can be slightly slower to simulate. The other constraints, *Simple Constrains (page 2–727)*, cannot be grouped and therefore are more prone to instability in complex scenes, but are faster to simulate.

Constraint Spaces

In rigid body dynamics, each body has six degrees of freedom to move:

- three translational degrees of freedom
- three rotational degrees of freedom

Each type of reactor constraint can remove or limit one or more of these degrees of freedom for its constrained bodies.

Depending on the number and type of these limitations, we get different types of constraint, from the simple *Point-Point (page 2–750)* constraint to the much more complicated *Rag Doll (page 2–737)* constraint. For example, with a Point-Point constraint, the constrained objects

are completely free to rotate around the constraint pivot point, but have no linear freedom relative to each other in any direction; they are attached together at the point. However, with a Rag Doll constraint, the objects not only have no linear freedom, but their possible relative orientation is also restricted.

In each constraint, these angular and linear limits are defined in terms of the constraint's coordinate system or *constraint space*. Because a constraint restricts the movement of its objects relative to each other, a constraint also needs to maintain a mapping from each object's local space to the constraint space. reactor lets you manipulate the constraint space in each object's local space separately, as you'll see in the *Working With Constraint Spaces (page 2–726)* section.

Parents and Children

Each reactor constraint can have two objects: a *parent* object and a *child* object. Although two-bodied constraints actually restrict the possible movement of both bodies, it is often simpler to specify how one object is allowed to move relative to the other one, particularly when you are using limits. With reactor, you specify how the child object can move relative to the parent.

For some constraints, it makes no difference which object is the parent and which is the child. For instance, if you use a Point-Point constraint to attach two objects together at a common point, it works in exactly the same way regardless of how you specify the attached objects.

However, with the more complex Rag Doll constraint, which you typically use to model body joints, it's important to specify which object is the parent. This is because this constraint lets you specify a number of limits on how the child can move relative to the parent object. For example, when you move your torso, your arm always moves with it. So when modeling a shoulder joint, you would specify that the torso object is the parent, making the arm object the child. You can then specify limits on the arm's movement relative to the torso.

When you constrain an object to a point in world space rather than another object, the constrained object is the child object and there is no parent object.

Working With Constraint Spaces

For most constraints, by default the constraint space is aligned with the child body. This means that the constraint pivot/attachment point is positioned at the child object's pivot point, and the constraint space takes its orientation from the child body's local space. The exceptions to this are the constraints with two attachment points, Springs and Stiff Springs, and the Point-Path constraint, each of which has its own default alignment.

You can move the constraint space relative to each of the bodies by going to the constraint's Parent Space or Child Space sub-object level in the modifier stack, and then using the Move and/or Rotate tools. The space will then maintain its position and/or orientation relative to the corresponding object during the simulation.

reactor also provides some quick shortcuts for aligning constraint spaces, available with each constraint. These are:

- **Align Spaces To Child Body—**This is the default alignment for all constraints except Spring, the Stiff Spring variant of Point-Point, and Point-Path.
- **Align Spaces To Parent Body—**The constraint space is aligned with the parent body: The pivot point is positioned at the parent body's pivot point, and the constraint space takes its orientation from the parent object's local space.
- **Align Spaces To Child Space—**Aligns the parent and child constraint spaces with the child's

constraint space. So, for instance, if you move the constraint space relative to the child, you can use this option to move the parent's constraint space into alignment with it.

- **Align Spaces To Parent Space—**Aligns the parent and child constraint spaces with the parent's constraint space. So, for instance, if you move the constraint space relative to the parent, you can use this option to move the child's constraint space into alignment with it.

Certain constraints have additional alignment options:

- **Align Spaces To Each Body—**(Spring and Point-Point > Stiff Spring only) This option aligns each local constraint space (and hence the spring attachment point) with the relevant body, so each body's attachment point is at its pivot. It also sets the spring length to the current distance between the nodes. This is the default alignment for springs and stiff springs.
- **Align Parent Space To Path—**(Point-Path constraint only) This option aligns the parent space with the local space of the shape used for the constraint's path. This is the default parent space alignment for the Point-Path constraint.

Simple Constraints

The simple constraints provided with reactor are:

- *Spring (page 2–727)*
- *Linear Dashpot (page 2–730)*
- *Angular Dashpot (page 2–732)*

Unlike cooperative constraints, the simple constraint does not require you to add it to a *Constraint Solver (page 2–736)* helper. Instead, reactor adds all valid simple constraints in a scene to the simulation by default. A two-bodied simple constraint is valid if it has two rigid bodies attached, while a single-bodied simple constraint is valid if it has one rigid body attached.

Spring 

Create panel > Helpers > reactor > Spring

Menu bar > reactor > Create Object > Spring

reactor toolbar > Create Spring button

The Spring helper lets you create a spring-like effect between two rigid bodies in the simulation, or between a rigid body and a point in space. During the simulation, the spring exerts forces on the attached bodies in an attempt to maintain its rest length. So, for instance, if the objects are pulled apart so that the spring attachment points are further apart than the rest length, the spring works to bring them back together again.

You can configure a spring's behavior by specifying its stiffness, damping, and rest length. reactor also lets you choose whether the spring acts under extension (when the attachment points are pulled further apart) or compression (when the attachment points are pushed closer together) or both (the default). A spring that acts only on extension behaves as if the objects were joined with a rubber band.

You do not have to add a Spring constraint explicitly to a simulation, as all valid springs in a scene are added to the simulation by default. A two-bodied spring is valid if it has two rigid bodies attached, while a single-bodied spring is valid if it has one rigid body attached. When not selected, invalid springs will be colored red in the viewport.

Procedures

To create a spring:

- Choose one of the above options, and then click in any viewport to add the Spring.

Note: The icon's position has no effect on the constraint's behavior.

To attach objects to the spring:

1. Create the Spring and the object(s) to connect using it.
2. If you do not want to assign a second body to the spring, turn off the Parent check box on the spring's Properties rollout if necessary.
3. Click the Child pick button and then select the object to use as the child object in one of the viewports. This attaches the Spring to the body at the body's pivot point.
4. If the spring is two bodied, repeat step 2, using the Parent pick button to set the parent object for the spring.

To create the spring and attach the objects in a single step:

- See the *Scripts (page 2–817)* topic.

To change the attachment positions for the bodies:

1. With the Spring constraint selected, open the Modify panel.
2. In the modifier stack, open the Spring's sub-object list.
3. Highlight the option corresponding to the object whose attachment point you want to move: Parent Space or Child Space. If the Spring is single-bodied, Parent Space allows you to move the world attachment position.
4. The appropriate space becomes active in the viewport and can be moved using the Move tool. The attachment point maintains its position relative to its corresponding object; for example, if you move the child object, its space, and hence its attachment point, moves with it. The attachment point doesn't actually have to be on the object.

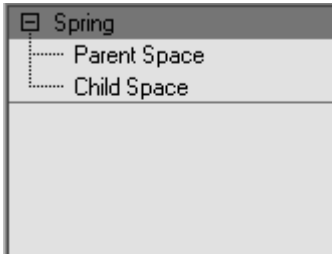
   If you move the attachment points so that they're inside or beyond the Spring rest length, you will see the attachment points represented by large box markers and the rest length denoted by small box markers. The difference between the two is displayed in red.
5. To reset the attachment point to the object's pivot, reattach the object to the Spring. Alternatively, if you click the Align Spaces To Each Body button, both attachment points will be reset to their respective body's pivot.

To align the parent and child constraint spaces:

1. Ensure that the constraint is selected in one of the viewports.
2. In the Align Spaces To group, click the button for your chosen alignment. You can find out more about the alignment options in the *Working With Constraint Spaces (page 2–726)* section.

Interface

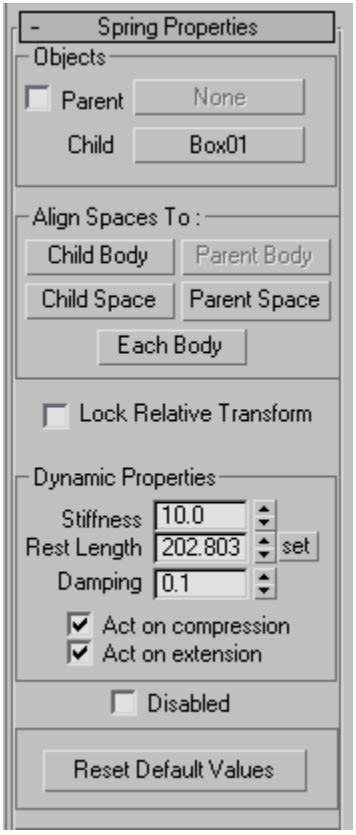
Modifier Stack



Child Space—At this sub-object level, you can select and move the constraint space for the child object, including the point relative to the child where the Spring is attached to it. When you assign a child body to the Spring, the Child Space attachment point is set to the child's pivot point. If you alter the position of the attachment point, it maintains its new position relative to the child during the simulation.

Parent Space—At this sub-object level, you can select and move the constraint's second attachment point. If the spring is two-bodied, then this is the point relative to the parent where the spring is attached to it. When you assign a parent object to the Spring, the Parent Space attachment point is set to the parent object's pivot point. If you alter the position of the attachment point, it maintains its new position relative to the parent. If the spring is single-bodied, then the Parent Space's attachment point represents the point in the world where the spring is attached.

Properties rollout



Parent—When on, the Spring has two bodies and you can designate an object to be the parent body. When off, the Spring is single-bodied and the point occupied by the Parent Space sub-object is the parent.

You assign the parent object by clicking this button and then selecting a rigid body from any viewport. Thereafter the button displays the name of the parent body.

Child—Displays the name of the second rigid body attached to the Spring. You assign the child object by clicking this button and then selecting a rigid body from any viewport.

Align Spaces To—Use these options to align the bodies' local constraint spaces. You can find out more about each option in the *Working With Constraint Spaces (page 2–726)* section.

Lock Relative Transform—When on, the relative transform between the child and parent constraint spaces is locked: If you move either space in the viewport, the other space moves along with it.

Stiffness—The strength of the spring; governs the force the spring applies to each attached body for each unit difference between its current length and its rest length. Values greater than or equal to 0.0 are valid. Default=10.0.

Rest Length—The spring applies forces to the attached objects in an attempt to reach this length. You can specify a figure in world units for this value, or you can use the Set button. The spring's rest length is displayed in the viewports as the length of the displayed line between the two small box markers. The large box markers are the objects' attachment points.

Set—Click this button to set the spring's rest length to the current distance between the two objects' attachment points.

Damping—Determines how quickly the oscillation of the spring settles down. It governs the forces applied to the connected bodies due to the relative velocities of their connection points. Valid

damping values are those greater than or equal to 0.0. The default damping value is 1.0.

Tip: Typically, a Damping value of 1/10th of the stiffness yields good results.

Act on compression—When on, the spring applies forces to the attached objects when its length is less than its Rest Length. Default=on.

Act on extension—When on, the spring applies forces to the attached objects when its length is greater than its Rest Length.

Disabled—When on, the spring is not used in the simulation.

Reset Default Values—Click this to set Stiffness, Rest Length, Damping, Act On Extension, and Act On Compression to their respective default values.

Linear Dashpot 

Create panel > Helpers > reactor > L Dashpot

Menu bar > reactor > Create Object > Linear Dashpot

reactor toolbar > Create Linear Dashpot button

Dashpots typically server to cushion impact. An example of a real-world linear dashpot is the hydraulic cylinder in an automobile shock absorber. The Linear Dashpot constraint lets you constrain two rigid bodies together in the simulation, or to constrain one body to a position in world space. It behaves like a heavily damped spring with zero rest length. You can specify the strength and damping, and whether collisions between the attached bodies are disabled.

reactor lets you specify a dashpot attachment point in each body's local space. During the simulation, the dashpot exerts impulses on the attached bodies in an attempt to make these points match up in world space, thus keeping the bodies in the same positions relative to each other. The bodies are still free to rotate around the attachment point.

You do not need to add a Linear Dashpot explicitly to a simulation, as all valid constraints in a scene are added to the simulation by default. A two-bodied dashpot is valid if it has two rigid bodies attached, while a single-bodied dashpot is valid if it has one rigid body attached. When not selected, invalid dashpots are colored red in the viewport.

Procedures

To create a Linear Dashpot:

- Choose one of the above options, and then click in any viewport to add the dashpot.

  Note: The icon's position has no effect on the constraint's behavior.

To attach objects to the Linear Dashpot:

1. Create the Linear Dashpot and the objects to connect using it.
2. On the linear dashpot's Properties rollout, click the Child pick button, and then select the object to use as the child object in one of the viewports. By default, this attaches the linear dashpot to the body at the body's pivot point.
3. If you do not want to assign a second body to the dashpot, turn off the Parent check box, if necessary.
4. If the dashpot is two-bodied, repeat step 2, using the Parent pick button to set Parent for the Linear Dashpot.

To create the constraint and attach objects in a single step:

- See the *Scripts (page 2–817)* topic.

To change the attachment positions for the bodies:

1. With the dashpot selected, open the Modify panel.

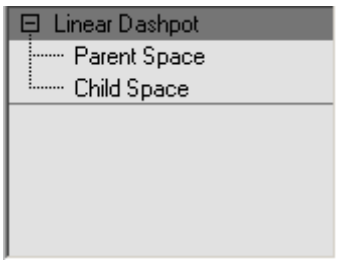
2. In the modifier stack open the dashpot's sub-object list.
3. Highlight the sub-object corresponding to the object in whose space you want to move the attachment point: Parent Space or Child Space. If the dashpot is single-bodied, Parent Space allows you to move the world attachment position.
4. The appropriate point becomes active in the viewport and can be moved using the Move tool. The attachment point maintains its position relative to its corresponding object; for example, if you move the child object, its space, and hence its attachment point, moves with it. The attachment point doesn't have to actually be on either object.
5. To reset the attachment point to the child object's pivot, reattach the objects to the dashpot, or use the Align Spaces To Child Body option.

To align the parent and child constraint spaces:

1. Ensure that the constraint is selected in one of the viewports.
2. In the Align Spaces To group, click the button for your chosen alignment. You can find out more about the alignment options in the *Working With Constraint Spaces (page 2–726)* section.

Interface

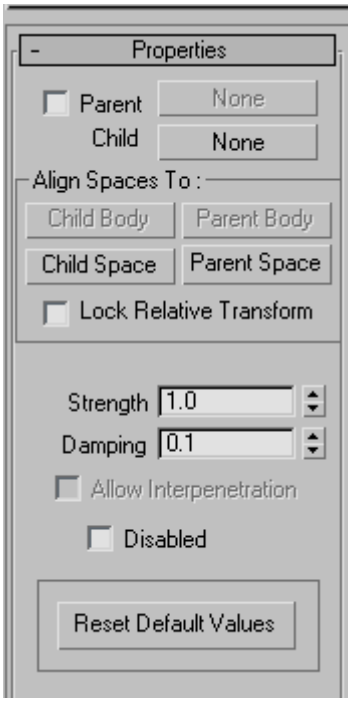
Modifier Stack



Child Space—At this sub-object level, you can select and move the dashpot attachment point for the child body. When you assign a child body to the dashpot, the Child Space attachment point is set to the child object's pivot point. If you alter the position of the Child Space attachment point, it maintains its new position relative to the child during the simulation.

Parent Space—At this sub object level, you can select and move the dashpot attachment point for the parent body. When you assign a parent object to the dashpot, the Parent Space attachment point is set to the parent object's pivot point. If you alter the position of the Parent Space attachment point, it maintains its new position relative to the parent during the simulation. If the dashpot is single-bodied, then the Parent Space's attachment point represents the point in the world where the dashpot will be attached.

Properties rollout



Parent—When on, the dashpot has two bodies and you can designate an object to be the parent body. When off, the dashpot is single-bodied and the point occupied by the Parent Space sub-object is the parent.

You assign the parent object by clicking this button and then selecting a rigid body from any viewport. Thereafter the button displays the name of the parent body.

Child—Displays the name of the second rigid body attached to the dashpot. You assign the child object by clicking this button and then selecting a rigid body from any viewport.

Align Spaces To—Use these options to align the bodies' local constraint spaces. You can find out more about each option in the *Working With Constraint Spaces (page 2–726)* section.

Lock Relative Transform—When on, the relative transform between the child and parent constraint spaces is locked: If you move either space in the viewport, the other space moves along with it.

Strength—Governs the size of the impulse the dashpot applies to each attached body, taking into account the distance between the attachment points. The strength value is mass-dependent. For example, a Strength value of 10 generates a different behavior when connecting two 5 kg bodies than when connecting two 50 kg bodies. Values greater than or equal to 0 are valid for this property. Default=1.

Damping—Determines how quickly the oscillation of the linear dashpot settles down. Damping governs the impulse applied to the connected bodies due to the relative velocities of their dashpot connection points. Valid damping values are those greater than or equal to 0 and less than or equal to 100000. Default=0.1.

Tip: Typically, a damping value of 1/10th of the stiffness yields good results.

Allow Interpenetration—When on, collisions are disabled between the dashpot's objects, so that they can pass through each other during the simulation. Defaults=off.

Disabled—When on, the dashpot is not added to the simulation.

Reset Default Values—Click this to set Strength, Damping, and Disabled to their respective default values.

Angular Dashpot 

Create panel > Helpers > reactor > A Dashpot

Menu bar > reactor > Create Object > Angular Dashpot

reactor toolbar > Create Angular Dashpot button

Dashpots typically serve to cushion impact. An example of a real-world angular dashpot is a device connected to a door to keep it from slamming shut. You can use the Angular Dashpot constraint to constrain the relative orientation of two rigid bodies, or the absolute orientation of a rigid body in world space. When simulating, the dashpot exerts angular impulses on the bodies to which it is attached in an attempt to maintain the specified rotation between the objects. You can specify the dashpot's strength and damping, and whether collisions between the system's bodies are disabled.

An angular dashpot has two sets of axes as sub-objects. For two-bodied dashpots, these are specified as offset rotations from the dashpot's bodies. For single-bodied dashpots, one sub-object is an offset rotation and the other is a world rotation. In a simulation, the dashpot tries to maintain a common rotation for these axes.

You do not need to add angular dashpots explicitly to a simulation, as all valid dashpots in a scene are added to the simulation by default. A dashpot is valid if it has the correct number of rigid bodies

attached and is enabled. When not selected, invalid dashpots are colored red in the viewport.

Procedures

To create an Angular Dashpot:

- Choose one of the above options, and then click in any viewport to add the dashpot.

  Note: The icon's position has no effect on the constraint's behavior.

To attach objects to the Angular Dashpot:

1. Create the dashpot and the objects to connect using it.
2. In the dashpot's Properties rollout, click the Child pick button and then select the object to use as the child object in one of the viewports. By default, this aligns the Angular Dashpot's target orientation with the child body's local space.
3. If you do not want to assign a second body to the dashpot, turn off the Parent check box if necessary.
4. If the dashpot is two-bodied, repeat Step 2, using the Parent pick button to set Parent for the angular dashpot.

To create the constraint and attach objects in a single step:

- See the *Scripts (page 2–817)* topic.

To change the offset rotations for the bodies:

1. With the Angular Dashpot selected, open the Modify panel.
2. In the modifier stack, open the dashpot's sub-object list.
3. Highlight the sub-object that corresponds to the offset rotation you want to change: Child Space or Parent Space. If the dashpot is single-bodied, Parent Space represents the world rotation that the dashpot works to align the Child Space with.
4. The selected axis sub-object becomes active in the viewports and can be rotated using the Rotate tool. The axis maintains its offset rotation so that if you rotate the object it's linked to, it also rotates.
5. To realign the rotation to the child body's local space, reattach the object to the dashpot. Alternatively, click the Align Spaces To Child Body button.

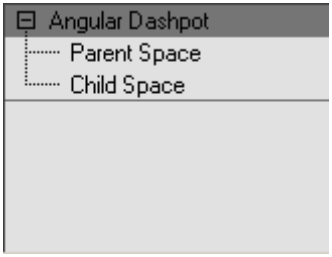
To align the parent and child constraint spaces:

1. Ensure that the constraint is selected in one of the viewports.
2. In the Align Spaces To group, click the button for your chosen alignment.

   You can find out more about the alignment options in the *Working With Constraint Spaces (page 2–726)* section.

Interface

Modifier Stack



Child Space—At this sub-object level, you can select and rotate the offset rotation for the child object. This is the rotation that the angular dashpot tries to maintain for the child body relative to the parent. When you assign a child object to the dashpot, reactor aligns Child Space to the child body's local space. If you alter the rotation of this sub-object, it maintains its new rotation relative to the child body.

Parent Space—At this sub-object level, you can select and rotate the offset rotation for the parent body. This is the rotation that the dashpot tries to maintain for the parent body relative to the child. When you assign the parent body to the dashpot, this rotation is aligned with the child body's local space. If you alter the rotation of this sub-object, it maintains its new rotation relative to the parent. If the dashpot is single-bodied, this rotation represents the world rotation to which the dashpot tries to align the Child Space axis.

Properties rollout



Parent—When on, the dashpot has two bodies and you can designate an object to be the parent body. When off, the dashpot is single-bodied and the point occupied by the Parent Space sub-object is the parent.

You assign the parent object by clicking this button and then selecting a rigid body from any viewport. Thereafter the button displays the name of the parent body.

Child—Displays the name of the second rigid body attached to the dashpot. You assign the child object by clicking this button and then selecting a rigid body from any viewport.

Align Spaces To—Use these options to align the bodies' local constraint spaces. You can find out more about each option in the *Working With Constraint Spaces (page 2–726)* section.

Lock Relative Transform—When on, the relative transform between the child and parent constraint spaces is locked: If you move either space in the viewport, the other space moves along with it.

Strength—Governs the size of the impulse the dashpot will apply to each attached body, taking into account the difference in their relative offset rotations (Child Space and Parent Space). The strength value is mass dependent – for example, a strength value of 10 will generate a different behavior when connecting two 5 kg bodies than when connecting two 50 kg bodies. Values greater than or equal to 0.0 are valid for this property. Default=1.0.

Damping—This value affects how quickly the oscillation of the angular dashpot settles down. It governs the impulse applied to the connected bodies due to their relative angular velocities. Valid damping values are those greater than or equal to 0. Default=0.1.

Tip: Typically, a damping value of 1/10th of the stiffness yields good results.

Allow Interpenetration—When on, collisions are disabled between the dashpot's objects: During the simulation, they can pass through each other. Default=off

Disabled—When on, the dashpot is not added to the simulation.

Reset Default Values—Click this to set Strength, Damping, and Disabled to their respective default values.

Cooperative Constraints

Cooperative constraints are more stable than *simple constraints (page 2–727)*, but can be more computationally expensive to simulate. For a cooperative constraint to take part in the simulation, you first need to add it to a *Constraint Solver (page 2–736)*. The solver acts as a container for the cooperative constraint in a particular Rigid Body Collection, and performs all the calculations necessary for the constraints to work together.

The following constraint types are cooperative constraints:

- *Hinge (page 2–747)*
- *Point-Point constraint (page 2–750)*
- *Point-Path constraint (page 2–762)*
- *Car-Wheel constraint (page 2–757)*
- *Prismatic constraint (page 2–754)*
- *Rag Doll constraint (page 2–737)*

Common Properties

The following properties are shared by all cooperative constraints:

Strength and Tau

These properties govern the impulses applied to the constraint's bodies in order to maintain the constraint, and so how strongly the constraint works to restrict their movement.

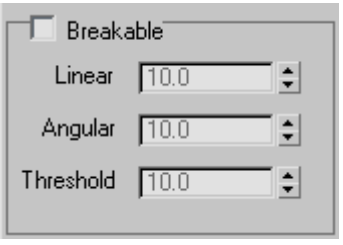
Strength—The constraint applies an impulse to each attached body in order to maintain the constraint, based on the relative velocities of the attached bodies. The Strength parameter governs the percentage of this impulse the constraint will apply to each object. If this is sufficient, then the constraint is maintained.

Tau—If the constraint drifts; that is, if the attached objects reach a state where the constraint isn't satisfied, then reactor applies a corrective force to rectify this drift. An impulse is calculated to rectify the drift in a constraint system; the Tau parameter governs the percentage of the corrective impulse the constraint applies to each object.

Note: Tau is unavailable in the Havok 3 engine.

Breakable Constraints

A breakable constraint stops working when a specified threshold is exceeded. For example, you could use a breakable constraint to make a door that flies off its hinges when kicked hard enough. You make a constraint breakable by turning on its Breakable check box. It then ceases to exert impulses on its constrained bodies if its limits are exceeded during the simulation.



Breakable—Turn on to make the constraint breakable.

The Linear and Angular parameters apply only to Havok 1, and the Theshold parameter applies only to Havok 3.

Linear—The maximum linear impulse the breakable constraint can apply before it breaks. Available in Havok 1 only.

Angular—The maximum angular impulse the breakable constraint can apply before it breaks. Available in Havok 1 only.

Threshold—A generic impulse parameter that governs the overall breakability of the constraint. Larger values make the constraint harder to break. Available in Havok 3 only.

Constraint Solver 

Create panel > Helpers > reactor > CSolver

Menu bar > reactor > Create Object > Constraint Solver

reactor toolbar > Create Constraint Solver button

The Constraint Solver acts as a container for the *cooperative constraints (page 2–735)* in a particular *Rigid Body Collection (page 2–723)*, and performs all the calculations necessary for the constraints to work together.

For a cooperative constraint to be simulated in a scene, it must be included in a valid Constraint Solver and any of the rigid bodies it contains should be in the Rigid Body Collection associated with the solver. For a Constraint Solver to be valid it should be associated with a valid Rigid Body Collection. When not selected, an invalid Constraint Solver is red in the viewport.

You can disable an entire Constraint Solver, thus disabling all of its constraints. You can also set a deactivation threshold for all the constraints in the solver, which lets you specify when you would like all the relevant rigid bodies to be deactivated.

Procedures

To create a Constraint Solver:

- Choose one of the above options, and then click in any viewport to add the Constraint Solver.

  Note: The icon's position has no effect on the constraint's behavior.

To assign a Rigid Body Collection to the solver:

1. On the Properties rollout click the RB Collection button.

   This enters a pick mode in which you can choose only a Rigid Body Collection.

2. In one of the viewports click the collection to assign to the Constraint Solver. The collection should contain the rigid bodies for the constraints in the solver.

To add constraints to a solver:

1. Create a Constraint Solver and some cooperative constraints to add to the solver.
2. Add constraints to the solver using either (or both) of these methods:
   - On the Constraint Solver > Properties rollout, click the Pick button, and then in the viewport click a constraint helper to add.
   - On the Properties rollout, click the Add button, and then use the dialog to specify one or more constraints to add.

Interface

Properties rollout



RB Collection button—Displays the name of the Rigid Body Collection assigned to the solver, if any. To assign a Rigid Body Collection to the solver, click this button and then selecting a collection in any viewport.

Highlight—Click this button to cause the constraints in the Constraints list to momentarily display in the viewports as if selected.

Constraints—Lists the names of the constraints currently in the solver.

Pick—Lets you add a constraint to the solver. Click this button, then in the viewport move the cursor over the constraint you would like to add to the solver. If the object can be used as a constraint, the cursor changes from an arrow to a cross and you can select the object to add it to the solver.

Add—Lets you add one or more constraints from the scene to the solver. Click the button to open up the Select New Constraints To Add dialog. Make a selection in the provided list, and then add the constraints to the solver by pressing the Select button.

Delete—Lets you remove constraints from the solver. Highlight the constraints to remove from the solver in the Constraints list and then click this button.

Deactivation Threshold—Lets you specify a special deactivation threshold for all the rigid bodies constrained by the solver's constraints. Objects in a Constraint Solver are activated and deactivated together; this parameter specifies how aggressive the deactivation is. A value of 0.0 disables deactivation. For more information about deactivation, see *Havok 1 World/Havok 3 World Rollout (page 2–808)*.

Disabled—When on, disables all the constraints in this solver.

Reset Default Values—Sets Deactivation Threshold and Disabled to their respective default values.

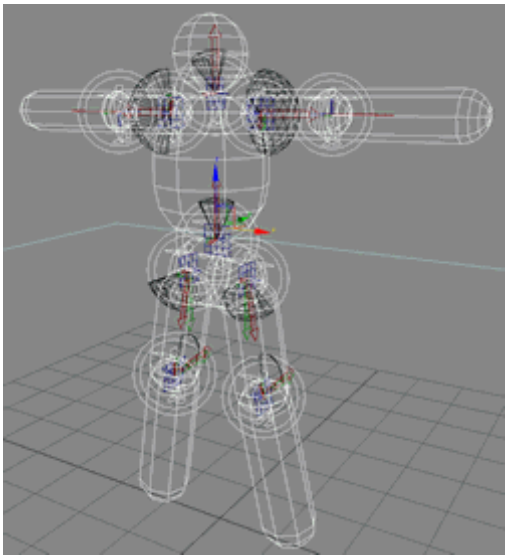
Rag Doll Constraint 

Create panel > Helpers > reactor > Ragdoll

Menu bar > reactor > Create Object > Rag Doll Constraint

reactor toolbar > Create Rag Doll Constraint button

The Rag Doll constraint lets you realistically simulate the behavior of body joints, such as hips, shoulders, and ankles. Once you decide the degree of movement a joint should have, you can model it by specifying limiting values for the Rag Doll constraint.



The Rag Doll constraint uses the parent body as a reference when defining the limits for the motion of the child body. For example, when your torso moves, your arm always moves with it. So when modeling a shoulder joint, you would usually designate the torso as the parent body and the upper arm as the child body. You can then specify limits on the arm's movement relative to the torso, as in the above illustration: The arm is allowed to rotate relative to the torso within the grey limited cone. You can also limit the child body's ability to twist.

You restrict how a Rag Doll constraint's child can move relative to the parent using the constraint's *Twist, Cone, and Plane limits (page 2–742)*. As with the other constraints, the constraint space for the Rag Doll joint is defined in each body's local space; you use this constraint space to define your limits. The constraint space is defined as follows:

- The origin as the attachment point of the constraint between the child and parent bodies.
- A twist axis, which is the axis around which the child body can twist relative to the parent body. You use this axis to specify twist limits and the child body's permitted cone of movement.
- A plane and plane normal centered on the attachment point; the plane serves to define further limits that restrict the child body's movement within the cone.

reactor provides special limit visualizations to help you choose appropriate values for the Rag Doll.

For a detailed explanation of how to define those limits, see *Defining Rag Doll Limits (page 2–743)*.

reactor simulates a Rag Doll constraint if it has the correct number of rigid bodies attached and is included in a valid *Constraint Solver (page 2–736)*. When not selected, invalid rag dolls are red in the viewport.

To find out how to create an example character quickly using the Rag Doll constraint, see *Scripts (page 2–817)*.

Procedures

To create a Rag Doll constraint:

- Choose one of the above options, and then click in any viewport to add the Rag Doll.

  Note: For a two-bodied Rag Doll, the icon's initial position has no effect on the Rag Doll's behavior. For a single-bodied Rag Doll, it defines the initial position of the world attachment point. You can change this afterwards.

To attach objects to the Rag Doll:

1. Create the rag doll and the objects to connect using it.
2. In the rag doll's Properties rollout, click the Child pick button. Now, in one of the viewports, select the object to use as the child object. This attaches the constraint to the selected body at the body's pivot point. The twist axis for the constraint in the child's space is aligned to the x-axis of the child object and initially, the plane normal will be aligned to the y-axis.

This constraint is more intuitive to use if the child and parent bodies are attached correctly. For example, if connecting an upper arm to a shoulder, the shoulder should be the parent and the upper arm the child.

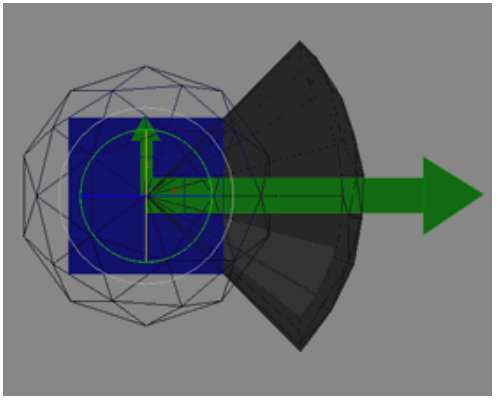
3. If you do not want to assign a parent to the constraint, turn off the Parent check box.
4. If the constraint is two bodied, repeat Step 2, using the Parent pick button to set the parent body for the constraint. By default, the constraint space in the parent's local space is also aligned with the child body.

To create the constraint and attach objects in a single step:

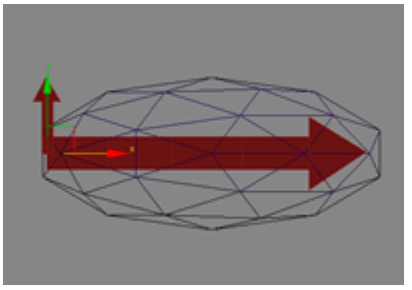
- See the *Scripts (page 2–817)* topic.

To change the Rag Doll positions and orientations for the bodies:

1. With the Rag Doll selected, open the Modify panel.
2. In the modifier stack open the sub-object list.
3. Access the sub-object level corresponding to the object whose constraint representation is to be altered: Parent Space or Child Space. If the rag doll is single-bodied, Parent Space represents the world attachment position and rotation of the rag doll's axes.



4. The space, represented by just the twist axis for the child and by the twist axis, plane normal and limit visualization for the parent, becomes active in the viewport and can be moved and rotated using the Move and Rotate tools respectively. The space maintains its position and rotation relative to its corresponding object. For example, if you move or rotate the child, Child Space moves relative to it.



5. To reset the space for an object to its default, reattach the object to the rag doll or use the Align To Child Body button.
6. You can also rotate the plane normal independent of the twist axis. To do this, set the sub-object level to Plane Rotation and you will be able to rotate the plane normal with respect to the parent's twist axis.

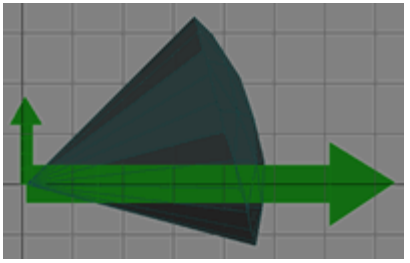
To change the Cone and Plane limits:

1. Ensure that you can see the limit visualization for the Rag Doll; this helps you to choose appropriate limits for your desired behavior. Visual representation of the limits is displayed at the location of the parent's sub-object. This provides you with a graphical representation of the limited area of movement for the child relative to the parent.
2. With the rag doll selected, open the Properties rollout.

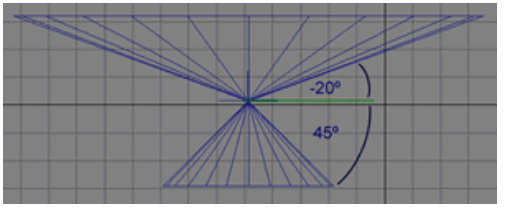
   Note: You can use this rollout to specify how the limits are represented in the visualization, as well as changing the limits.

3. Change the Cone > Min and Max settings.

   When displayed, the cone remains symmetrical and is rotated about the plane normal by the average offset of Min and Max. For example, if Min=-10.0 and Max=20.0, the average offset is (-10+20)*0.5=5 with a cone angle of 15. In effect, we have a 15-degree symmetrical cone [-15, 15] rotated by 5 degrees, resulting in limits of [–10, 20]. If the plane normal is not at 90 degrees to the twist axis, non-intuitive cone rotations can result. The rotation of the twist axis for the attached object is limited to the volume of this cone. You can view the cone on its own by ensuring that only the Show Cone option in the Properties rollout > Display group is on.



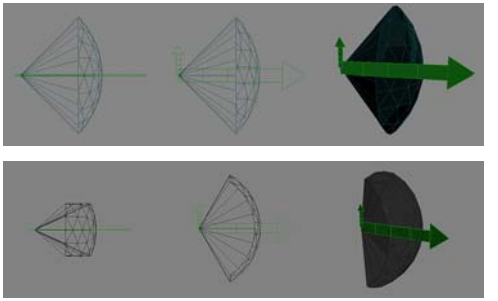
4. You can also try changing the Plane > Min and Max settings. These are limited between [-90, 0] and [0,90] respectively. Taking the Min and Max limits and sweeping those angles around the plane normal creates cones. The rotation of the attached object is then limited to within the region between the plane and these cones.



   You can also rotate the plane independent of the twist axis. To display the plane limits, turn on Display rollout Show Plane Limits.

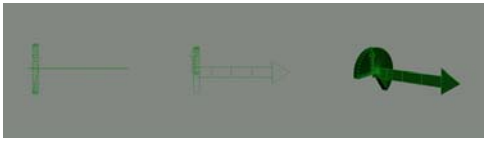
5. By combining Cone and Plane limits you generate valid rotational regions for the child object relative to the parent object. This is the region represented by the Show Volume option.

   Note: Cone and plane limits limit the movement independently. As a result, it is possible to have ineffective plane limits, where the cones generated by the plane limits do not intersect the cone generated by the cone limits. Use the plane limits to reduce the level of movement provided by the cone.



To change the Twist limits:

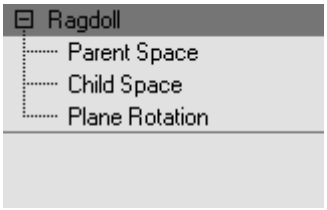
1. Select the Rag Doll helper object and open the Properties rollout.
2. The Twist settings limit the rotation of the child object about the twist axis relative to the parent object. To fix the twist of one object relative to the other, set the twist limits to [0,0].
3. Display the Twist limits by turning on Display rollout > Show Twist.



For a detailed explanation of how to define these limits, see *Defining Rag Doll Limits (page 2–743)*.

Interface

Modifier Stack



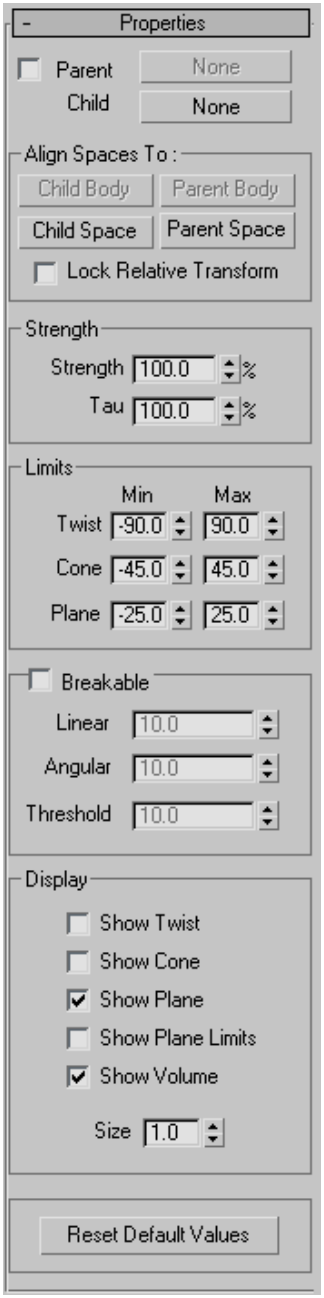
Parent Space—At this sub-object level, you can move and rotate the representation of the Rag Doll for the parent body. The constraint axes are defined in the parent body's space, which means that if you alter the position or rotation of the parent, the space moves as well and maintain its offset rotation and translation with respect to the parent. When you assign a parent body to the Rag Doll, Parent Space is also aligned with the child body's local space, with the attachment point at the child body's pivot point, the twist axis aligned with the child's X axis, and the plane normal aligned to the child's Y axis. If the Rag Doll is single-bodied, then Parent Space represents the world orientation and position of the constraint.

Child Space—At this sub-object level, you can move and rotate the representation of the Rag Doll for the child body. The constraint axes are defined in the child body's space, which means that if you alter the position or rotation of the child, the space moves as well and maintain its offset rotation and translation with respect to the child. When you assign a child body to the Rag Doll, the Child Space is positioned so that the attachment point coincides with the child body pivot point, and the twist axis is aligned with the child's X axis.

Note: While the Child Space sub-object has its own representation of the plane and plane normal, these are not displayed at this sub-object level, as the plane normal can be edited only in the parent body's space.

Plane Rotation—This sub-object level allows independent rotation of the plane normal with respect to the Parent Space. This means that your plane normal doesn't have to be perpendicular to your twist axis, allowing for a greater ranger of constraint setups. However, it should be noted that non-orthogonal systems affect the way reactor handles non-symmetrical cone limits.

Properties rollout



Parent—When on, the Rag Doll has two bodies and you can designate an object to be the parent body. When off, the Spring is single-bodied and the point occupied by the Parent Space sub-object is the parent.

You assign the parent object by clicking this button and then selecting a rigid body from any viewport. Thereafter the button displays the name of the parent body.

Child—Displays the name of the second rigid body attached to the Rag Doll. You assign the child object by clicking this button and then selecting a rigid body from any viewport.

Align Spaces To—Use these options to align the bodies' local constraint spaces. You can find out more about each option in *Working With Constraint Spaces (page 2–726)*.

Lock Relative Transform—When on, the relative transform between the child and parent constraint spaces is locked: If you move either space in the viewport, the other space moves along with it.

Strength group

Strength/Tau—Govern the impulses applied to the constraint's bodies in order to maintain the constraint, and so how strongly the constraint works to restrict their movement. You can find out more about these parameters in *Strength and Tau (page 2–735)*.

Limits group

Twist Min—The minimum rotation about the twist axis the constraint allows between the reference and attached objects' constraint spaces. For example, if Min=–5 then, with respect to the constraint spaces, the child cannot rotate more than 5 degrees counterclockwise about the twist axis.

Twist Max—The maximum rotation about the twist axis the constraint allows between the reference and attached objects' constraint spaces. For example, if Max=5 then, with respect to the

constraint spaces, the child cannot rotate more than 5 degrees clockwise about the twist axis.

Cone Min—The minimum value, in degrees, used to specify the limit cone. This value combines with the Cone Max value to generate a symmetrical cone. This cone is rotated about the plane normal by an offset calculated by combining these limits. The rotation of the child's twist axis is then limited to the volume of this cone with respect to the parent.

For example, if the Min=–10 and Max=20, the range is 30 degrees. A 30-degree cone [–15,15] is created about the twist axis and is then rotated by the offset between the limits (–10+20)*0.5=5. This results in the effective limits being [–15,15] rotated by 5 degrees, thus final limits of [–10, 20].

Cone Max—The maximum value, in degrees, used to specify the limit cone. This is combined with the Cone Min value to generate a symmetrical cone. This cone is rotated about the plane normal by an offset calculated by combining these limits. The rotation of the child's twist axis is then limited to being within this cone with respect to the parent.

Plane Min—The angle between the plane and a reference line in the range [-90,0]. This line is swept about the plane normal to create a cone. The child's twist axis is then restricted from entering this cone. If this cone intersects the cone produced by the cone limits, it further inhibits the motion of the child with respect to the parent.

Plane Max—The angle between the plane and a reference line in the range [0,90]. This line is swept about the plane normal to create a cone. The child's twist axis is then restricted from entering this cone. If this cone intersects the cone produced by the cone limits, it further inhibits the motion of the child with respect to the parent.

Breakable group

Breakable—When on, the constraint is breakable. If its breakable limits are exceeded during simulation it will cease to exert impulses on the attached objects. You can find out more about breakable constraints in *Breakable Constraints (page 2–735)*.

Display group

Show Twist—When on, Twist limits display as part of the Parent Space visualization. The representation is a cylinder section showing the restricted motion of the attached object with respect to rotation about the twist axis.

Show Cone—When on, the entire symmetrical cone produced by the Cone > Min and Max limits displays as part of the Parent Space visualization.

Show Plane—When on, representations of the plane and plane normal display as part of the Parent Space visualization.

Show Plane Limits—When on, the cones formed by the plane limits display as part of the Parent Space visualization.

Show Volume—If true, a representation of the allowed region of movement based on a combination of the plane and cone values will be calculated and displayed as part of the Parent Space visualization.

Size—Allows you to change the size of the limit visualization in the viewport.

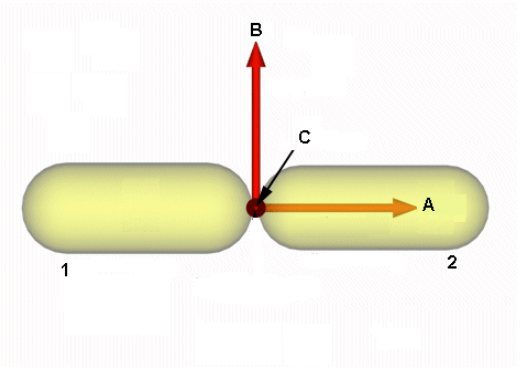
Reset Default Values—Resets the parameters for this constraint to their default values.

Defining Rag Doll Limits

The *Rag Doll constraint (page 2–737)* is useful for simulating complex joints, like those found in some human and animal joints (shoulders, hips), that cannot be properly simulated by simpler

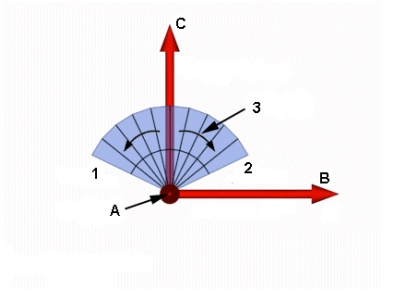
constraints like *Hinge (page 2–747)* or *Point-Point (page 2–750)* constraints.

Because they represent complex joints, the rotational limits use a complex parameterization that can benefit from some explanation. This section covers the different axes, planes, and cones used for that parameterization. We'll begin with a simple picture of two objects connected by a rag doll constraint; you can think of a shoulder joint, where the parent body is the torso and the child body the upper arm:

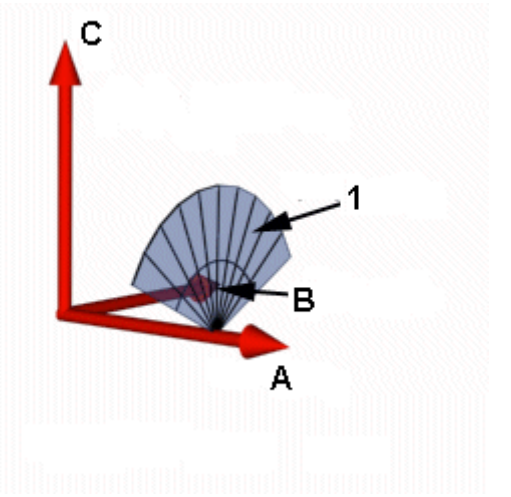


A. Twist Axis
B. Plane Axis
C. Twist X Plane Axis (outward)
1. Parent Body
2. Child Body

The illustration shows three axes of importance: the *Twist Axis*, the *Plane Axis* and the *Twist X Plane Axis*. The twist axis is probably the easiest to visualize and you can do so by simply stretching out your arm and rotating it so that your thumb changes from pointing upwards to downwards. Your movement is just a series of rotations along the twist axis of the shoulder and elbow. In general, the twist axis should follow the length of the child body in a joint. In reactor you can specify the extent to either side of the neutral position the constraint can twist; this need not be symmetrical.

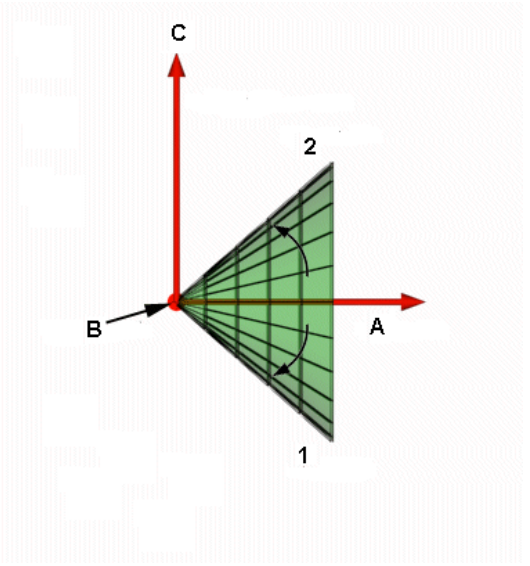


A. Twist Axis (outward)
B. Plane Axis
C. Twist X Plane Axis
1. Twist Min.
2. Twist Max.
3. Twist Range

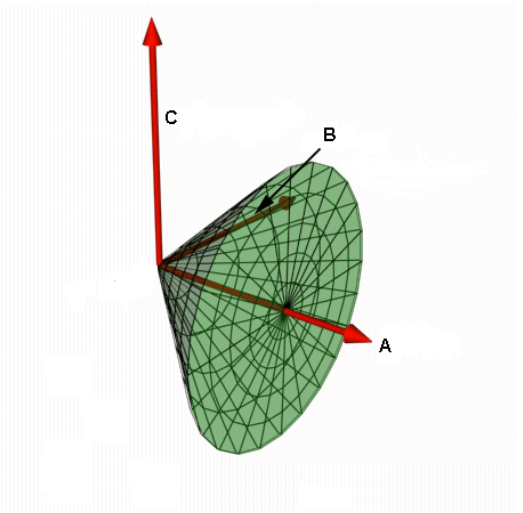


A. Twist Axis
B. Plane Axis
C. Twist X Plane Axis
1. Twist Range

The other two axes are interlinked and together control the volume the child body may occupy. So first we'll look at the *cone* angles that allow you to specify the volume that the child may move through:

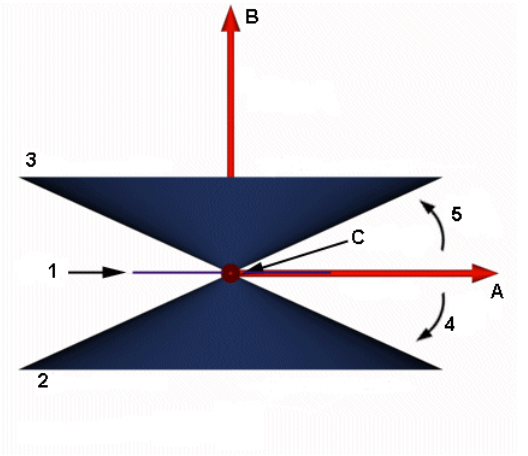


A. Twist Axis
B. Plane Axis (inward)
C. Twist X Plane Axis
1. Cone Min.
2. Cone Max.

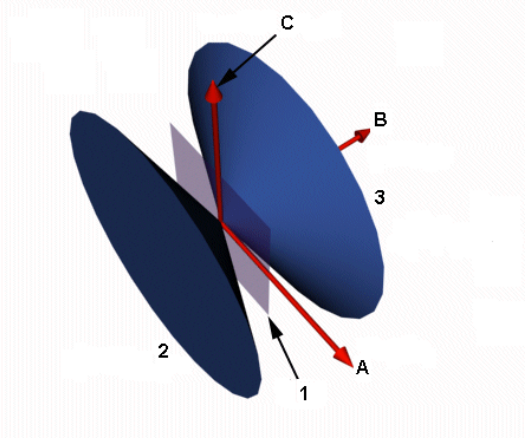


A. Twist Axis
B. Plane Axis
C. Twist X Plane Axis

If you imagine the constraint representing your shoulder, then the green volume is the range of allowed positions your upper arm may occupy. However, we can refine the allowed volume further by using the plane axis to enforce *plane limits*. These limits are used to generate two further cones:

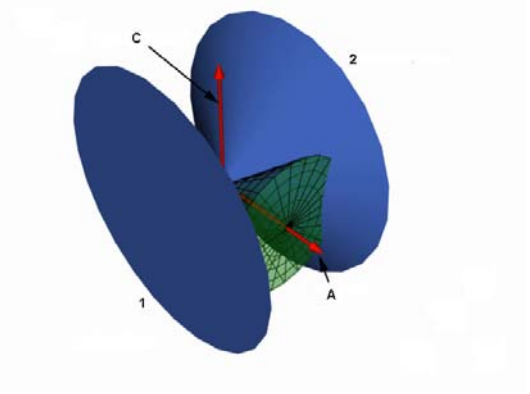


A. Twist Axis
B. Plane Axis
C. Twist X Plane Axis (outward)
1. Plane
2. Plane Min. Cone
3. Plane Max. Cone
4. Plane Min.
5. Plane Max.

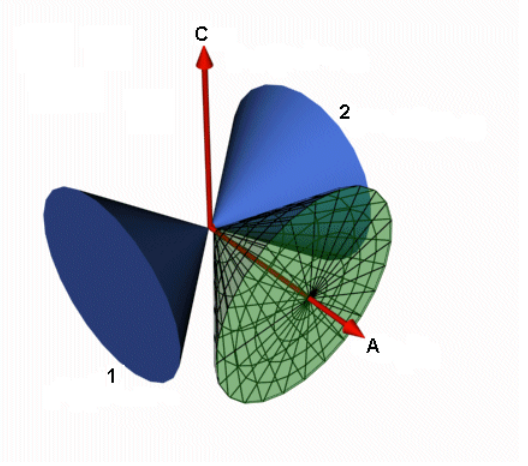


A. Twist Axis
B. Plane Axis
C. Twist X Plane Axis
1. Plane
2. Plane Min. Cone
3. Plane Max. Cone

This allows us to create two different scenarios: one where the cones generated by the plane limits intersect with cone produced by the cone angles, and another where they do not:

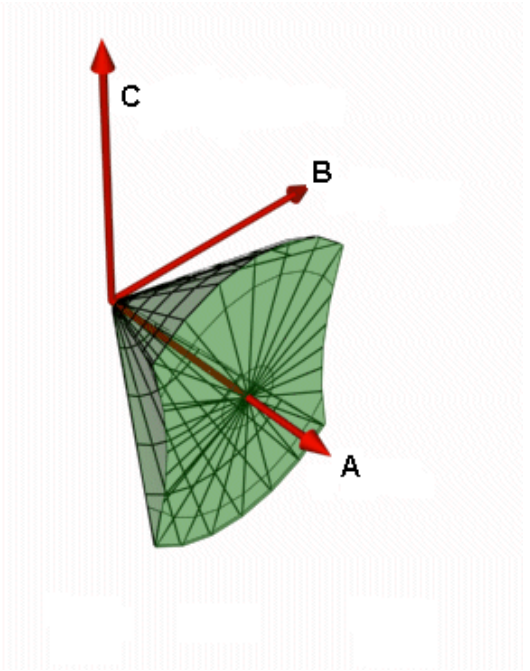


A. Twist Axis
C. Twist X Plane Axis
1. Plane Min. Cone
2. Plane Max. Cone



A. Twist Axis
C. Twist x Plane Axis (outward)
1. Plane Min. Cone
2. Plane Max. Cone

If the volumes do not intersect, the plane limits have no effect in the constraint. However, by allowing the cones to intersect, and deeming the overlapping volume as invalid body positions, we can limit the allowed positions for the child body even further:



A. Twist Axis
B. Plane Axis
C. Twist X Plane Axis

Back to *Rag Doll Constraint (page 2–737)*.

Hinge Constraint

Create panel > Helpers > reactor > Hinge
Menu bar > reactor > Create Object > Hinge Constraint
reactor toolbar > Create Hinge Constraint button

The Hinge constraint allows you to simulate a hinge-like action between two bodies. reactor lets you specify an axis in local space for each body, with a position and a direction. During the simulation, the two axes attempt to match position and direction, thereby creating an axis around which the two bodies can rotate. Alternatively, you can hinge a single body to an axis in world space.

You can also limit the rotation between the attached objects relative to the hinge axis, allowing you, for instance, to create a door that can open only to a specified extent. You define this limitation with respect to an axis perpendicular to the hinge axis for each body.

A Hinge constraint is simulated only if it has the correct number of rigid bodies attached and is included in a valid *Constraint Solver (page 2–736)*. When not selected, an invalid Hinge is red in the viewport.

Procedures

To create a Hinge constraint:

- Choose one of the above options, and then click in any viewport to add the Hinge constraint.

  Note: The icon's position has no effect on the constraint's behavior.

To attach objects to the Hinge:

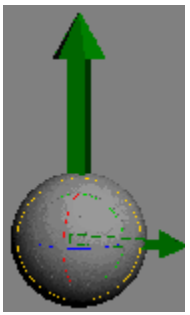
1. Create the Hinge and the objects to connect using it.
2. On the Hinge's Properties rollout, click the Child pick button, and then select the object to use as the child object in one of the viewports. By default, this attaches the constraint so that its axis goes through the child body pivot point and the axes are aligned with the child body's local space. The constraint axis is aligned with the child's Z axis, and the perpendicular axis used for rotation limits is aligned to the child's X axis.
3. If you do not want to assign a second body to the Hinge, turn off the Parent check box if necessary.
4. If the Hinge is to be two-bodied, repeat step 2, using the Parent pick button to set Parent for the Hinge.

To create a Hinge and attach objects in a single step:

- See the *Scripts (page 2–817)* topic.

To change the Hinge positions and rotations for the bodies:

1. With the Hinge helper object selected, open the Modify panel.
2. In the modifier stack open the sub-object list for the Hinge.
3. Access the sub-object level corresponding to the object whose hinge axes you want to move: Child Space or Parent Space. If the Hinge is single-bodied, Parent Space represents the world attachment position and rotation of the axes.
4. The space, represented by a pair of axes, the larger of which represents the hinge axis, the smaller of which is used to specify limits, becomes active in the viewport. It can be moved and rotated using the Move and Rotate tool respectively. The space maintains its position and rotation relative to its corresponding object. For example, if you move or rotate the child body, the Child Space moves with it.



5. To reset the axes to their default alignment, reattach the objects to the Hinge or click the Align Spaces To > Child Body button.

To limit the rotation between the attached objects:

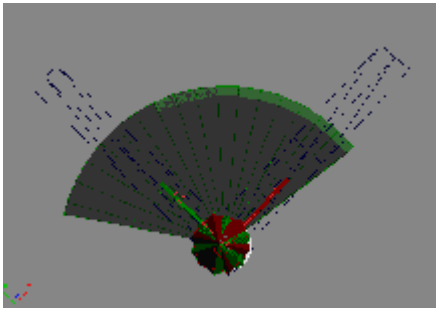
1. With the Hinge selected, open the Properties rollout.
2. Turn on the Limited check box.

   The Hinge limits display around the Parent Space axis.

   The default angle limits are –90.0 and 90.0. This means that in a situation where the sub-objects are fully aligned for the objects, the objects will be able to rotate 90 degrees in either direction about the shared hinge axis.

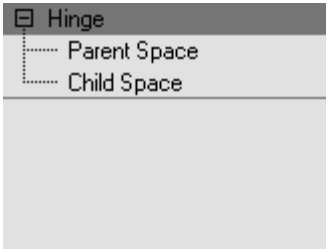
   The child object is free to rotate within the indicated limits relative to the parent during the simulation.
3. Change the minimum and maximum rotations the constraint maintains between the objects by adjusting the Min Angle and Max Angle values.

   The limit display updates to reflect your changes, to help you choose appropriate values. Upon simulation you can see that the rotation between the objects has been limited.



Interface

Modifier Stack

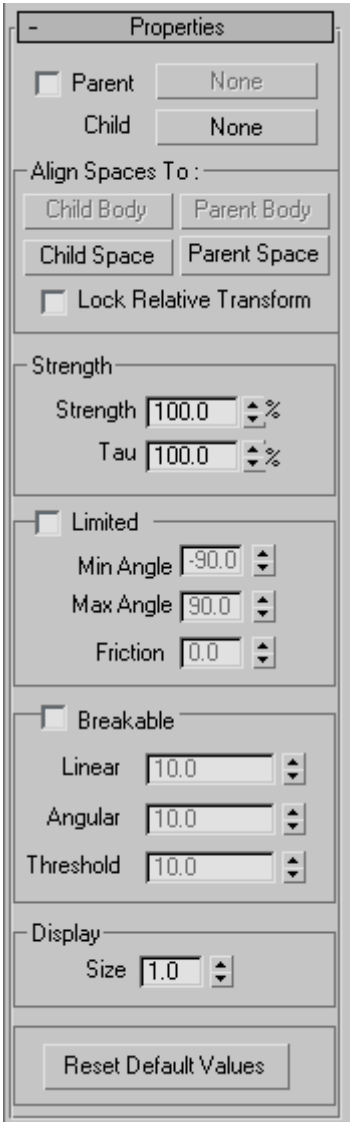


Parent Space—At this sub-object level, you can move and rotate the representation of the Hinge for

the parent body. The Hinge axes are defined in the parent body's space, which means that if you alter the position or rotation of the parent, the space moves as well and maintains its offset rotation and translation with respect to the parent. When you assign a parent body to the Hinge, Parent Space is also aligned with the child body's local space, with the Hinge axis going through the child body's pivot point. If the Hinge is single-bodied, then Parent Space represents the world orientation and position of the Hinge.

Child Space—At this sub-object level, you can move and rotate the representation of the Hinge for the child body. The Hinge axes are defined in the child body's space, which means that if you alter the position or rotation of the child, the space moves as well and maintains its offset rotation and translation with respect to the child. When you assign a child body to the Hinge, the Child Space hinge axis is positioned so that it goes through the child body pivot point, and the axes are aligned with the child body's local space: The Hinge axis is aligned with the child's Z axis, and the perpendicular axis used for rotation limits is aligned to the child's X axis.

Properties rollout



Parent—When on, the Hinge has two bodies and you can designate an object to be the parent body. When off, the Hinge is single-bodied and the point occupied by the Parent Space sub-object is the parent.

You assign the parent object by clicking this button and then selecting a rigid body from any viewport.

Thereafter the button displays the name of the parent body.

Child—Displays the name of the second rigid body attached to the Hinge. You assign the child object by clicking this button and then selecting a rigid body from any viewport.

Align Spaces To—Use these options to align the bodies' local constraint spaces. You can find out more about each option in *Working With Constraint Spaces (page 2–726)*.

Lock Relative Transform—When on, the relative transform between the child and parent constraint spaces is locked: If you move either space in the viewport, the other space moves and rotates along with it, and vice-versa.

Strength group

Strength/Tau—Govern the impulses applied to the constraint's bodies in order to maintain the constraint, and so how strongly the constraint works to restrict their movement. You can find out more about these parameters in *Strength and Tau (page 2–735)*.

Limited group

Limited—When on, limits the rotation about the hinge axis between the attached objects. When this value is true, the child object will only be free to rotate relative to the parent within the range specified by the Min Angle and Max Angle values. You can also specify a friction value for a limited Hinge.

Min Angle—The minimum rotation the constraint allows between the two hinge spaces. For example, if Min Angle is –5.0 then the child body cannot rotate more than 5 degrees counter clockwise about the hinge axis relative to the parent body.

Max Angle—The maximum rotation the constraint allows between the two hinge spaces. For example, if Max Angle is 5.0 then the child body cannot rotate more than 5 degrees clockwise about the hinge axis relative to the parent body.

Friction—The level of friction applied to the objects as they try to rotate about the hinge axis. It is only applied to limited Hinges.

Breakable group

Breakable—When on, the constraint is breakable. If its breakable limits are exceeded during simulation, it ceases to exert impulses on the attached objects. You can find out more about breakable constraints in *Breakable Constraints (page 2–735)*.

Display group

Size—Lets you change the size of the Hinge limit display in the viewport.

Reset Default Values—Returns the Strength, Tau, Min Angle, Max Angle, Friction, Linear, Angular, Threshold, and Display settings to their default values.

Point-Point Constraint 

Create panel > Helpers > reactor > Point-Point

Menu bar > reactor > Create Object > Point-Point Constraint

reactor toolbar > Create Point-Point Constraint button

The Point-Point (point-to-point) constraint lets you attach two objects together, or an object to a point in world space. It forces its objects to try to share a common point in space. The objects can rotate freely relative to each other, but always have the attachment point in common. When you set up the constraint, the point is defined in the object space of each object involved. During the simulation the constraint tries to apply forces to the objects so that the two pivot points defined by the two objects match.

reactor also provides two variations on the Point-Point constraint: Limited and Stiff Spring. With a Limited Point-Point constraint, you can specify limits for how much the child object can rotate relative to the parent. A Stiff Spring constraint is similar to a standard Point-Point constraint with one important exception: It holds the constrained bodies apart at a specified distance, as if they were attached at each end of an invisible rod. The Stiff Spring variant requires a point in the body space of each of two bodies, and a constant distance by which reactor must keep these two points separated. Each body is free to rotate around its point.

The constraint is simulated when it has the correct number of rigid bodies attached and is included in a valid *Constraint Solver (page 2–736)*. When not selected, an invalid Point-Point constraint is red in the viewport.

Procedures

To create a Point-Point constraint:

- Choose one of the above options, and then click in any viewport to add the Point-Point constraint.

  Note: The icon's position has no effect on the constraint's behavior.

To attach objects to the Point-Point constraint:

1. Create the constraint and the objects to connect with it.
2. On the constraint's Properties rollout, click the Child pick button, and then select the object to use as the child object in one of the viewports. By default, this attaches the constraint to the body at the body's pivot point.
3. If you do not want to assign a second body to the constraint, turn off the Parent check box if necessary.
4. If the constraint is to be two-bodied, make sure Parent is on and then use the Parent pick button to set the parent for the constraint. By default, the attachment point for the parent is also aligned with the child body's pivot point, which means that the bodies will maintain their current positions relative to each other in the simulation.

To create the constraint and attach objects in a single step:

- See the *Scripts (page 2–817)* topic.

To change the attachment positions for the bodies:

1. With the constraint selected, open the Modify panel.
2. In the modifier stack open the constraint's sub-object list.
3. Access the sub-object level corresponding to the object in whose space you want to move the attachment point: Parent Space or Child Space. If the constraint is single-bodied, Parent Space allows you to move the world attachment position.
4. The corresponding point becomes active in the viewports and can be moved using the Move tool.

   The attachment point maintains its position relative to its corresponding object. For example, if you move the child object, its space, and hence its attachment point, moves with it. The attachment point doesn't have to actually be on either object.
5. To reset the attachment point to the child object's pivot, reattach the objects to the constraint, or click the Align Spaces To > Child Body button.

To limit the rotation between the attached objects:

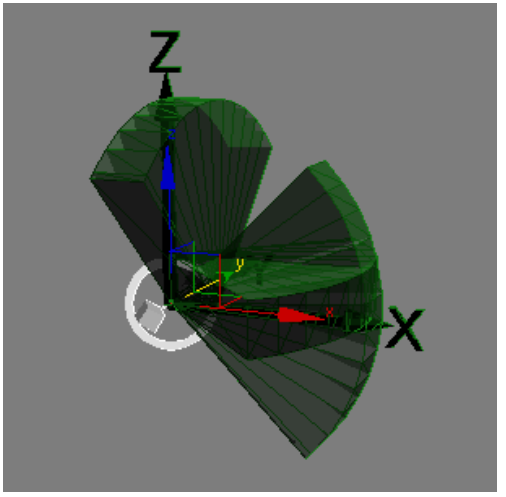
1. With the constraint selected, open the Properties rollout.

2. In the Constraint Type group, choose Limited.

   By default, the child body can rotate a total of 90 degrees around each axis relative to the parent during the simulation: 45 degrees clockwise and 45 degrees counterclockwise.

3. You can then choose which limits to display around the Parent Space constraint axis by turning on each one's check box.

4. If you then change the limits, the limit display updates to reflect your changes to help you choose appropriate values. Upon simulation you can see that the rotation between the objects has been limited.



5. You can also change the orientation of the axis to be limited by rotating the Child Space and Parent Space sub-objects.

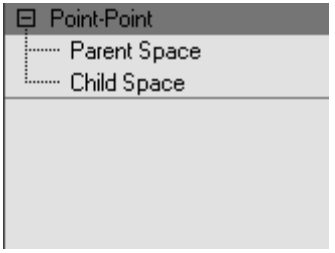
To create a Stiff Spring constraint:

1. With a Point-Point constraint selected, open the Properties rollout.
2. In the Constraint Type group, choose Stiff Spring.
3. You can now specify a Length value for the stiff spring. The default is 0.0: the same as a Point-Point constraint.

   During the simulation, reactor works to maintain this distance between the constrained objects' attachment points. To use the current distance between the attachment points as the length, click the Set button.

Interface

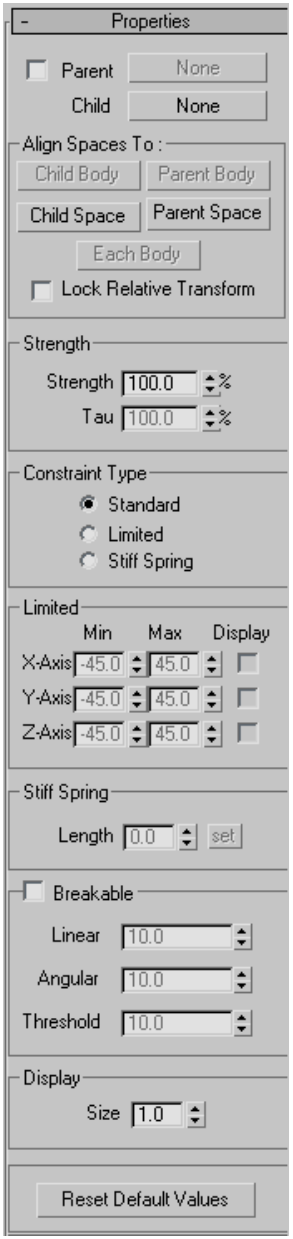
Modifier Stack



Parent Space—At this sub-object level, you can select and move the constraint attachment point in the parent body's local space. When you assign a parent object to the constraint, the Parent Space attachment point is set to the child object's pivot point. If you alter the position of the Parent Space attachment point, it maintains its new position relative to the parent during the simulation. If the constraint is single-bodied, the Parent Space's position represents the point in the world where the constraint is attached.

Child Space—At this sub-object level, you can select and move the constraint attachment point in the child body's local space. When you assign a child body to the constraint, the Child Space attachment point is set to the child object's pivot point. If you alter the position of the Child Space attachment point, it maintains its new position relative to the child during the simulation.

Properties rollout



Parent—When on, the constraint has two bodies and you can designate an object to be the parent body. When off, the constraint is single-bodied and the point occupied by the Parent Space sub-object is the parent.

You assign the parent object by clicking this button and then selecting a rigid body from any viewport. Thereafter the button displays the name of the parent body.

Child—Displays the name of the second rigid body attached to the constraint. You assign the child object by clicking this button and then selecting a rigid body from any viewport.

Align Spaces To—Use these options to align the bodies' local constraint spaces. You can find out more about each option in *Working With Constraint Spaces (page 2–726)*.

Lock Relative Transform—When on, the relative transform between the child and parent constraint spaces is locked: If you move either space in the viewport, the other space moves and rotates along with it, and vice-versa.

Strength group

Strength/Tau—Govern the impulses applied to the constraint's bodies in order to maintain the constraint, and so how strongly the constraint works to restrict their movement. You can find out more about these parameters in *Strength and Tau (page 2–735)*.

Constraint Type group

These options let you choose the type of Point-Point constraint to add to your scene.

- **Standard—**The default option.
- **Limited—**This type allows you to limit how the child object can rotate relative to the parent.
- **Stiff Spring—**This type allows you to specify a distance for reactor to maintain between the constraint attachment points.

Limited (Limited constraints only)—Lets you specify how much the child object can rotate relative to

the parent around the x, y, and z axes - for each axis there is a minimum (counterclockwise) and maximum (clockwise) limit. You can use the check boxes to switch on or off the limit displays in the viewport.

Length (Stiff Springs only)—The distance that reactor should maintain between the constraint attachment points.

Set—Sets Length to the current distance between the attachment points.

Breakable group

Breakable—When on, the constraint is breakable. If its breakable limits are exceeded during simulation, it ceases to exert impulses on the attached objects. You can find out more about breakable constraints in *Breakable Constraints (page 2–735)*.

Display group

Size—Lets you change the size of the constraint limit display in the viewport.

Reset Default Values—Returns the constraint's parameters to their default values.

Prismatic Constraint 

Create panel > Helpers > reactor > Prismatic

Menu bar > reactor > Create Object > Prismatic Constraint

reactor toolbar > Create Prismatic Constraint button

The Prismatic constraint serves as a constraint between two rigid bodies, or a rigid body and the world, that allows its bodies to move relative to each other along one axis only. Rotations, as well as the remaining two translation axes, are fixed. For example, you could use a Prismatic constraint when creating a forklift truck.

A Prismatic constraint is simulated if it has the correct number of rigid bodies attached and is included in a valid *Constraint Solver (page 2–736)*. When not selected, an invalid prismatic constraint is red in the viewport.

Procedures

To create a Prismatic constraint:

- Choose one of the above options, and then click in any viewport to add the Prismatic constraint.

  Note: The icon's position has no effect on the constraint's behavior.

To attach objects to the prismatic constraint:

1. Create the Prismatic constraint and the objects to connect using it.
2. On the constraint's Properties rollout, click the Child pick button and then, in one of the viewports, select the object to use as the child.

   This attaches the constraint to the body at the body's pivot point. The constraint's sliding axis is aligned to the Z axis of the object.
3. If you do not want to assign a second body to the prismatic constraint, turn off the Parent check box, if necessary.
4. If the constraint is to be two-bodied, make sure Parent is on and then use the Parent pick button to set the parent for the constraint.

To create the constraint and attach objects in a single step:

- See the *Scripts (page 2–817)* topic.

To change the Prismatic constraint positions and rotations for the bodies:

1. With the constraint selected, open the Modify panel.
2. In the modifier stack open the constraint's sub-object list.

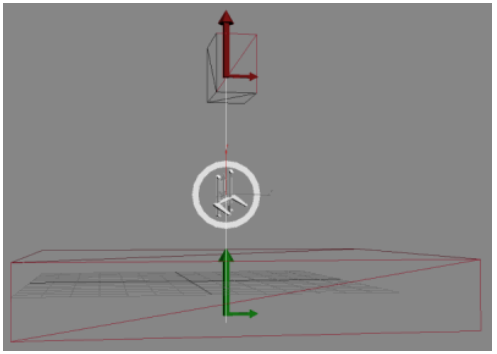
3. Access the sub-object level corresponding to the object whose attachment point you want to move: Child Space or Parent Space. If the constraint is single-bodied, Parent Space represents the world attachment position and rotation of the constraint axes.
4. The space, represented by a pair of axes, the larger of which represents the sliding axis, becomes active in the viewport and can be moved or rotated using the Move or Rotate tool, respectively. The space maintains its position and rotation relative to its corresponding object. For example, if you move or rotate the child body, Child Space moves with it.
5. To realign the constraint space to the child body, reattach the object to the Prismatic constraint, or choose the Align To Child Body option.

To limit the movement of the child body along the sliding axis:

1. With the constraint selected, open the Properties rollout.
2. Turn on the Limited check box.

   Both Limit parameters (Min and Max) are set by default to 0.0. This means that from a situation where the sub-objects are fully aligned for the objects, the child can move 0 units in either direction along the shared sliding axis relative to the parent.

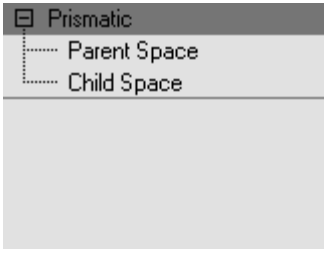
   The limits are displayed as a line in Parent Space.



3. Change the Min Limit and Max Limit settings as necessary.

Interface

Modifier Stack

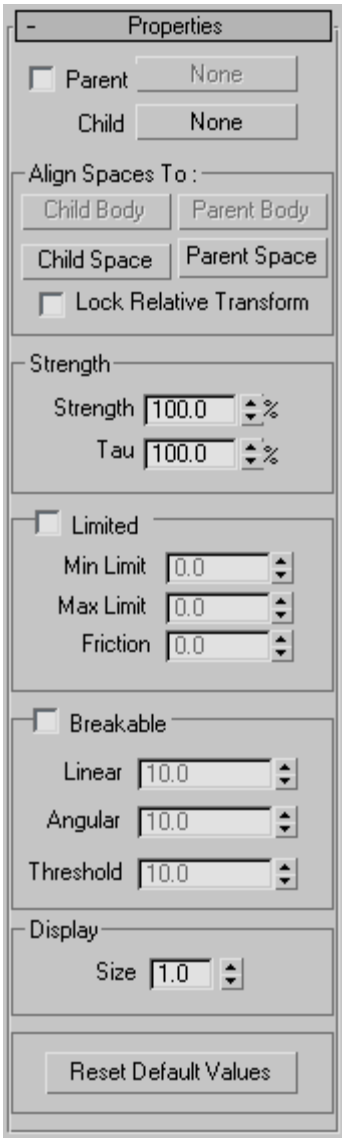


Parent Space—At this sub-object level, you can move and rotate the representation of the Prismatic constraint for the parent body. The constraint axes are defined in the parent body's space, which means that if you alter the position or rotation of the parent, the space moves as well and maintains its offset rotation and translation with respect to the parent. When you assign a parent body to the constraint, Parent Space is also aligned with the child body's local space, with the constraint axis going through the child body's pivot point. If the constraint is single-bodied, then Parent Space represents the world orientation and position of the constraint.

Child Space—At this sub-object level, you can move and rotate the representation of the Prismatic constraint for the child body. The constraint

axes are defined in the child body's space, which means that if you alter the position or rotation of the child, the space moves as well and maintains its offset rotation and translation with respect to the child. When you assign a child body to the constraint, the Child Space axis is positioned so that it goes through the child body pivot point, and the axes are aligned with the child body's local space: The prismatic sliding axis is aligned with the child's Z axis.

Properties rollout



Parent—When on, the constraint has two bodies and you can designate an object to be the parent body. When off, the constraint is single-bodied and the point occupied by the Parent Space sub-object is the parent.

You assign the parent object by clicking this button and then selecting a rigid body from any viewport. Thereafter the button displays the name of the parent body.

Child—Displays the name of the second rigid body attached to the constraint. You assign the child object by clicking this button and then selecting a rigid body from any viewport.

Align Spaces To—Use these options to align the bodies' local constraint spaces. You can find out more about each option in *Working With Constraint Spaces (page 2–726)*.

Lock Relative Transform—When on, the relative transform between the child and parent constraint spaces is locked: If you move either space in the viewport, the other space moves and rotates along with it, and vice-versa.

Strength group

Strength/Tau—Govern the impulses applied to the constraint's bodies in order to maintain the constraint, and so how strongly the constraint works to restrict their movement. You can find out more about these parameters in *Strength and Tau (page 2–735)*.

Limited group

Limited—When on, limits the movement of the child object along the sliding axis; the child object can move relative to the parent only within the range specified by the Min Limit and Max Limit values. You can also specify a Friction value for a limited prismatic constraint.

Min Limit—The minimum distance the constraint allows between the two spaces. For example, if Min Limit is –5.0 then the child body then the child object cannot move more than 5 units backwards relative to the parent along the sliding axis.

Max Limit—The maximum distance the constraint allows between the two spaces. For example, if Max Limit is 5, then the child object will not be able to move more than 5 units forward relative to the parent along the sliding axis.

Friction—The level of friction applied to the objects as they try to move along the sliding axis. It is applied only to limited Prismatic constraints.

Breakable group

Breakable—When on, the constraint is breakable. If its breakable limits are exceeded during simulation, it ceases to exert impulses on the attached objects. You can find out more about breakable constraints in *Breakable Constraints (page 2–735)*.

Display group

Size—Lets you change the size of the constraint space display in the viewport.

Reset Default Values—Returns the constraint's parameters to their default values.

Car-Wheel Constraint 

Create panel > Helpers > reactor > Carwheel
Menu bar > reactor > Create Object > Car-Wheel Constraint
reactor toolbar > Create Car-Wheel Constraint button

You can use this constraint to attach a wheel to another object; for instance, a car chassis. You can also constrain a wheel to a position in world space. During the simulation, the wheel object is free to rotate around a spin axis defined in each object's space. Linear motion is allowed for the wheel along a suspension axis. You can also add limits to the wheel's movement along this axis. The constraint's child body always acts as the wheel, while the parent acts as the chassis.

The Car-Wheel constraint also has spin parameters. If these are nonzero, the constraint turns the wheel during the simulation. You can specify a target velocity and a target gain: the maximum angular impulse that the motor can apply to the rigid body in order to achieve the target velocity.

A Car-Wheel constraint is simulated if it has the correct number of rigid bodies attached and is included in a valid *Constraint Solver (page 2–736)*. When not selected, an invalid Car-Wheel constraint is red in the viewport.

Procedures

To create a Car-Wheel constraint:

- Choose one of the above options, and then click in any viewport to add the Car-Wheel constraint.

  Note: The icon's position has no effect on the constraint's behavior.

To attach objects to the Car-Wheel constraint:

1. Create the constraint and the objects to connect using it.
2. On the constraint's Properties rollout, click the Child pick button, and then, in one of the viewports, select the object you want to use as the wheel.

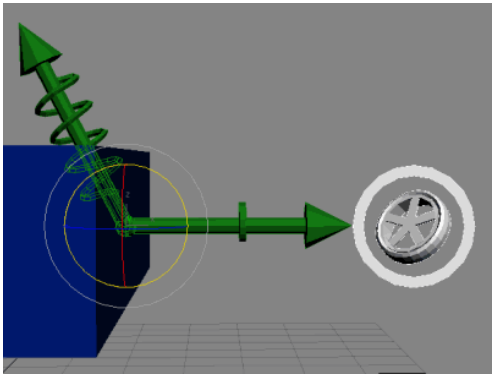
   This attaches the constraint to the wheel at the wheel's pivot point. The spin axis is aligned to the X axis of the wheel object and the suspension axis is aligned to the Z axis.
3. If you do not want to assign a chassis to the Car-Wheel constraint, turn off the Parent check box if necessary.
4. If the Car-Wheel constraint is to be two-bodied, repeat step 2, using the Parent pick button to specify the chassis object for the constraint.

To create the constraint and attach objects in a single step:

- See the *Scripts (page 2–817)* topic.

To change the positions and rotations for the bodies:

1. With the constraint selected, open the Modify panel.
2. In the modifier stack open the constraint's sub-object list.
3. Access the sub-object level corresponding to the object whose constraint space you want to move: Parent Space for the chassis or Child Space for the wheel. If the constraint is single-bodied, Parent Space represents the world attachment position and rotation of the constraint axes.
4. The space, represented by just a spin axis for the wheel and by the spin and suspension axes for the chassis, becomes active in the viewport and can be moved or rotated using the Move or Rotate tool, respectively. The space maintains its position and rotation relative to its corresponding object. For example, if you move or rotate the wheel, Child Space moves with it.



5. To reset the attachment point to the child's pivot, reattach the objects to the car-wheel constraint or choose the Align Spaces To Child Body option.

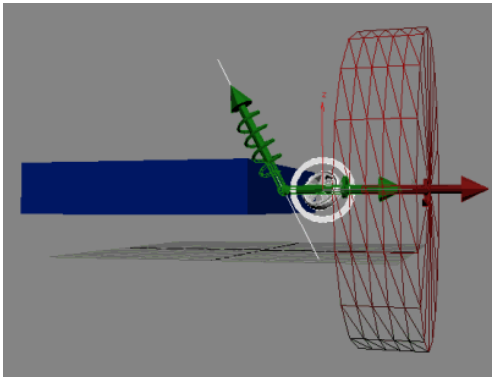
6. You can also rotate the suspension axis independent of the spin axis. To do this, highlight the Suspension Rotation sub-object level; this lets you rotate the suspension axis with respect to the chassis's spin axis.

To change the suspension limits:

1. With the car-wheel constraint selected, open the Properties rollout.

   The default limit settings in the Suspension Parameters group are Min Limit=0.0 and Max Limit=0.0 . This means that in a situation where the constraint spaces are fully aligned for the objects, the wheel cannot move along the suspension axis.

   The limits define the allowed motion of the wheel along the suspension axis with respect to the chassis, and are displayed as a line with respect to the chassis. This means that from a stable position where the sub-objects are aligned, if the limits are –5 and 10, the wheel will be able to move 10 units away from chassis along the suspension axis in the direction of the icon's arrow and 5 units in the opposite direction.



2. Change the limits and observe the difference in the simulation.

3. You can also change Friction for the Car-Wheel constraint; this inhibits the wheel's motion along the suspension axis.

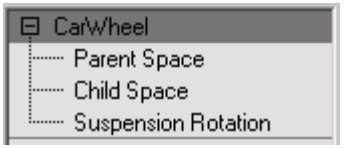
To spin the wheel:

These controls are found in the Spin Parameters group.

1. Set the Velocity value to the required value. This is specified in radians per second.
2. Set the Gain. This is the maximum angular impulse the constraint can apply to maintain this angular velocity.

Interface

Modifier Stack



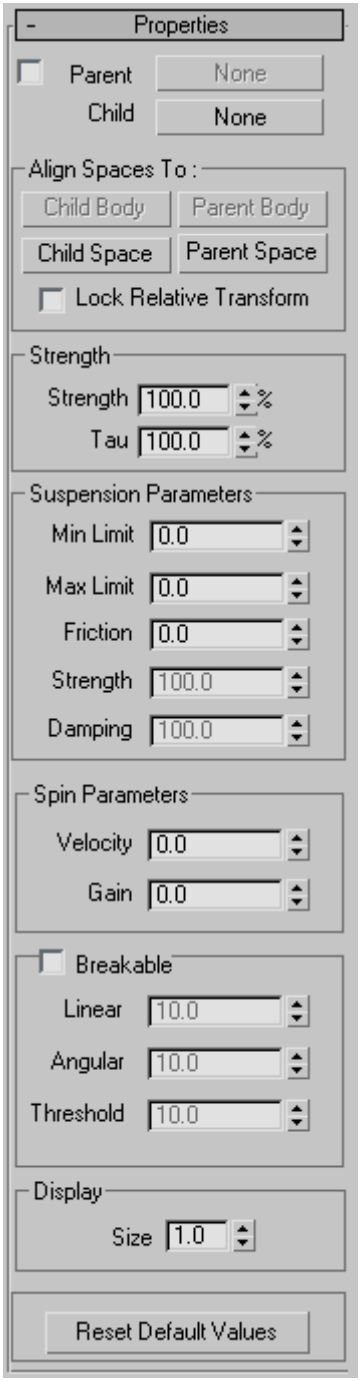
Parent Space—At this sub-object level, you can move and rotate the representations of spin and suspension axes for the chassis. This describes where on the chassis the wheel is attached and also the directions of the spin and suspension axes for the chassis. This is defined in chassis space, which means that if you alter the position or rotation of the chassis the axes will move as well and maintain their offset rotation and translation with respect to the chassis. When you assign a chassis to the car-wheel constraint, Chassis Space's translation is set to the wheel's pivot point, the spin axis is aligned to the wheel's X axis, and the suspension axis is aligned to its Z axis. If the car-wheel constraint is single-bodied, then Chassis Space represents the world orientation and position of the constraint.

Child Space—At this sub-object level, you can move and rotate the representation of the constraint for the wheel. This describes both the location on the wheel that the constraint considers as its

attachment point and the rotation of the spin axis for the wheel. This is defined in wheel space, which means that if you alter the position or rotation of the wheel the space moves as well and maintains its offset rotation and translation with respect to the wheel. When you assign a wheel body to the Car-Wheel constraint, Wheel Space's translation is set to the wheel's pivot point and the wheel's spin axis is aligned to its X axis.

Suspension Rotation—This level allows independent rotation of the suspension axis with respect to the Chassis Space. This means that your suspension axis doesn't have to be perpendicular to your spin axis, allowing for a greater range of constraint setups.

Properties rollout



Parent—When on, the constraint has two bodies and you can designate an object to be the optional chassis rigid body attached to the Car-Wheel constraint. When off, the constraint is single-bodied and the point occupied by the Parent Space sub-object is the parent.

You assign the chassis by clicking this button and then selecting a rigid body from any viewport. Thereafter the button displays the name of the chassis object.

Child—Displays the name of the second rigid body attached to the constraint. You assign the child object by clicking this button and then selecting a rigid body from any viewport.

Align Spaces To—Use these options to align the bodies' local constraint spaces. You can find out more about each option in *Working With Constraint Spaces (page 2–726)*.

Lock Relative Transform—When on, the relative transform between the child and parent constraint spaces is locked: If you move either space in the viewport, the other space moves and rotates along with it, and vice-versa.

Strength group

Strength/Tau—Govern the impulses applied to the constraint's bodies in order to maintain the constraint, and so how strongly the constraint works to restrict their movement. You can find out more about these parameters in *Strength and Tau (page 2–735)*.

Suspension Parameters group

Min Limit—The distance the constraint allows the wheel to move along the suspension axis relative to the chassis, in the direction opposite that indicated by the icon's suspension arrow. For example, if Min Limit=–5, then the wheel cannot move more than five units backwards along the suspension axis with respect to the constraint's spaces.

Max Limit—The distance the constraint allows the wheel to move along the suspension axis relative to the chassis, in the direction indicated by the icon's suspension arrow. For example, if Max Limit=10 then with respect to the constraint's spaces the wheel cannot move more than 10 units forward along the suspension axis.

Friction—The level of friction applied to the objects as they try to move relative to each other along the suspension axis. Available with Havok 1 only.

Strength—The strength value of the suspension spring. Higher values cause less vertical movement. Available with Havok 3 only.

Damping—The damping value of the suspension spring. Higher values suppress vertical oscillation. Available with Havok 3 only.

Spin Parameters group

Velocity—This value indicates the angular velocity in radians per second that the constraint will seek to achieve for the wheel.

Gain—This value indicates the maximum angular impulse the constraint can apply to attain the specified angular velocity.

Breakable group

Breakable—When on, the constraint is breakable. If its breakable limits are exceeded during simulation, it ceases to exert impulses on the attached objects. For more information, see *Breakable Constraints (page 2–735)*.

Display group

Size—Lets you change the size of the constraint space display in the viewport.

Reset Default Values—Returns the constraint's parameters to their default values.

Point-Path Constraint 

Create panel > Helpers > reactor > Point-Path

Menu bar > reactor > Create Object > Point-Path Constraint

reactor toolbar > Create Point-Path Constraint button

The Point-Path constraint allows you to constrain two bodies so that the child is free to move along a specified path relative to the parent. Alternatively, you can create a single-bodied version of the constraint, where the constrained body can move along a path in world space. The child body's orientation is not restricted by the constraint.

You could, for example, use a Point-Path constraint to simulate a bead on a wire. In most cases, you don't need to specify a parent rigid body for a Point-Path constraint, unless the motion of the path itself should follow another rigid body.

The constraint is simulated if it has the correct number of rigid bodies attached, a specified path shape, and is included in a valid *Constraint Solver (page 2–736)*. When not selected, invalid constraints are red in the viewport.

Procedures

To create a Point-Path constraint:

- Choose one of the above options, and then click in any viewport to add the Point-Path constraint.

  Note: The icon's position has no effect on the constraint's behavior.

To specify objects and a path for the constraint:

1. Create the constraint and the bodies to constrain.
2. Create a line or curve in the viewport using one of the Shapes tools on the Create panel.
3. In the constraint's Properties rollout, click the Child pick button, and then select the object you want to use as the child object in one of the viewports.
4. If you do not want to assign a second body to the constraint, turn off the Parent check box if necessary.
5. If the constraint is to be two-bodied, repeat step 2, using the Parent pick button to set the parent for the constraint.
6. Use the Path pick button to designate the path.

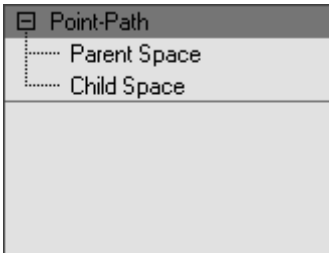
   The constraint icon moves to the designated path. By default, this constraint type is aligned so that the parent constraint space is aligned with the path shape's local space, while the child space is aligned with the child body's local space.

To change the path position and orientation for the bodies:

1. With the constraint selected, open the Modify panel.
2. In the modifier stack open the constraint's sub-object list.
3. Access the sub-object level corresponding to the object whose constraint space you want to move: Child Space or Parent Space. If the Point-Path constraint is single-bodied, Parent Space represents the world position and rotation of the path.
4. The space, represented by a set of axes and the path, becomes active in the viewport and can be moved or rotated using the Move or Rotate tool, respectively. The space maintains its position and rotation relative to its corresponding object. For example, if you move or rotate the child body, Child Space moves with it.

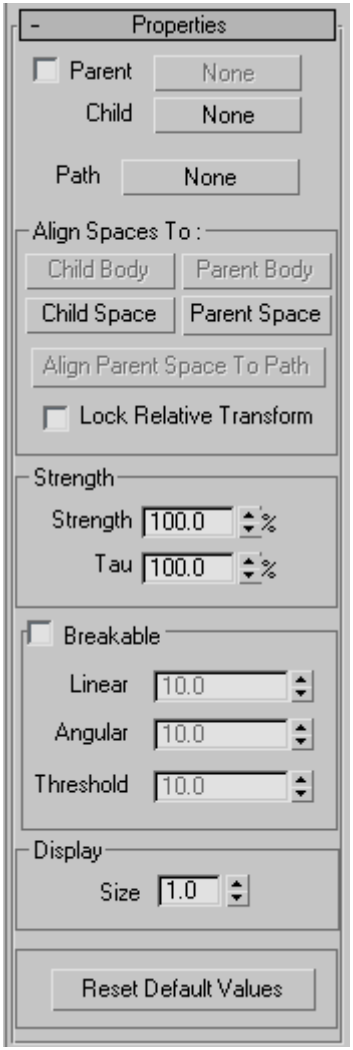
Interface

Modifier Stack



Parent Space—At this sub-object level, you can move and rotate the representation of the constraint for the parent body. The axes and path are defined in the parent body's space, which means that if you alter the position or rotation of the parent, the space moves as well and maintains its offset rotation and translation with respect to the parent. When you assign a parent body to the constraint, Parent Space is aligned with the path's local space. If the constraint is single-bodied, then Parent Space represents the world orientation and position of the path.**Child Space—**At this sub-object level, you can move and rotate the representation of the constraint for the child body. The axes and path are defined in the child body's space, which means that if you alter the position or rotation of the child, the space moves as well and maintains its offset rotation and translation with respect to the child. When you assign a child body to the constraint, the Child Space is aligned with the child body's local space.

Properties rollout



Parent—When on, the constraint has two bodies and you can designate an object to be the parent body. When off, the constraint is single-bodied and the point occupied by the Parent Space sub-object is the parent.

You assign the parent object by clicking this button and then selecting a rigid body from any viewport. Thereafter the button displays the name of the parent body.

Child—Displays the name of the second rigid body attached to the constraint. You assign the child object by clicking this button and then selecting a rigid body from any viewport.

Align Spaces To—Use these options to align the bodies' local constraint spaces. You can find out more about each option in *Working With Constraint Spaces (page 2–726)*.

Lock Relative Transform—When on, the relative transform between the child and parent constraint spaces is locked: If you move either space in the viewport, the other space moves and rotates along with it, and vice-versa.

Strength group

Strength/Tau—Govern the impulses applied to the constraint's bodies in order to maintain the constraint, and so how strongly the constraint works to restrict their movement. You can find out more about these parameters in *Strength and Tau (page 2–735)*.

Breakable group

Breakable—When on, the constraint is breakable. If its breakable limits are exceeded during simulation, it ceases to exert impulses on the attached objects. You can find out more about breakable constraints in *Breakable Constraints (page 2–735)*.

Display group

Size—Lets you change the size of the constraint space display in the viewport.

Reset Default Values—Returns the constraint's parameters to their default values.

reactor Objects

Plane 

Create panel > Helpers > reactor > Plane

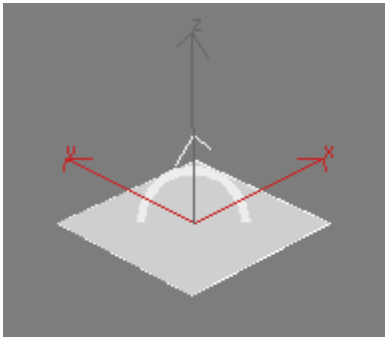
Menu bar > reactor > Create Object > Plane

reactor toolbar > Create Plane button

The reactor Plane object is a type of rigid body that acts as a fixed, infinite plane in the simulation. It shouldn't be confused with the standard 3ds Max plane, which can also be used as a rigid body provided its *simulation geometry property (page 2–719)* is set to Concave Mesh.

Like other rigid bodies, the Plane can have display proxies assigned to it and can be contained in a compound rigid body. You can also assign friction and elasticity values to it. However, it is always fixed in the simulation (the *Mass (page 2–718)* and *Unyielding (page 2–718)* properties are ignored).

The reactor Plane acts in only one direction, as shown in the viewport by an arrow pointing away from its solid surface. This means that rigid bodies approaching the plane from the "wrong" direction pass through it. You can, of course, also use two opposing planes.



The reactor Plane as rendered in a viewport

Procedures

To create a Plane:

1. Choose one of the above options, and then click in any viewport.

   The plane's normal is aligned with the viewport's Z axis.

2. Use the Move and Rotate tools to adjust its position and orientation, as necessary.

Interface

Properties rollout



Show Normal—When on, the Plane normal is displayed as an arrow in the viewport.

Motor 

Create panel > Helpers > reactor > Motor
Menu bar > reactor > Create Object > Motor
reactor toolbar > Create Motor button

The Motor helper lets you apply a turning force to any non-fixed rigid body in the scene. You can specify the angular velocity for the target and the maximum angular impulse the Motor can apply to achieve this velocity.

You do not need to add a Motor explicitly to a simulation, as all valid motors in a scene are added to the simulation by default. A Motor is valid if its Rigid Body property has been set to a valid rigid body in the scene. When not selected, invalid motors are colored red in the viewport.

Procedures

To create a motor:

- Choose one of the above options, and then click in any viewport to add the motor.

  Note: The icon's position has no effect on the motor's behavior.

To set the rigid body for the motor:

1. Create the motor and the rigid body to attach the motor to.
2. On the Properties rollout, click the Rigid Body pick button and then, in one of the viewports, select the object you would like to motor to act on.

   The motor can apply rotational forces about one of the object's axes. The default rotation axis is the object's Z axis. The rotation axis and rotation direction are displayed in the viewport.

Interface

Motor Properties rollout

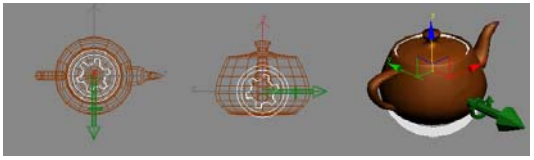


Rigid Body—Displays the name of the rigid body that the motor applies forces to. You assign a motor's rigid body by clicking this button and then selecting a rigid body from any viewport.

Ang Speed—The target angular velocity that the motor attempts to achieve for the rigid body. It is specified in radians per second, so the default value of 3.14 (pi) results in 180 degrees of rotation per second or half a revolution per second.

Gain—The maximum angular impulse that the motor can apply to the rigid body in order to achieve the target velocity.

Rotation Axis—The local axis for the rigid body about which the motor applies impulses. This axis has a visual representation when the motor helper is selected.



Disabled—When on, the motor is not added to the simulation.

Size—The size of the rotation axis in the viewport.

Reset Default Values—Returns Ang Speed, Gain, and Rotation Axis to their default values.

Toy Car

Create panel > Helpers > reactor > Toy Car
Menu bar > reactor > Create Object > Toy Car
reactor toolbar > Create Toy Car button

The reactor Toy Car is a quick way to create and simulate a simple car without having to set up each constraint separately.

The Toy Car helper lets you choose a chassis and wheels for your car, tweak various properties such as the strength of its suspension, and specify whether you would like reactor to turn its wheels during the simulation. reactor then sets up all the necessary constraints to simulate the car.

A Toy Car is simulated if it has a rigid body assigned as a chassis and at least one rigid body assigned as a wheel. When not selected, an invalid Toy Car helper is red in the viewport.

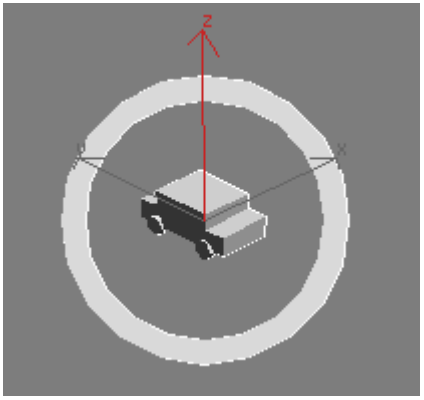
Procedures

To set up a toy car, follow these procedures in order, and make sure the car has a surface to travel on, such as a level plane. For getting started, we suggest you use Icon Orientation and ignore the procedure about Common Local Orientation.

To create a Toy Car helper:

- Choose one of the above menu options, and then click in any viewport to add the Toy Car.

  Note: The icon's initial position has no effect on the car's behavior.



To set up the rigid bodies for a Toy Car:

1. Create and position the rigid bodies you want to connect using the Toy Car helper.

   A toy car must have a chassis and at least one wheel.

2. Use the Rigid Body Properties dialog to assign a mass value to each car element.

   The Mass values you specify affect how the toy car behaves. Ideally, the mass of the chassis should not be much greater than the total wheel mass, and the density of the bodies should be

such that the Analyze World utility generates no warnings.

To assign rigid bodies to a Toy Car:

1. Create the Toy Car helper and the objects to connect using it, as described above.
2. On the Properties rollout, click the Chassis pick button, and then, in one of the viewports, select the object to use as the chassis.

   The Toy Car icon moves to the chosen rigid body. By default, the icon is aligned with the chassis's local space. reactor can use this orientation as a guide for the directions of the suspension and wheel axes.
3. Add your car's wheels to the helper. You can add wheels to the toy car either by picking or using a selection list:
   - Click the Pick button and then select the object to use as a wheel.
   - Click the Add button and then use the dialog to select one or more objects to use as wheels.

   By default, reactor uses the icon orientation (taken from the chassis) to provide the directions of the wheels' spin and suspension axes. Each spin axis is aligned with the icon's X axis and passes through the relevant wheel's pivot point. The suspension axes are aligned with the icon's Z axis, and also pass through each wheel's pivot point.

   Note: For this to work, make sure that the wheels are oriented correctly relative to the icon before simulating.

To change the wheel axis orientation (Icon Orientation):

1. Use the Rotate tool to change the toy car icon's orientation so that it matches that of the actual car model.
2. Ensure that Toy Car Properties rollout > Toy Car Orientation is set to Icon Orientation (the default).

   The wheel axes are aligned with the icon's new orientation during the simulation.

To change the wheel axis orientation (Common Local Orientation):

1. Ensure that the wheel and chassis local spaces are correctly aligned in world space.

   The initial world space orientation of the wheels relative to the chassis/icon is unimportant when using this option, as reactor reorients the bodies so that their local spaces continue to match during the simulation.
2. On the Toy Car Properties rollout, set Toy Car Orientation to Common Local Orientation.
3. Use the radio buttons to change the Suspension and Spin directions relative to the bodies' local spaces. By default, the Spin axes are aligned with the wheel and chassis's local X axes, and the suspension axes with their Z axes.

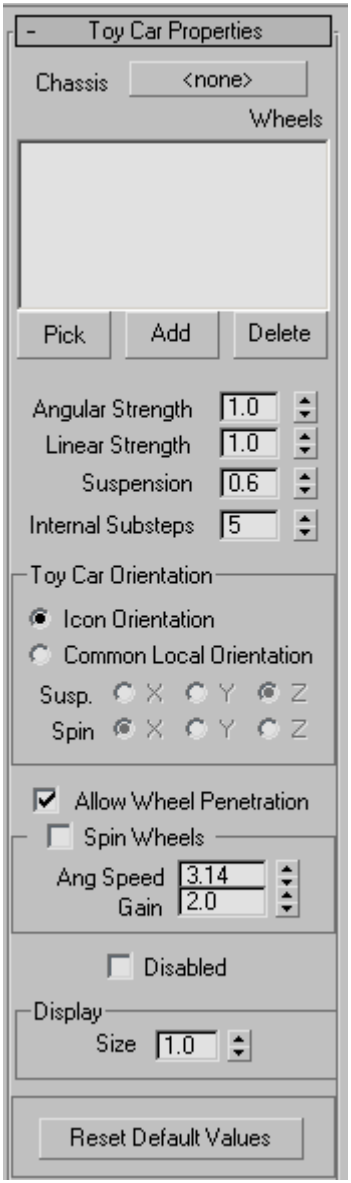
To spin the Toy Car wheels:

1. On the Toy Car Properties rollout, turn on the Spin Wheels check box.
2. Change the default Ang (Angular) Speed and Gain values as necessary.

   Angular Speed is the target angular velocity that reactor attempts to achieve for the wheels. The Gain is the maximum angular impulse that can be applied to the wheels in order to achieve the target velocity.

Interface

Toy Car Properties rollout



Chassis—Displays the name of the rigid body that serves as the chassis. You assign the chassis object by clicking this button and then selecting a rigid body from any viewport.

Wheels—Lists the names of the toy car's wheels.

Pick—Lets you add a wheel to the car. Click this button, and then in the viewport position the cursor over the object you would like to use as a wheel. If the object can be used as a rigid body, the cursor changes from an arrow to a cross and you can select the object to add it to the toy car.

Add—Lets you add one or more rigid bodies from the scene to act as wheels for the toy car. Click the button to open the Select Car Wheels dialog. Make a selection in the provided list, then add the wheels to the car by clicking the Select button.

Delete—Lets you remove wheels from the car. Highlight the wheels to remove from the car in the Wheels list and then click this button.

Note: The following four parameters are available with Havok 1 only. With Havok 3, use the *Havok 3 Wheel Params rollout (page 2–769)* settings instead.

Angular Strength—The strength of the corrective forces applied to keep the wheels upright and pointing forward.

Linear Strength—The strength of the forces applied to keep the wheels in position relative to the chassis.

Suspension—The suspension strength for the car. The lower the value, the more the chassis bounces when forces are applied to it.

Internal Substeps—Allows the car to be simulated using more substeps per keyframe than the rest of the simulation.

Increasing the Internal Substeps value can be useful if you are simulating a car with unusual parameters, which might require more substeps for stability.

Toy Car Orientation group

Choose how reactor orients the bodies in relation to the spin and suspension axes.

Icon Orientation—The Toy Car icon's initial orientation provides the directions of the wheels' spin and suspension axes. Each spin axis is aligned with the icon's X axis and pass through the relevant wheel's pivot point. The suspension axes are aligned with the icon's Z axis, and also pass through each wheel's pivot point. The wheels' own local orientations are unimportant.

Note: Animating the orientation of the Toy Car icon has no effect on the direction of motion during the simulation; reactor uses only the orientation at the start of the simulation.

Common Local Orientation—With this option, you align the bodies' local spaces in world space, and specify the Spin and Suspension axes relative to this common orientation. The orientation of the rigid bodies in world space before simulating is unimportant.

This means, for instance, that you can set up your car with its wheels turned sideways and they still spin around the correct axes when simulating. This option allows you to change the spin and suspension axes separately. By default, the spin axes are aligned with the wheel and chassis' local X axes, and the suspension axes with their Z axes.

Allow Wheel Penetration—When on, reactor disables collision detection between the chassis and the wheels.

Spin Wheels—When on, the car's wheels to spin during the simulation and an arrow appears on the Toy Car viewport icon showing the direction of motion.

- **Ang Speed—**The target angular velocity that reactor attempts to achieve for the wheels when Spin Wheels is on. It is specified in radians per second, so the default value of 3.14 (pi) results in 180 degrees of rotation per second or half a revolution per second.
- **Gain—**The maximum angular impulse that reactor can to the wheels in order to achieve the target velocity.

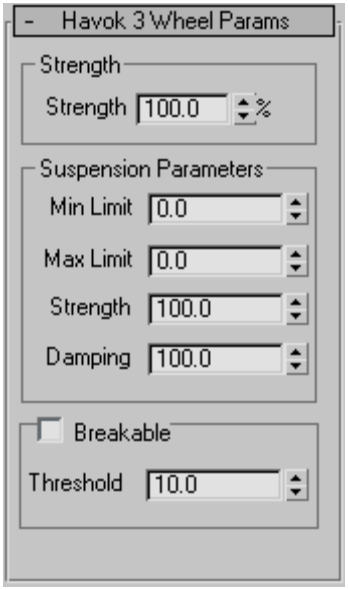
Disabled—When on, reactor removes the Toy Car helper from the simulation.

Display Size—The size of the Toy Car helper icon as displayed in the viewport.

Reset Default Values—Returns the Toy Car parameters to their default values.

Havok 3 Wheel Params rollout

These parameters apply only when the simulation engine is Havok 3, and reflect the implementation of a subset of Car-Wheel constraint parameters in the Toy Car object in Havok 3.



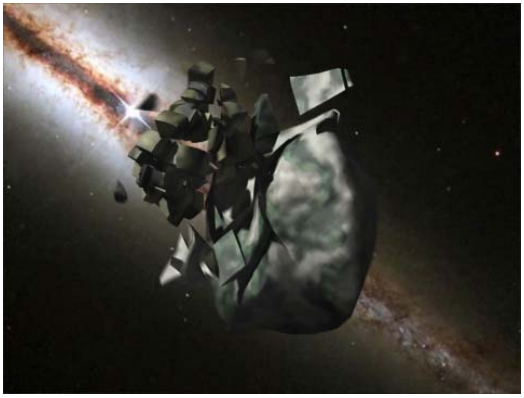
These parameters replicate the *Car-Wheel constraint parameters of the same names (page 2–761)*.

Fracture 

Create panel > Helpers > reactor > Fracture

Menu bar > reactor > Create Object > Fracture

reactor toolbar > Create Fracture button



The Fracture helper object simulates the breaking of a rigid body into a number of smaller pieces as the result of an impact. To do this, you need to supply the pieces that are glued together to create the whole object. This is known as "pre-fracturing." reactor does not break the pieces you supply into smaller pieces.

Tip: An excellent tool for pre-fracturing your objects is the *ProCutter compound object (page 1–388)*.

Rigid bodies that are part of a Fracture helper move together as a single body. When a rigid body that belongs to a Fracture helper collides with another body, the collision information is analyzed and if a threshold is exceeded, the rigid body is removed from the fracture helper. Once the rigid body has been removed, it moves independently of the fracture object and is free to collide with the rigid bodies that are still part of the fracture object. This can often result in a chain reaction of pieces breaking off.

You can optionally enable the use of connectivity. Connectivity analyzes the bodies belonging to the Fracture helper to find bodies that are in contact with each other before simulation occurs. This builds up an internal graph of which bodies are connected to which other bodies. Thus groups of rigid bodies that are connected to each other, but disconnected with other bodies in the fracture helper, can move independently as separate fracturable objects. As a result, you can get the effect of multiple fracturable objects in a single Fracture helper.

After a fracture event occurs, reactor once again analyzes this connectivity graph to create new chunks if they exist. If you see floating disconnected bodies moving strangely because they are invisibly still part of a larger fracture object, be sure to turn on connectivity. A body is considered to be in contact with another body if the closest distance between them is less than the collision tolerance.

To allow the fracture feature to work reasonably well, reactor uses a special collision detection technique: penetration depth calculation. Unlike normal rigid bodies in reactor, bodies that are part of a fracture helper are allowed to continue to collide with each other, even when they are in a interpenetrating state. To see how the non-fracture bodies are treated during collision detection, see the *Fracture Tips (page 2–773)* topic. Penetration depth calculation, especially in the context of the simulation paradigm used in reactor, is computationally expensive, so you might notice slower performance when the fracture bodies are penetrating each other. Stability in these cases can also become an issue.

reactor provides a number of techniques for getting your fracture simulation to behave more realistically If you find that it is unstable. For instance, the fractured object might appear to

explode after a single piece breaks loose. You can find these in the *Fracture Tips (page 2–773)* topic.

Procedures

To create a Fracture helper:

- Choose one of the above menu options, and then click in any viewport to add the Fracture helper.

  Note: The icon's position has no effect on the fracture object's behavior.

To add rigid bodies to a Fracture helper:

1. Create a Fracture helper in the scene and some objects to use as rigid bodies.

   You can add rigid bodies to the helper in two ways: by picking or by using a selection list.

2. With the Fracture helper selected, use either (or both) of the following methods:

   - Click the Pick button on the Properties rollout and then select the select the rigid body to add to the helper.
   - Click the Add button and then use the dialog to select one or more objects to use as fracture elements.

   The names of rigid bodies you add appear in the helper's Pieces list.

To create the Fracture helper and add objects in a single step:

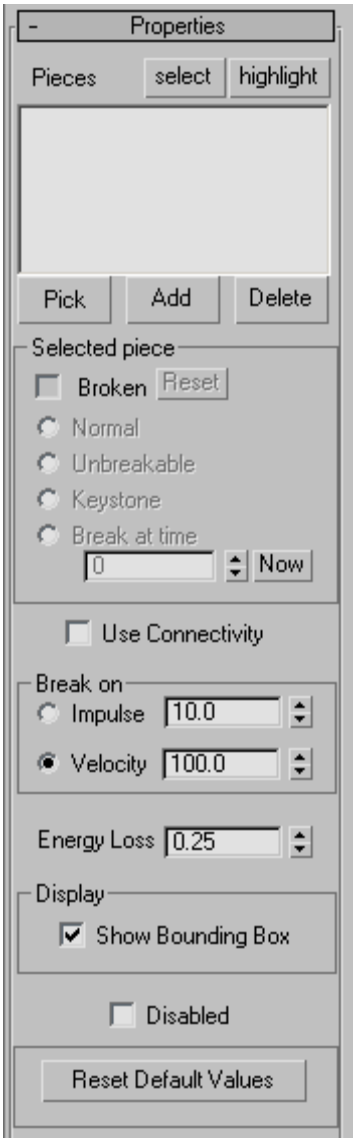
- See the *Scripts (page 2–817)* topic.

To prevent the Fracture object from exploding:

- See the *Fracture Tips (page 2–773)* topic.

Interface

Properties rollout



Highlight—Click this button to cause the objects in the Pieces list to momentarily display as if selected.

Pieces—Lists the names of the objects currently part of the Fracture object. To highlight multiple contiguous list items, click and then Shift+click.

To highlight multiple non-contiguous list items or remove highlighting, use Ctrl+click.

Pick—Lets you add an object to the Fracture helper. Click this button, and then in the viewport position the cursor over the object to add to the helper. If the object can be used as a fracture piece, the cursor changes from an arrow to a cross and you can select the object to add it to the helper.

Add—Lets you add one or more objects from the scene to the helper. Click the button to open the Select Rigid Bodies For Fracture dialog. Make a selection in the provided list, and then press the Select button to add the objects to the helper.

Delete—Lets you remove objects from the helper. In the Pieces list, highlight the bodies to remove from the helper and then click this button.

—

Selected piece

You can assign these properties to pieces highlighted in the Pieces list.

Broken—This read-only check box shows whether any of the highlighted pieces have broken off at the current frame. For example, say you've created an animation from frame 0 to frame 50 using Fracture, where piece A broke off at frame 10. Highlighting piece A and moving the time slider to frame 10 or after automatically turns on Broken. Moving the time slider to frame 9 or before turns off Broken. This is because reactor keeps track of the times pieces break off. That way, if you continue the animation at frame 50, reactor knows it shouldn't add piece A to the fracture object, so it keeps moving independently. You can restore a highlighted piece to its "unbroken" state with the Reset button (see following).

Reset—This button is available if any of the pieces broke off during a previous animation (reactor stores the times when pieces broke off). Clicking this button tells reactor to ignore the times they broke off in the previous animation.

Choose the behavior for highlighted pieces:

- **Normal—**The default behavior option for a fracture piece.
- **Unbreakable—**The highlighted piece never breaks away from the Fracture object, even if it experiences an intense collision.
- **Keystone—**If this piece breaks, then all pieces break.
- **Break At Time —**The selected piece breaks at the time you specify. For this kind of piece the Broken check box reflects the state of the piece at the current frame. Thus, if you set a piece to break at frame 5, the Broken check box is on at frame 5 and later, and off at frames before 5.

Use Connectivity —When on, groups of pieces that are connected to each other move independently of pieces that they are not connected to. This is like having multiple fracture objects contained in one fracture helper. The collision detector determines that two pieces are connected to each other if their distance from each other is less than the world's *Collision Tolerance*.

Break on

Choose the method by which reactor determines whether fracturing should occur:

- **Impulse—**When a piece of the fracture object experiences a collision where the impulse is greater than the threshold specified, it breaks off from the fracture object. The impulse of the collision is proportional to both the relative velocity of the collision and the mass of the objects involved in the collision. The higher the mass of the object that hit the piece, the larger the impulse.
- **Velocity—**When a piece of the fracture object experiences a collision where the relative velocity is greater than the threshold specified,

it breaks off from the fracture object. The size or mass of the object is not be taken into account.

Energy Loss—The amount of extra kinetic energy is lost in the collision due to the breaking of the fracture bonds. This has the effect of dampening the collision somewhat and also transferring energy to the remaining, unbroken pieces of the fracture object. To simulate a brittle object, with maximum shatter effect, keep this value low. To simulate breakage of a material more like wood or concrete, use a high Energy Loss value.

Tip: Increasing this value can improve the performance and stability of the simulation.

Display: Show Bounding Box —When on, the viewports display a box that contains all the pieces of the fracture object.

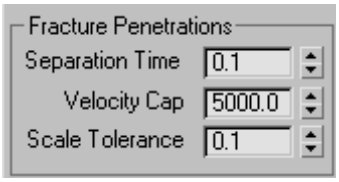
Disabled —When on, the Fracture helper has no influence on the rigid bodies that are part of it.

Reset Default Values—Returns the settings to their default values.

World rollout > Fracture Penetrations group

Several additional values govern fracture behavior. These are global in that they apply to every rigid body added to a Fracture object, and are found in the Utility panel > reactor > World rollout.

These values control the behavior of penetration depth collision detection and response, and apply only to objects that are included in a Fracture helper.



Separation Time—Forces are applied to penetrating bodies that are strong enough so the bodies become separated within the time specified. To simulate an explosion effect, make this value very small and reduce the Velocity Cap value.

Velocity Cap—reactor won't apply penetration recovery forces if they cause the relative velocity between two penetrating bodies to exceed the specified value.

Scale Tolerance—The collision tolerance for bodies that use the penetration depth algorithm is determined by multiplying the *Collision Tolerance value* by the value specified. If this value is negative, the collision tolerance will extend inside the bodies, effectively shrinking the collision geometry.

For more information about using these values, see the *Fracture Tips (page 2–773)* topic.

Fracture Tips

After you procedurally slice up an object for fracture, the new pieces fit snugly among each other. When one piece breaks loose after a fracture event, it might be in contact with other pieces that are still part of the non-fractured body. This can throw the broken piece into unstable oscillations against two or more opposing pieces. The result is a chain-reaction throughout the fracture body, with all the pieces flying off into space. It might appear as though some of the pieces simply wink out of existence. Unlike normal rigid bodies in reactor, Fracture pieces are allowed to exist in a state of interpenetration, where restoring forces are applied to separate them. Many objects sitting in penetrating states next to each other can result in an unstable system.

To make the system more stable and less prone to exploding:

Try one or more of the following:

- Reduce the *Scale Tolerance value* in the Fracture Penetrations group of the reactor utility World

rollout. This value can be as low as –1.0, but –0.1 should suffice in most cases. This scales the *Collision Tolerance (page 2–808)* for pieces of a fracture object. A negative Scale Tolerance value effectively shrinks the object as it is perceived by the collision-detection system. Pieces that were snug against each other now have a comfortable safety zone where no collisions take place. This is tied to the Collision Tolerance value, so changing Collision Tolerance changes the effective tolerance used when fracture pieces collide with other fracture pieces. It dramatically reduces unstable oscillations. This is the most effective thing to try when you wish to avoid explosions.

- Manually offset the pieces from each other, leaving a physical gap between all pieces. The gap should be as big as Collision Tolerance * Scale Tolerance.
- Increase the Energy Loss value on the Fracture object's Properties rollout. This is the percent of energy that can be "lost" in the fracture event by transferring momentum from the broken piece to the unbroken Fracture object. Increasing it increases the relative velocity between the newly fractured piece and the remaining unfractured pieces. Ultimately this reduces the chances of a chain-reaction of unstable oscillations.
- Increase the Separation Time value and/or decrease the Velocity Cap in the *Fracture Penetrations (page 2–773)* group of the World rollout. This controls how aggressively reactor pushes the pieces out of a state of interpenetration. After doing this you may notice visible interpenetrations.
- Slice up your object in a different configuration. Avoid creating very small pieces, especially if they are right next to very large pieces.
- Turn on the Use Connectivity check box on the Properties rollout of the Fracture helper. When connectivity is enabled, clumps of pieces that are disconnected from other clumps but still internally glued to each other are set free to move independent of the other clumps. This serves to create gaps between fractured pieces earlier, so unstable oscillations won't happen.

Storing and Accessing Collisions

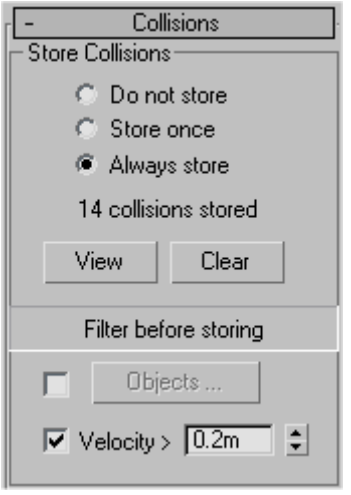
Utility panel > reactor > Collisions rollout

reactor lets you store information about all rigid body collisions that occur during the simulation. You can access this information via MAXScript or save it to a text file. The information includes the objects involved, the point of collision, and the relative velocity during collision. You can then use it to generate particles or other effects in your animation when objects collide, trigger sounds, and so on.

You can activate and configure collision storing using the reactor utility's *Collisions rollout (page 2–810)*.

Interface

Store Collisions group



Store Collisions—Use these options to store collision information when creating an animation. For each collision that occurs during the simulation, reactor can record the simulation time at which the collision occured, the objects involved, the point of collision, and the relative velocity during the collision.

- **Do not store—**reactor stores no collision data.
- **Store once—**reactor stores collision data next time an animation is created.
- **Always store—**reactor always stores collision data when creating an animation.

 # collisions stored—This read-only field reports how many collisions were stored in reactor last time an animation was created.

View—Opens the *Stored Collisions dialog (page 2–775)*, which shows all the currently stored collision information.

Clear—Deletes all stored collision information.

Filter Before Storing

These options let you provide more specific details about the collision information you want to record.

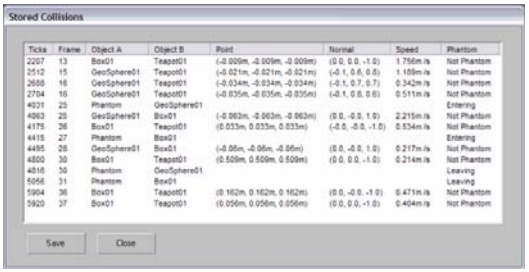
Objects—When on, lets you store only collisions involving particular objects. Use the Objects button to designate those objects.

Click this button to open the *Filter Bodies In Collisions dialog (page 2–776)*. This allows you to select the pairs of objects whose collisions interest you. Only collisions involving those pairs of objects are stored.

Velocity —When on, reactor stores only collisions that occur above a certain velocity. You can specify the threshold velocity using the provided field. reactor ignores this filter for *phantom collisions*.

Define Collision Pairs—Allows you to enable or disable collision detection between specified pairs of objects. You can find out more about this in the *Collisions Rollout (page 2–810)* topic.

Stored Collisions Dialog



Collisions List Box—This list box shows all the information stored about the collisions ocurred during the last animation. The information is organized in different columns:

- **Ticks—**The time, in ticks (1/4800th of a second) when the collision happened.
- **Frame / SMPTE / MM:SS:Ticks—**The time, in the current time units (which can be frames, ticks, SMPTE, etc.), when the collision happened.
- **Object A—**The first rigid body involved in the collision.
- **Object B—**The second rigid body involved in the collision.
- **Point—**The location of the collision, in world coordinates.
- **Normal—**The unit normal of the collision, specified from B to A.
- **Speed—**The relative velocity of the two objects along the specified normal direction.
- **Phantom—**Specific information generated by *phantom rigid bodies*. There are three possible values:
  - **Not Phantom—**This collision wasn't generated by a phantom rigid body.
  - **Entering—**The time the phantom rigid body started penetrating (entered) the other rigid body.

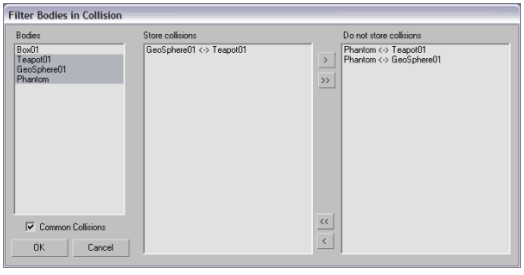
- **Leaving**—The time the phantom rigid body finished penetrating (exited) the other rigid body.

Close—Closes the dialog.

Save—Allows you to save the stored information in a text file. This is an example of the format:

```
....
Time : 4175
  A : Box01
  B : Teapot01
  Point : ( 3.25862 -56.7632 10.0549 )
  Normal : ( 1.28512e-007 3.18186e-007 -1 )
  NRV : 53.4184
  Phantom : Not PhantomTime : 4415
  A : Phantom
  B : Box01
  Point : ( 0 0 0 )
  Normal : ( 0 0 0 )
  NRV : 0
  Phantom : EnteringTime : 4495
  A : GeoSphere01
  B : Box01
  Point : ( -6.02472 18.1558 10.1478 )
  Normal : ( -2.48431e-006 1.51463e-006 1 )
  NRV : 21.6887
  Phantom : Not Phantom
....
```

Filter Bodies In Collisions dialog



Bodies—Lists the names of the reactor bodies in the scene. Highlighting an object populates the Store Collisions and Do Not Store Collisions lists for that object with respect to the other bodies in the scene. If you highlight more than one object in this list, how the lists are populated depends on the value of Common Collisions.

Common Collisions—When on, if more than one rigid body is highlighted in the Bodies list, the Store and Do Not Store lists are populated from the list of possible pairs that can be made using only your selected objects. When off, the lists contain all possible pairs from the scene that contain at least one of your selected objects.

For example, take a scene containing four rigid bodies: Box01, Box02, Box03 and Box04. If Common Collisions is on, and you highlight Box01 and Box02 in the Bodies list, then the only possible pair that includes only your selected objects is Box01<->Box02.

If you highlight Box01, Box02, and Box03, then the possible pairs using these objects are Box01<->Box02, Box02<->Box03, and Box01<->Box03. If highlight the same objects with Common Collisions off, then the lists also include the highlighted objects paired with the remaining objects in the collection. This means that the available pairs would also now include Box01<->Box04, Box02<->Box04 and Box03<->Box04.

This option is selected by default.

Store Collisions—Collision data is stored for any pairs of objects in this list.

Do Not Store Collisions—Collision data is not stored for these pairs of objects.

>—Moves the pairs that are highlighted in the Store list to the Do Not Store list.

>>—Moves all the pairs in the Store list to the Do Not Store list.

<—Moves the pairs that are highlighted in the Do Not Store list to the Store list.

<<—Moves all the pairs in the Do Not Store list to the Store list.

OK—Closes the dialog, saving any changes.

Cancel—Closes the dialog without saving your changes.

Deformable Bodies

You can use *rigid bodies (page 2–716)* in reactor to model any real-world object whose shape doesn't change over time. However, what if you want to simulate an object whose geometry does change over the course of the simulation, such as a cloak, hair, foam bricks, or perhaps a slithering tentacle? reactor allows you to model these with a second category of objects, called *deformable bodies*. The geometry (vertices) of deformable bodies can change over time, driven either by reactor during the simulation or by existing animation in 3ds Max, allowing the objects to bend, flex, and stretch while affecting and being affected by the rest of objects in the world simulation.

Important: The Havok 3 engine does not support deformable bodies in reactor; you can use these only with the Havok 1 engine.

In addition to creating entirely deformable objects, you can combine deformable bodies with rigid bodies, for instance to add secondary motion to a simulated character. Secondary motion could include swirling clothing, wobbling flesh, or a swinging tail. Deformable objects are also useful for environmental effects such as swinging ropes and chains, curtains, and flags with dynamic wind.

Generally, you create a deformable body in reactor by first creating a mesh or spline that models the object's basic shape, and then applying a special modifier. You can then specify additional physical properties for the object. reactor includes four main types of deformable bodies, each of which is dealt with in its own section:

- *Cloth (page 2–778)*, a deformable two-dimensional triangular mesh
- *Soft Body (page 2–783)*, a deformable three-dimensional closed triangular mesh
- *Rope (page 2–789)*, a deformable one-dimensional chain of vertices.
- *Deforming Mesh (page 2–793)*, a deformable mesh whose vertices have already been animated, for instance the skin on a character rig.

As with rigid bodies, you must add deformable bodies to a collection in order to be added to the simulation. Each deformable body type has its own corresponding collection type.

Also like rigid bodies, you can constrain the possible movement of deformable bodies: see *Constraining Deformable Bodies (page 2–795)*.

Note: Deformable bodies are more complex than *rigid bodies (page 2–716)*, so reactor requires more time to simulate them.

Working with Deformable-Body Vertices

reactor includes several tools for working with individual vertices in deformable bodies. The topics covering these are:

- *Fixing Vertices in World Space (page 2–796)*
- *Keyframing Vertices (page 2–797)*
- *Attaching Vertices to a Rigid Body (page 2–798)*
- *Attaching Vertices to Deforming Meshes (Skin) (page 2–799)*
- *Soft Selection (page 2–800)*

Cloth



Cloth objects in reactor are two-dimensional deformable bodies. Using cloth objects you can simulate flags, curtains, clothing (skirts, capes, shirts), banners and even materials like paper and sheet metal.

This section contains two topics:

- *Cloth Modifier (page 2–778)*: Apply it to each mesh to be deformed as cloth.
- *Cloth Collection (page 2–781)*: A container for cloth objects.

Cloth Modifier

Menu bar > reactor > Apply Modifier > Cloth Modifier

reactor toolbar > Apply Cloth Modifier button

The Cloth modifier lets you to turn any geometry into a deformable mesh, allowing you to simulate the behavior of objects such as curtains, clothes, metal sheets, and flags. You can specify a number of special properties for cloth objects, including stiffness and how the object folds.

To add a cloth object to the simulation, you need to add it to a *Cloth Collection (page 2–781)*.

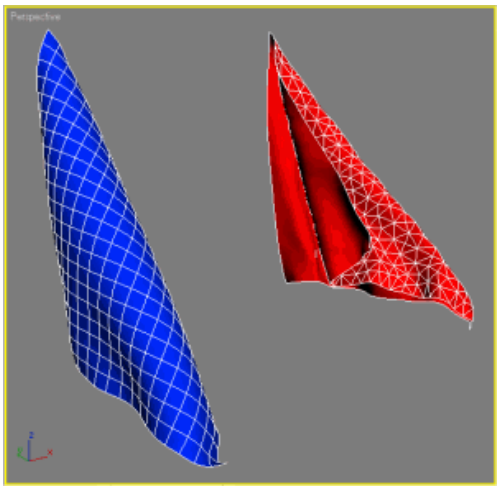
Procedures

To create a Cloth object:

1. Create the mesh object that you would like to simulate as cloth.

   Note: In reactor the underlying topology of this object can influence the cloth's behavior. For instance, cloth tends to fold along contiguous edges; highly tessellated meshes also stretch more, etc. Irregular triangulations (for example, a Delaunay triangulation of a 3ds Max NURBS surface) results in isotropic behavior (same behavior across all directions), avoiding artificial creases and folds around specific directions; it can therefore produce more realistic-looking results for pieces of clothing. Regular triangulations (like that of a standard 3ds Max plane) lead to anisotropic behavior (tendency to folding and creasing is different depending on the direction), this of course might be irrelevant or even desirable in some situations.

   Tip: A quick, easy way to create a Delaunay mesh for reactor Cloth is to apply the *Garment Maker modifier (page 1–607)* to a shape.



Different tessellation produces different cloth behavior.

2. With the object selected, choose one of the above options.

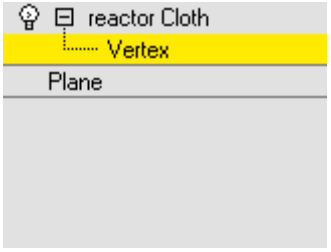
   The Cloth modifier appears in the object's modifier stack.

To edit a cloth object's physical properties:

1. Select the appropriate object in the scene.
2. On the Modify panel, ensure that the reactor Cloth modifier is highlighted in the modifier stack.
3. Use the Properties rollout to edit the properties, which are described in the Interface section, following.

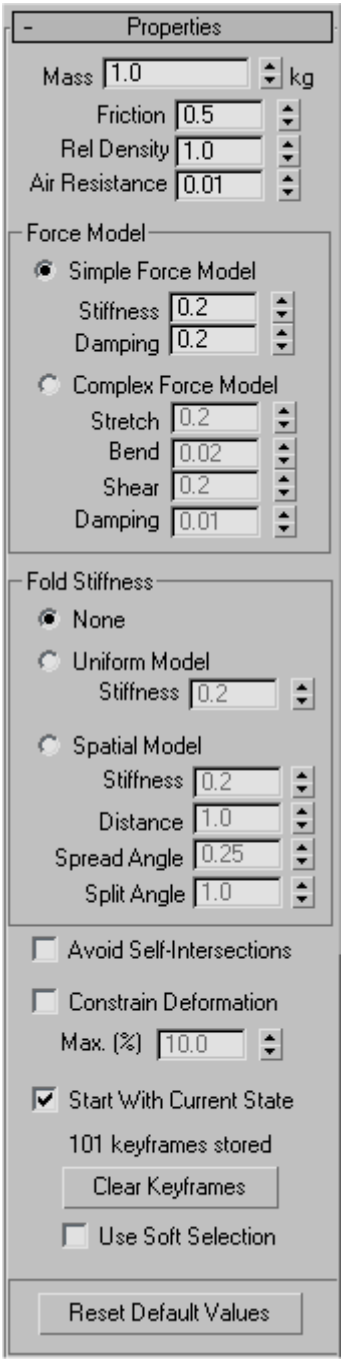
Interface

Modifier Stack



Vertex—Allows you to select individual vertices of the cloth object to apply and modify *deformable constraints (page 2–795)*.

Cloth Properties rollout



Mass—The Cloth object's mass in kilograms. This affects the behavior during collisions with other objects, and also the stretching caused by *attached rigid bodies (page 2–798)*: the higher the mass of the cloth, the less stretching induced by the rigid body.

Friction—The coefficient of friction for the cloth's surface. As with rigid bodies, this affects how smoothly the cloth moves relative to surfaces it contacts. The Friction values for both objects in contact combine to produce a coefficient for the interaction.

Rel Density—Cloth has no volume, so reactor cannot calculate its density. Instead, reactor provides a buoyancy property for Cloth objects, which reflects its relative density. The default value is 1.0, the density of water. This value is important only if you are going to float or immerse the cloth in water.

Air Resistance—The extent to which the cloth loses energy as it moves.

Force Model—Choose the way forces in the simulation affect your cloth object:

- **Simple Force Model—**This default method is suitable for most situations.
  - **Stiffness—**The stiffness of the cloth.
  - **Damping—**How quickly the cloth dissipates energy as it changes shape.
- **Complex Force Model—**This more-accurate model of cloth dynamics is more computationally expensive to simulate. It provides parameters to specify the shear and stretch stiffness separately, as well as a physically accurate out-of-plane bend property.
  - **Stretch—**The cloth resistance to stretching.
  - **Bend—**The cloth resistance to bending.
  - **Shear—**The cloth resistance to shearing.
  - **Damping—**How quickly the cloth dissipates energy as it changes shape.

Fold Stiffness —Fold stiffness governs the cloth resistance to folding.

- **None—**By default, Cloth objects have no fold stiffness, and behave like a loose material such as silk.
- **Uniform Model—**This model adds fold stiffness uniformly across the surface of the cloth, regardless of its topology.
  - **Stiffness—**The fold stiffness value.
- **Spatial Model—**This more-complex model of fold stiffness can be useful if you are simulating a complicated piece of cloth, such as a dress. It allows you, for instance, to add fold stiffness only to sections of the cloth that are flat.
  - **Stiffness—**The fold stiffness value.
  - **Distance—**The degree of fold stiffness per unit area. A Distance value of 2 refers to twice the average triangle edge length in the cloth, so a large Distance value results in a large concentration of fold stiffness per unit area.
  - **Spread Angle—**Controls the degree to which fold stiffness is added to cloth when it has a non-flat shape. The angle relates to the threshold angle between two parts of the cloth, below which stiffness is added. If this value is 0.0, fold stiffness is added only to flat parts of the cloth.
  - **Split Angle—**The degree to which fold stiffness is added; specifically, how fold stiffness is concentrated along the grid lines of the cloth's original mesh.

Avoid Self-Intersections—When on, the cloth will not intersect with itself during the simulation. This results in a more realistic-looking simulation, but can increase simulation time.

Constrain Deformation—When on, limits the extent to which the cloth can stretch.

- **Max—**A percentage value that indicates how stretchy the cloth is: The lower the value, the less stretching reactor allows.

Start With Current State—When on, the Cloth object starts the simulation using the current state stored in the modifier. This can be useful if you have, for example, draped the cloth around something in the Preview Window and then updated the viewport using Update MAX. When off, the cloth starts with the state it had originally, below the modifier. Available only when keyframes are stored.

 # Keyframes Stored —This read-only shows the number of keyframes, if any, stored for the Cloth object. reactor stores keyframes for the object if you create a reactor animation, or if you use Update MAX in the Preview Window.

Clear Keyframes—Deletes any stored keyframes for the Cloth object.

Use Soft Selection—Allows you to use soft selection to smooth the transition between keyframed and simulated vertices for this deformable object. For more information, see *Soft Selection (page 2–800)*.

Reset Default Values—Resets the values for the Cloth object to their defaults.

Constraints rollout



These options allow you to create different types of deformable constraints for the body. For more information, see *Constraining Deformable Bodies (page 2–795)*.

Cloth Collection

Create panel > Helpers > reactor > CLCollection
Menu bar > reactor > Create Object > Cloth Collection
reactor toolbar > Create Cloth Collection button

The Cloth Collection is a reactor helper object that acts as a container for Cloth objects. Once you have added a Cloth Collection to your scene, you can add Cloth objects (object with a *Cloth modifier (page 2–778)*) in the scene to the collection.

When you run the simulation, reactor examines the cloth collections in the scene, and, provided

that the collections are not disabled, adds the cloth objects they contain to the simulation.

Procedures

To create and use a Cloth Collection:

1. Create some Cloth objects by applying the reactor Cloth modifier to standard objects.
2. Choose one of the above commands, and then click in any viewport to add the Cloth Collection.

   Note: The icon's position has no effect on the collection's behavior.

   The collection icon is added to the scene. You can add Cloth objects to the collection in two ways: by picking or by using a selection list.
3. Add objects to the collection using either (or both) of these methods:
   - On the Properties rollout, click the Pick button, and then in the viewport click an object to add.
   - On the Properties rollout, click the Add button, and then use the Select Cloths dialog to specify one or more objects to add.

To create the collection and add objects in a single step:

- See the *Scripts (page 2–817)* topic.

Interface

Cloth Collection Properties rollout



Highlight—Causes the objects in the collection to flash momentarily in the viewports.

Cloth Entities—Lists the names of the objects currently in the Cloth Collection.

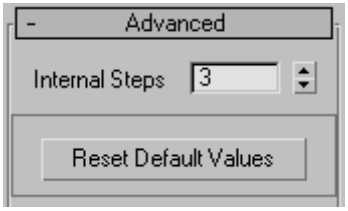
Pick—Adds an object to the cloth collection. Click this button, then in the viewport move the cursor over an object to add to the collection. If it is a Cloth object, the cursor changes from an arrow to a cross and you can select the object to add it to the collection.

Add—Adds one or more objects from the scene to the collection. Click the button to open the Select Cloths dialog. Highlight one or more entries in the list, and then click the Select button to add the objects to the collection.

Delete—Removes objects from the collection. In the Cloth Entities list, highlight the bodies to remove from the collection and click this button.

Disabled—When on, the collection and the bodies it contains are not added to the simulation.

Advanced rollout

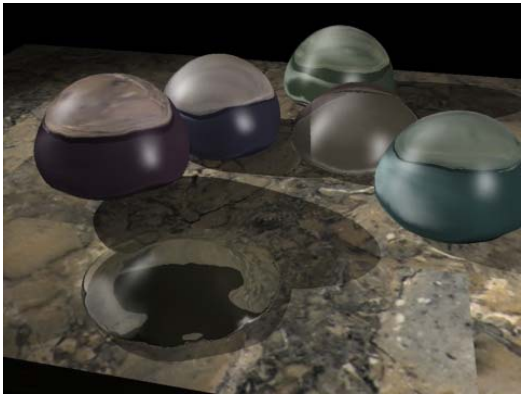


Internal Steps—Because deformable objects are more complex to simulate, it is usually necessary to perform more *simulation steps (page 2–709)* to increase stability. This parameter specifies the number of steps reactor performs internally to simulate objects in this collection for each simulation substep taken globally during the simulation. For example, at 60 fps, simulating at one keyframe (step) per frame, 10 substeps, at three internal substeps for this collection, reactor simulates the objects in this collection using 60x10x3=1,800 steps per second, while the simulation in general uses 60x10=600 steps every second. This allows the objects in this collection to be simulated more accurately without slowing down the rest of the simulation. In general terms, you can think of it in this way:

- Precision (stability) of overall simulation=Number of Substeps (as specified on the *Preview and Animation rollout (page 2–806)*)
- Precision (stability) of objects in this collection=(Precision of overall simulation)x(Internal Substeps)

Reset Default Values—Restores Internal Steps to its default value.

Soft Bodies



Soft bodies are three-dimensional *deformable bodies (page 2–777)*. Like *Cloth (page 2–778)* objects, they modify meshes. The main difference between cloth as soft bodies is that soft bodies have a notion of shape : a soft body tries, to some extent, to keep its original shape.

You can use soft bodies to simulate squashy objects like a partially deflated beach ball, a water bottle, jelly, or fruit. They are also useful for adding realistic secondary motion to your objects and characters: floppy ears, noses, tails, etc.

reactor provides two methods for simulating soft bodies:

- The mesh-based method uses the vertices in the mesh to operate (as with cloth and rope).
- *FFD soft bodies (page 2–786)* manipulate the control points in an FFD grid instead.

Depending on the complexity of your objects and the desired effect, you can use either method. The following sections go into more detail about them.

- *Soft Body Modifier (page 2–784)*: Apply it to each mesh to be deformed as soft body.
- *FFD-based soft bodies (page 2–786)*: Soft bodies deformed by an FFD modifier

- *Soft Body Collection (page 2–788)*: A container for soft bodies.

Soft Body Modifier

Modify panel > Modifier List > reactor SoftBody
Menu bar > reactor > Apply Modifier > Soft Body Modifier
reactor toolbar > Apply Soft Body Modifier button

The Soft Body modifier lets you turn a rigid body into a closed, deformable, 3D triangular mesh, thus creating objects that can be flexed, bent, and squashed during the simulation. You can specify physical properties for soft bodies, including stiffness, mass, and friction.

reactor provides two types of soft body: mesh-based soft bodies, where reactor deforms the underlying mesh directly; and Freeform Deformation (FFD) bodies, in which reactor encases the original mesh in a simpler lattice that it uses to update the mesh. For more information about creating FFD bodies, see *FFD Soft Bodies (page 2–786)*.

To add a soft body to the simulation, you need to add it to a *Soft Body Collection (page 2–788)*.

Procedures

To create a soft body:

1. Create the mesh to use as the basic, non-deformed state of your soft body. You can turn any mesh with a fixed number of vertices into a soft body.
2. With the object selected, choose one of the above commands.

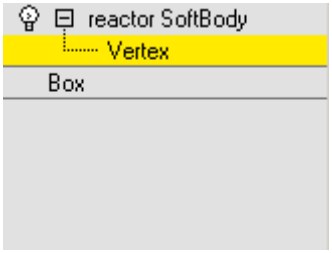
   The reactor SoftBody modifier appears in the object's modifier stack.

To edit a soft body's physical properties:

1. Select an appropriate object in the scene.
2. On the Modify panel, ensure that the reactor SoftBody modifier is highlighted in the modifier stack.
3. Use the Properties rollout to edit the properties, which are described in the Interface section, following.

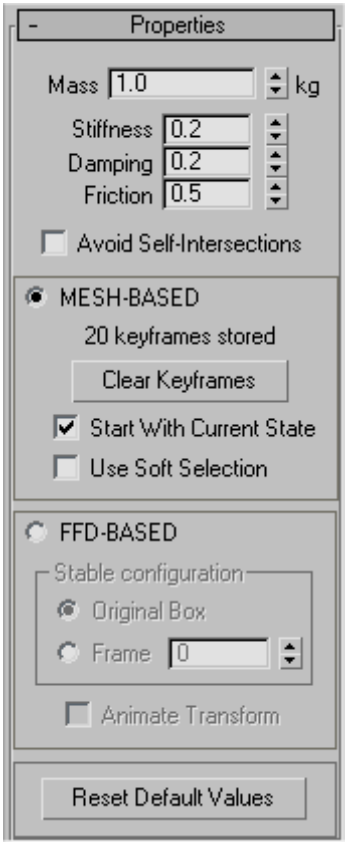
Interface

Modifier Stack



Vertex—Allows you to select individual vertices of the soft body to which to apply and modify *deformable constraints (page 2–795)*.

Soft Body Properties rollout



Mass—The mass of the soft body in kilograms. This affects the soft body behavior upon colliding with other objects, its buoyancy when interacting with water, and the stretching caused by *attached rigid bodies (page 2–798)*: the higher the mass of the soft body, the less stretch induced by the rigid body.

Stiffness—The stiffness of the soft body: the stiffer it is, the harder it is to deform.

Damping—The damping coefficient for the oscillation of the soft body's compression and expansion.

Friction—The coefficient of friction for the soft body's surface. As with rigid bodies, this affects how smoothly the soft body will move relative to surfaces it's in contact with. The friction values for both objects are combined to produce a coefficient for the interaction.

Avoid Self-Intersections—When on, the soft body does not intersect with itself during the simulation. This results in a more realistic simulation, but can increase simulation time.

[soft body type]—

- **Mesh-Based—**By default, soft bodies are mesh-based - the modifier directly modifies the underlying mesh. This option is suitable for most simple objects, such as balls and bricks. For more complex meshes (greater than 200 triangles), however, this approach can be time-consuming and may slow down the simulation. A good alternative is to use *FFD soft bodies (page 2–786)*.
  - **# keyframes stored—**This indicates the number of keyframes stored for the soft body, if any. reactor stores keyframes for the object if you create a reactor animation, or use Update MAX in the Preview Window.
  - **Clear Keyframes—**Clears any stored keyframes for this soft body.
  - **Start With Current State—**The soft body starts the simulation using the current state stored in the modifier. This can be useful if you have, for example, deformed the soft body in the Preview Window and then updated the viewport using Update MAX. When off, the soft body will start with the state it had originally (below the modifier).
  - **Use Soft Selection—**Allows you to use soft selection to smooth the transition between keyframed and simulated vertices for this deformable object. For more information, see *Soft Selection (page 2–800)*.
- **FFD-Based—**Uses the FFD version of the Soft Body modifier. For more information about

FFD bodies, including descriptions of these parameters, see *FFD Soft Bodies (page 2–786)*.

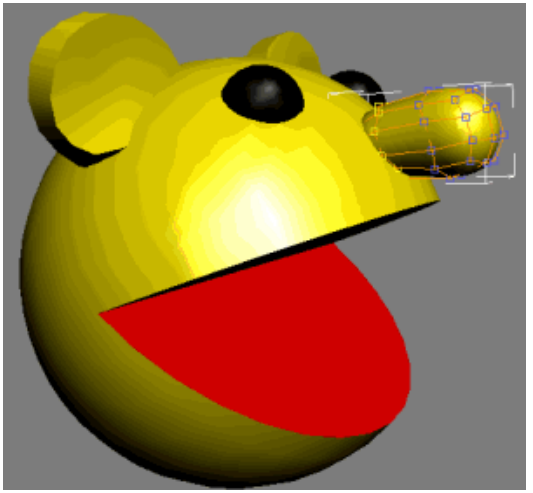
Reset Default Values—Restores the default values for parameters in the modifier.

Constraints rollout



These options let you create different types of deformable constraint for the body. For more information, see *Constraining Deformable Bodies (page 2–795)*.

FFD Soft Bodies



reactor provides two different types of soft body: mesh-based soft bodies, where the underlying mesh is directly deformed, and Freeform Deformation (FFD) bodies. With an FFD body, reactor encases the original shape in a simpler FFD lattice. Then the simulation uses the lattice, rather than the object itself, as the soft body's shape. As the lattice deforms, it updates the original shape so that it, too, appears to deform.

FFD bodies are less accurately modeled than mesh-based soft bodies, but also less computationally expensive to simulate. Because of this, it's generally advisable to use an FFD-based soft body when deforming a complex mesh. As a general rule of thumb, if a mesh has more than 200 triangles, it's probably a good idea to use an FFD-based soft body.

Note: When you view an FFD soft body in the Preview Window, you will see the FFD rather than the underlying shape. This should still give you some idea as to how the body will behave in your final animation.

Procedures

To create an FFD soft body:

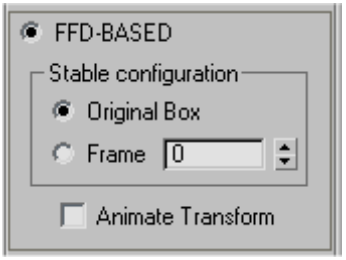
1. Create the mesh to use as the basic, non-deformed state of your soft body.
2. Apply an FFD modifier: the *FFD 2x2x2, 3x3x3, 4x4x4 (page 1–683)* or *FFD(box) (page 1–685)* modifier.

   Different FFD dimensions produce different behavior. When using FFD(box), set the dimensions to a number that leaves the FFD vertices evenly distributed over the lattice. You can also use the Conform To Shape command to make the FFD lattice better represent the underlying geometry.
3. Apply the *reactor Soft Body modifier (page 2–784)* on top of the FFD modifier.
4. in the Properties rollout of the Soft Body modifier, choose FFD-Based.

   You can now add the body to a *soft body collection (page 2–788)* and simulate it as an FFD soft body.

Interface

FFD Options



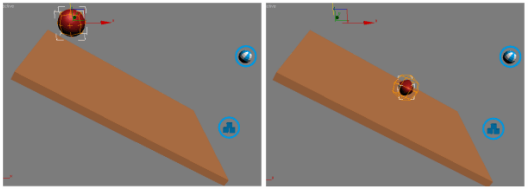
FFD-based—Specifies that you want to simulate this body as an FFD soft body.

Stable Configuration—Here you can specify which configuration (position of vertices) reactor considers to be the stable (non-deformed) configuration. The FFD soft body will tend to maintain that configuration.

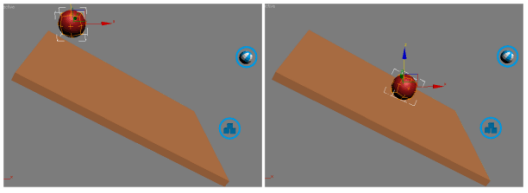
- **Original Box—**The stable configuration for an FFD soft body is the original, non-deformed FFD box; any animation or modification of the lattice is ignored. This is the default option.
- **Frame—**Lets you specify the state of the FFD state at a particular keyframe as the stable configuration for the object.

Animate Transform—When on, reactor animates both the FFD lattice *and* the transformation of the object. In some situations, the deformation calculated by the FFD modifier is more consistent if the object transformation follows the FFD lattice.

Important: Use this option only when the FFD encloses the entire mesh of the object, as the transformation applies to the whole mesh; that is, the whole object will be transformed.



Not animating the object transform can cause the FFD modifier to apply strange deformations.



Animating the object transform (turning on Animate Transform) fixes the problem.

Soft Body Collection 

Create panel > Helpers > reactor > SBCollection

Menu bar > reactor > Create Object > Soft Body Collection

reactor toolbar > create Soft Body Collection button

The Soft Body Collection is a reactor helper object that acts as a container for soft bodies. Once you have created a Soft Body Collection, you can add any soft bodies in the scene to the collection. For more information about soft bodies, see *Soft Body Modifier (page 2–784)*.

When you run the simulation, the soft body collections in the scene are examined, and, provided the collections are not disabled, reactor adds the soft bodies they contain to the simulation.

Procedures

To create and use a Soft Body Collection:

1. Create some soft bodies by applying the reactor SoftBody modifier to standard objects.
2. Choose any of the above commands, and then click in any viewport to add the Soft Body Collection.

   Note: The icon's position has no effect on the collection's behavior.

   The collection icon is added to the scene. You can add soft bodies to the collection in two ways: by picking or by using a selection list.
3. Add objects to the collection using either (or both) of these methods:
   - On the Properties rollout, click the Pick button, and then in the viewport click an object to add.
   - On the Properties rollout, click the Add button, and then use the Select Soft Bodies dialog to specify one or more objects to add.

To create the collection and add objects in a single step:

- See the *Scripts (page 2–817)* topic.

Interface

Soft Body Collection Properties rollout



Highlight—Click this button to cause the objects in the Soft Bodies list to momentarily display as if selected.

Soft Bodies—Lists the names of the objects currently in the Soft Body Collection.

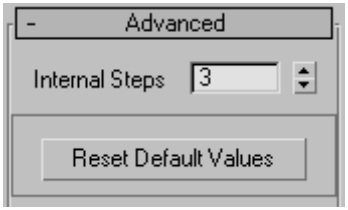
Pick—Adds an object to the Soft Body Collection. Click this button, and then in the viewport move the cursor over the object to add to the collection. If it is a soft body, the cursor will change from an arrow to a cross and you can select the object to add it to the collection.

Add—Adds one or more objects from the scene to the collection. Click the button to open the Select Soft Bodies dialog. Highlight one or more items in the list, and then click the Select button to add the objects to the collection.

Delete—Removes objects from the collection. In the Soft Bodies list, highlight the bodies to remove from the collection and click this button.

Disabled—When on, the collection and in turn the bodies it contains are not added to the simulation.

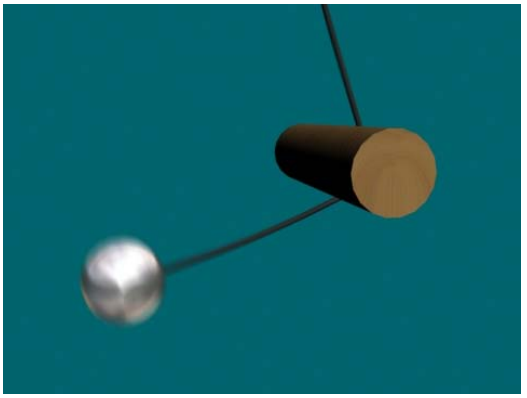
Advanced rollout



Internal Steps—Specifies how many substeps per keyframe are used to simulate the collection. Deformable bodies often require a higher level of simulation accuracy than rigid bodies, so you might need to tweak this value to get realistic results.

Reset Default Values—Resets Internal Steps to its default value.

Rope



The reactor Rope is a one-dimensional *deformable body (page 2–777)*. You can use it to simulate rope, string, hair, etc. Only *shapes (page 1–262)* can be simulated as ropes.

This section contains two topics:

- *Rope Modifier (page 2–789)*
- *Rope Collection (page 2–792)*

Rope Modifier

Modify panel > Modifier List > reactor Rope
Menu bar > reactor > Apply Modifier > Rope Modifier
reactor toolbar > Apply Rope Modifier button

You can create a reactor Rope using any spline object in 3ds Max. The Rope modifier turns the object into a deforming, one-dimensional chain of vertices. You can use rope objects to simulate ropes, as well as hair, chains, fringing, and other rope-like objects.

A rope must be added to a *rope collection (page 2–792)* in order to be simulated.

Procedures

To create a rope:

1. Create the spline-based shape to use to create the rope. As with all *deformable bodies (page 2–777)*, the underlying topology of this object will influence the rope's behavior. If the spline contains only two vertices, then the rope simulation geometry will have only a single section and will behave in a particularly un-rope-like fashion as a result. More vertices will allow the rope to bend easily and behave more realistically.
2. With the object selected, choose one of the commands listed above.

   The Rope modifier appears in the object's modifier stack.

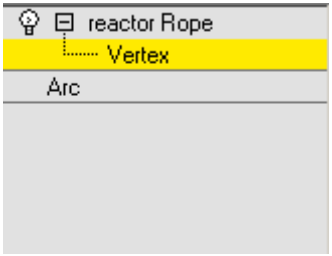
To edit a Rope's physical properties:

1. Select the Rope object.
2. On the Modify panel, ensure that the reactor Rope modifier is selected in the modifier stack.

3. Use the Properties rollout to edit the properties, which are described in the Interface section, following.

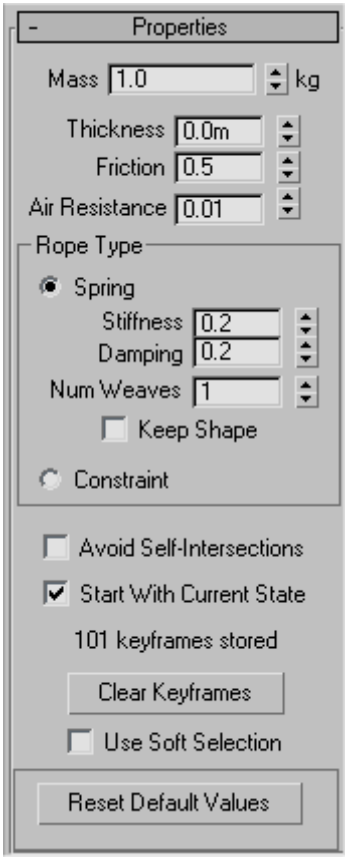
Interface

Modifier Stack



Vertex—Allows you to select individual vertices of the rope to apply and modify *deformable constraints (page 2–795)*.

Rope Properties rollout



Mass—The rope's mass in kilograms. This affects the rope behavior on collisions against other objects, its buoyancy when interacting with water, and the stretching caused by *attached rigid bodies (page 2–798)*: the higher the mass of the rope, the less stretching induced by the rigid body.

Thickness—Lets you specify a thickness for the simulated rope, as the shape used to create the rope has no inherent thickness. A rope with a thickness of 0.0 will not be visible in the simulation.

Friction—The coefficient of friction for the rope's surface. As with rigid bodies, this affects how smoothly the rope will move relative to surfaces it's in contact with. The Friction values for both

objects are combined to produce a coefficient for the interaction.

Air Resistance—The extent to which the rope loses energy as it moves.

Rope Type—

- **Spring—**The default rope type.
  - **Stiffness—**How much the rope can stretch.
  - **Damping—**How quickly oscillations settle down when the rope is compressed or expands.
  - **Num Weaves—**The extent to which the rope's inflexibility extends across vertices. For example, a value of 3 ensures that all vertices separated by three or fewer sections are subject to extra forces to stop bending. Large Num Weaves values produce an inflexible rope, whereas small values mean that only sections very close to each other will have restricted relative movement.
  - **Keep Shape—**When on, the rope tries to keep its original shape, such as a spiral, rather than returning to a straight line. The effectiveness of this depends on the number of weaves (the higher of number of weaves, the stronger the shape-keeping forces).
- **Constraint—**Uses a simpler, and hence less computationally expensive, model to simulate the rope. Ropes of this type are quite stiff, and cannot be configured in the same way as ropes with a Spring type.

Avoid Self-Intersections—When on, the rope will not intersect with itself during the simulation. This results in a more realistic simulation, but can increase simulation time.

Start With Current State—The rope starts the simulation using the current state stored in the modifier. This can be useful if you have, for example, draped the rope around something in the Preview Window and then updated the viewport using Update MAX. When off, the rope starts with the state it had originally, below the modifier in the stack.

 # Stored Keyframes—The number of keyframes, if any, stored for the rope. reactor stores keyframes for the rope if you create an animation, or if you use Update MAX in the Preview Window.

Clear Keyframes—Clears any stored keyframes for this rope.

Use Soft Selection—Allows you to use Soft Selection to smooth the transition between keyframed and simulated vertices for this deformable object. For more information, see *Soft Selection (page 2–800)*.

Reset Default Values—Resets the values for this rope to their defaults.

Constraints rollout



These options let you create a different types of deformable constraint for the body. For more

information, see *Constraining Deformable Bodies (page 2–795)*.

Rope Collection 

Create panel > Helpers > reactor > RPCollection

Menu bar > reactor > Create Object > Rope Collection

reactor toolbar > create Rope Collection button

The Rope Collection is a reactor helper object that serves as a container for ropes. Once you have placed a Rope Collection in your scene, you can add any ropes in the scene to the collection. You can find out more about ropes in the *Rope Modifier (page 2–789)* section.

When you run the simulation, the rope collections in the scene are examined, and, provided that the collections are not disabled, the ropes they contain are added to the simulation.

Procedures

To create and use a Rope Collection:

1. Choose any of the above commands, and then click in any viewport to add the Rigid Body Collection.

   Note: The icon's position has no effect on the collection's behavior.

   The collection icon is added to the scene. You can add rigid bodies to the collection in two ways: by picking or by using a selection list.

2. Add objects to the collection using either (or both) of these methods:

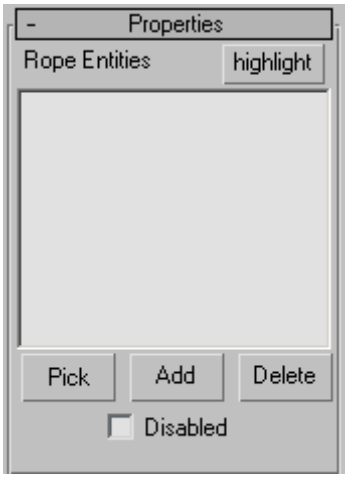
   - On the Properties rollout, click the Pick button, and then in the viewport click an object to add.
   - On the Properties rollout, click the Add button, and then use the Select Rigid Bodies dialog to specify one or more objects to add.

To create the collection and add objects in a single step:

- See the *Scripts (page 2–817)* topic.

Interface

Rope Collection Properties rollout



Highlight—Causes the objects in the collection to flash momentarily in the viewports.

Rope Entities—Lists the names of the objects currently in the Rope Collection.

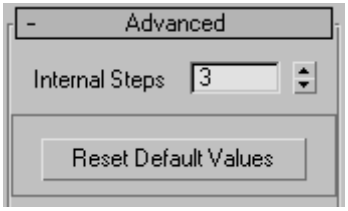
Pick—Adds an object to the Rope Collection. Click this button, and then in the viewport move the cursor over an object. If it is a Rope object, the cursor changes from an arrow to a cross and you can select the object to add it to the collection.

Add—Adds one or more objects from the scene to the collection. Click the button to open the Select Ropes dialog. Highlight one or more objects in the list, then click the Select button to add the objects to the collection.

Delete—Removes objects from the collection. In the Rope Entities list, highlight the bodies to remove from the collection and click this button.

Disabled—When on, the collection and the bodies it contains are not added to the simulation.

Advanced rollout



Internal Steps—Because deformable objects are more complex to simulate, it is usually necessary to perform more *simulation steps (page 2–709)* to increase stability. This parameter specifies how many steps will be taken internally to simulate objects in this collection for each simulation substep taken globally during the simulation. For example, at 60 fps, simulating at one keyframe (step) per frame, 10 substeps, at three internal substeps for this collection, the objects in this collection will be simulated using 60x10x3=1,800 steps every second , while the simulation in general uses 60x10=600 steps per second. This allows the objects in this collection to be simulated more accurately without slowing down the rest of the simulation. In general terms, you can think of it this way:

- Precision (stability) of overall simulation=Number of Substeps (as specified in the *Preview and Animation rollout (page 2–806)*)
- Precision (stability) of objects in this collection=Precision of overall simulation x (multiplied by) Internal Substeps

Reset Default Values—Resets Internal Steps to its default value.

Deforming Meshes



A deforming mesh is a mesh whose vertices' behavior has been keyframed. The skin of a skinned character, where any deformation comes from the underlying animated character rig, could be used as a deforming mesh in reactor.

Why use deforming meshes? Firstly, *rigid bodies (page 2–716)* and *deformable bodies (page 2–777)* can collide with a deforming mesh as if it was another reactor body. The deforming mesh won't be affected by the collision, but the rigid body/deformable body will. So, for instance, you can create deformable clothes for a deforming mesh that you can drape realistically around his body, or you can cause environmental objects to react to the deforming mesh.

You can also attach deformable bodies to deforming meshes. This allows you, for example, to attach physically-simulated hair, garments, or tentacles to a keyframed character. When the character moves, its simulated elements swirl, swing, or slither (as appropriate) along with it. This is known as secondary motion, and is an easy way to add realism to your characters without having to animate the behavior by hand.

Unlike the other reactor deformable body types, you do not need to apply a special modifier to a

deforming mesh or set physical properties for it. You just have to add the shape to a *Deforming Mesh Collection (page 2–794)*. This is because reactor itself does not deform the mesh's vertices during the simulation: The deforming mesh is simply controlled by its current animation.

For information about how to attach objects to deforming meshes , see *Constraining Deformable Bodies (page 2–795)*.

Deforming Mesh Collection

Create panel > Helpers > reactor > DMCollection
Menu bar > reactor > Create Object > Deforming Mesh Collection
reactor toolbar > create Deforming Mesh Collection button

The Deforming Mesh Collection is a reactor helper object that acts as a container for *deforming meshes (page 2–793)*. Once you have placed a Deforming Mesh Collection in your scene, you can add any deforming meshes in the scene to the collection. For more information about deforming meshes, see *Deforming Meshes (Skin) (page 2–793)*.

When you run the simulation, reactor examines the deforming mesh collections in the scene, and, provided the collections are not disabled, adds the deforming meshes they contain to the simulation.

Procedures

To create and use a Deforming Mesh Collection:

1. Create some objects to use as deforming meshes.
2. Choose any of the above commands, and then click in any viewport to add the Deforming Mesh Collection.

   Note: The icon's position has no effect on the collection's behavior.

   The collection icon is added to the scene. You can add deforming meshes to the collection in two ways: by picking or by using a selection list.
3. Add objects to the collection using either (or both) of these methods:
   - On the Properties rollout, click the Pick button, and then in the viewport click an object to add.
   - On the Properties rollout, click the Add button, and then use the dialog to specify one or more objects to add.

To create the collection and add objects in a single step:

- See the *Scripts (page 2–817)* topic.

Interface

Deforming Mesh Collection Properties rollout



Highlight—Click this button to cause the objects in the collection to flash momentarily in the viewports.

Deforming Meshes—Lists the names of the objects in the deforming mesh collection.

Pick—Adds an object to the deforming mesh collection. Click this button, and then in the viewport move the cursor over an object to add

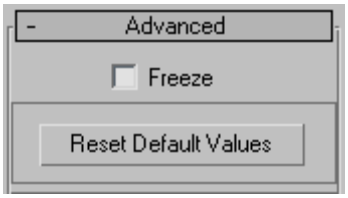
to the collection. If it can be used as a deforming mesh, the cursor will change from an arrow to a cross and you can select the object to add it to the collection.

Add—Adds one or more objects from the scene to the collection. Click the button to open the Select Deforming Meshes dialog. Highlight one or more objects in the list, and then click the Select button to add the objects to the collection.

Delete—Removes objects from the collection. In the Deforming Meshes list, select the bodies to remove from the collection and then click this button.

Disabled—When on, the collection and the bodies it contains are not added to the simulation.

Advanced rollout



Freeze—When on, the deforming meshes in this collection do not follow their animation during the simulation but instead remain static. This is useful, for instance, if you want to drape clothes around a deforming-mesh-skinned character in the *Preview window (page 2–815)* before using Update MAX - it's much easier to dress a character that isn't still walking around!

Reset Default Values—Resets the Freeze value to its default.

Deformable-Body Constraints

reactor modifier (Cloth/Soft Body/Rope) > Constraints rollout

As with rigid bodies, you might want to constrain the possible movement of your deformable bodies. For instance, you might want to fix cloth draperies to a position in world space, or attach deformable clothing and hair to a moveable character. However, because deformable bodies can change their shape over the course of the simulation, you can't just specify attachment points in the object's local space as you do with rigid bodies.

Instead, when you work with deformable bodies, you need to specify the vertices to constrain in the deformable mesh. You can then constrain the vertices using one of four deformable constraint types, which you can choose from the deformable body's deformable Constraints rollout.

This section has separate topics for each of the following deformable constraint types:

- *Fixing Vertices in World Space (page 2–796)*
- *Keyframing Vertices (page 2–797)*
- *Attaching Vertices to a Rigid Body (page 2–798)*
- *Attaching Vertices to Deforming Meshes (Skin) (page 2–799)*

For information about using the Soft Selection option to smooth the transition between simulated and keyframed vertices for deformable bodies, see *Soft Selection (page 2–800)*.

Note: You can constrain deformable bodies only to points in world space, rigid bodies, or deforming meshes. You can't attach deformable bodies to "live" deformable bodies. You can, however, attach deformable bodies to other deformable bodies once they have been simulated. For example, after simulating a piece of cloth, you can remove it from

the Cloth Collection and add it to a *Deforming Mesh Collection (page 2–794)*. You can then attach any deformable body (cloth, soft or rope) to it using the *Attach to Deforming Mesh constraint (page 2–799)*.

Procedures

The procedure for each deformable constraint type is described in its own topic.

Interface

Constraints rollout



Fix Vertices—Creates a *Fix Vertices to World constraint (page 2–796)* that fixes the chosen vertices to their current position in world space.

Keyframe Vertices—Creates a *Keyframe constraint (page 2–797)* that makes the chosen vertices follow their current animation in 3ds Max.

Attach to Rigid Body—Creates an *Attach to Rigid Body constraint (page 2–798)* between the chosen vertices and a rigid body. The vertices follow the animation (changes in position and rotation) of the *rigid body (page 2–716)*.

Attach to DefMesh—Creates an *Attach to Deforming Mesh constraint (page 2–799)* between the chosen vertices and a deforming mesh. The vertices will follow the animation (deformations) of the *deforming mesh (page 2–793)*.

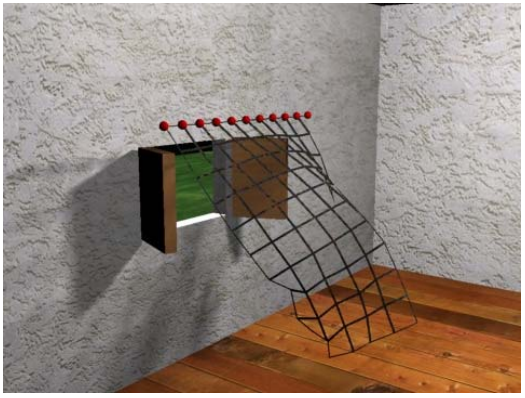
[constraints list]—Displays a list of the deformable constraints for the deformable body. To highlight a constraint for changing its parameters or deleting it, click its name in the list.

Delete Constraint—Deletes the highlighted constraint in the list.

Fixing Vertices in World Space

reactor modifier (Cloth/Soft Body/Rope) > Constraints rollout

The Fix Vertices constraint lets you fix vertices in a deformable body to their current positions in world space.



Procedures

To fix points in a deformable body in world space:

1. On the Modify panel, click Vertex in the deformable body modifier's sub-object list.

This lets you select individual vertices in the deformable body.

2. On the Constraints rollout, click Fix Vertices.

   A new Constrain To World constraint appears in the constraints list.

3. Ensure the constraint is highlighted in the constraints list.

4. Select the vertices to constrain.

   These vertices remain fixed in place during the simulation, while the rest of the mesh is subject to physical forces and deformation as usual.

5. When you are finished choosing vertices, it's advisable to exit the Vertex sub-object level. This prevents you accidentally deselecting some of the constrained vertices before simulating.

6. To change the set of fixed vertices, select the object, access the modifier's Vertex sub-object level, highlight its Constrain To World constraint in the list, and then select a different set of vertices.

Keyframing Vertices

reactor modifier (Cloth/Soft Body/Rope) > Constraints rollout

This constraint allows you to make the deformable body's chosen vertices follow their current animation.



Procedures

To keyframe points for a deformable body:

1. On the Modify panel, click Vertex in the deformable body modifier's sub-object list.

   This lets you select individual vertices in the deformable body.

2. On the Constraints rollout, click Keyframe Vertices.

   A new Keyframe constraint appears in the constraints list.

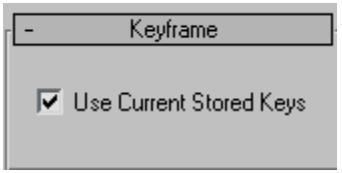
3. Ensure the constraint is highlighted in the constraints list; this makes the Keyframe rollout available.

4. Select the vertices to constrain in the deformable body. During the simulation, these vertices follow the deformable body's current animation, while the rest of the body is affected by the physical simulation as usual.

5. When finished, it's advisable to exit the Vertex sub-object level. This prevents you accidentally deselecting some of the constrained vertices before simulating.

Interface

Keyframe rollout

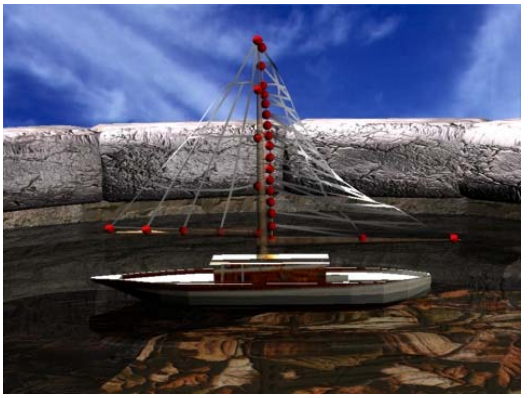


Use Current Stored Keys—When on, the animation used for these vertices is that *after* the reactor modifier is applied. In other words, if the modifier contains any keyframes, reactor uses those for the animation. When off, reactor evaluates the animation *before* the reactor modifier is applied, thus ignoring any stored keyframes.

Attaching Vertices to a Rigid Body

reactor modifier (Cloth/Soft Body/Rope) > Constraints rollout

The Attach To Rigid Body constraint allows you to fix vertices in a deformable body to a particular rigid body.



Procedures

To fix vertices in a deformable body to a rigid body:

1. On the Modify panel, click Vertex in the deformable body modifier's sub-object list.

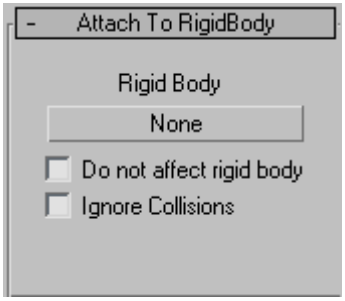
   This lets you select individual vertices in the deformable body.
2. On the Constraints rollout, click Attach To Rigid Body.

   The new constraint appears in the Constraints list.
3. Ensure the constraint is selected in the Constraints list; this makes the Attach To RigidBody rollout available.
4. On the Attach To RigidBody rollout, click the Rigid Body pick button and then select the body to which to constrain the vertices.
5. Select the vertices to constrain in the deformable body.

   During the simulation, these vertices maintain their positions relative to the rigid body.
6. When finished, it's advisable to exit the Vertex sub-object level, to keep from accidentally deselecting some of the constrained vertices before simulating.

Interface

Attach To Rigid Body rollout



Rigid Body—Displays the name of the rigid body to which the vertices are constrained. Choose a rigid body by clicking this button and then selecting a rigid body in one of the viewports.

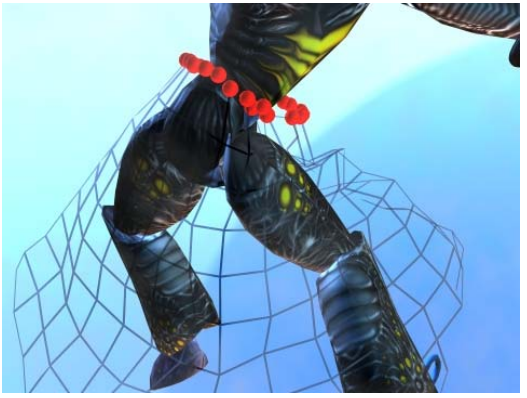
Do not affect rigid body—When on, the attached deformable body does not affect the behavior of the rigid body in the simulation. In other words,

the points in the deformable body will follow the rigid body, but no force is applied to the rigid body from the deformable.

Ignore Collisions—When on, disables collision detection between the rigid body and the deformable body.

Attaching Vertices to Deforming Meshes (Skin)

reactor modifier (Cloth/Soft Body/Rope) > Constraints rollout



The Attach To Deforming Mesh constraint allows you to fix vertices in a deformable body to a particular *deforming mesh (page 2–793)*. This option is useful for adding physically simulated elements such as hair and clothing to an skinned character.

Procedures

To fix vertices in a deformable body to a deforming mesh:

1. On the Modify panel, access the Vertex sub-object level of the deformable body's modifier.

   This lets you select individual vertices in the deformable body.

2. On the Constraints rollout, click Attach To DefMesh .

   The new constraint appears in the Constraints list.

3. Ensure the constraint is selected in the Constraints list; this makes the Attach To DefMesh rollout available.

4. On the Attach To DefMesh rollout, click the Deformable Mesh pick button and then select the body to which to constrain the vertices.

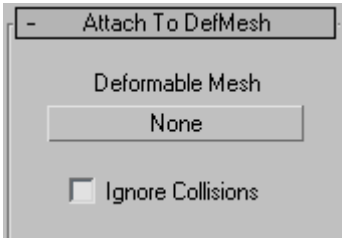
5. Select the vertices to constrain in the deformable body.

   During the simulation, these vertices will maintain their position relative to the deforming mesh.

6. When finished, it's advisable to exit the Vertex sub-object level, to keep from accidentally deselecting some of the constrained vertices before simulating.

Interface

Attach To DefMesh rollout



Deformable Mesh—Displays the name of the deforming mesh to which the vertices are constrained. Choose a deforming mesh by clicking this button and then selecting the object in one of the viewports.

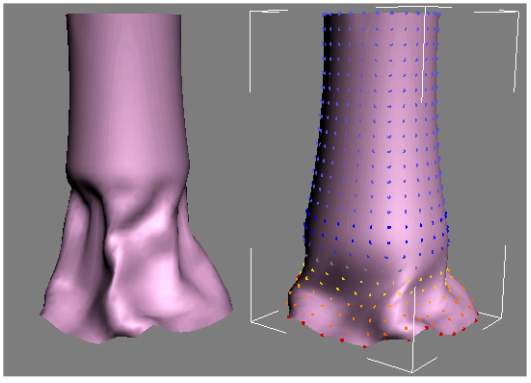
Ignore Collisions—When on, disables collision detection between the deforming mesh and the deformable body.

Soft Selection

You can use deformable constraints to specify that you want certain vertices of a deformable object to follow a user-specified animation, such as skinning, while the rest of the object is fully physically simulated. However, sometimes this can result in a visible "join" between the animated and simulated parts of a piece of cloth, soft body or rope.

To deal with this, the *Cloth (page 2–778)*, *Soft Body (page 2–784)*, and *Rope (page 2–789)* modifiers have a Use Soft Selection option. The option can help smooth transitions between the keyframed and simulated vertices.

The following illustration shows the effect of using this option. The cloth tube on the left has a keyframed top and a simulated bottom, with a visible join. The cloth on the right uses soft selection to blur the line between the two.

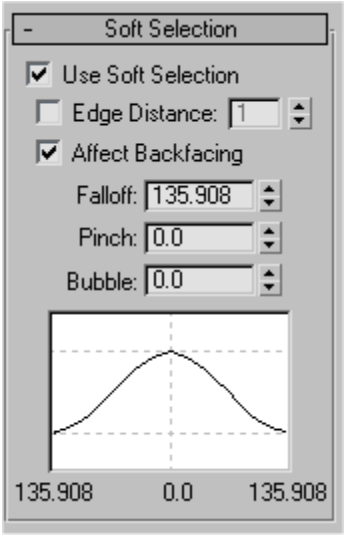


Procedures

To use soft selection when animating a deformable body:

1. Before you apply the reactor modifier to a deformable body, apply a modifier that has a Soft Selection option to your object, such as *Edit Mesh (page 1–634)* or *Mesh Select (page 1–719)*.

   You can also apply this modifier to the body later, but it's important that it be *below* the reactor modifier in the object's modifier stack.
2. Follow the instructions in *Keyframing Vertices (page 2–797)* to keyframe selected points for the deformable body, then create your reactor animation.
3. With the body selected, open the selection modifier in the modifier stack, and access the Vertex sub-object level.
4. Select the vertices that should appear fully physically simulated. For instance, for a piece of cloth where the top section is keyframed, you might want only the bottom vertices to appear to be fully affected by the physical simulation.
5. On the Soft Selection rollout, turn on Use Soft Selection.



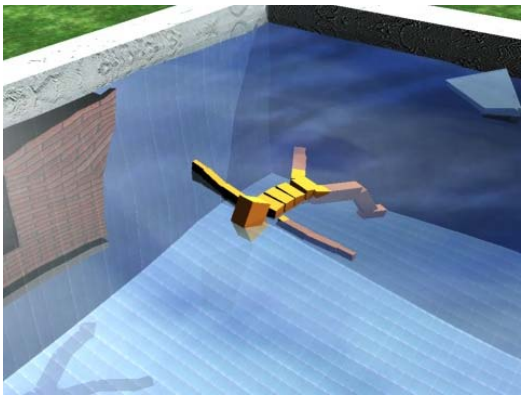
6. Set the Falloff value to specify a smooth falloff from the selected vertices to the unselected ones.

   This is represented visually in the viewport.
7. Now open the reactor modifier and turn on Use Soft Selection.

When you run the animation, your specified falloff will be used to blend smoothly between the keyframed and simulated vertices.



Water



The reactor Water object allows you to simulate the behavior of a water surface. Objects can interact with the water in physically realistic ways, creating waves and ripples. reactor calculates a buoyancy value for any objects that fall into the water using their mass and size, so that some objects sink and others float. You can even change the density of a water object, which affects how objects float in it.

Water in reactor is a *space warp (page 3–1014)* that is simulated as water. Space warps don't render, so you need to provide reactor with a geometrical representation of the water if you want it to appear in your final animation. You do this by binding a plane or other geometry to the space warp. Note that a representation of the water still appears in the Preview Window even if you haven't done this.

This section covers the following topics:

- *The Water Space Warp (page 2–801)*
- *Rendering Water (page 2–803)*

Water Space Warp

Create panel > Space Warps > reactor > Water
menu bar > reactor > Create Object > Water
reactor toolbar > Create Water button

You can use the Water space warp to simulate the behavior of a liquid surface in your reactor scene. You can specify a size for the water, and physical properties such as density, wave speed, and viscosity.

You don't need to add the water to a collection for it to take part in the simulation. However, while it will appear in the Preview Window, it will not appear in a rendered animation unless you bind the space warp to a plane or other geometry. You can find out how to do this in *Rendering Water (page 2–803)*.

Procedures

To create a Water space warp:

1. Choose one of the above commands, and then in any viewport drag out a Water object.

   Tip: For horizontal water, create the space warp in a Perspective or Top viewport.

2. Move and rotate the object as necessary using standard commands. For changing other properties, see the next procedure.

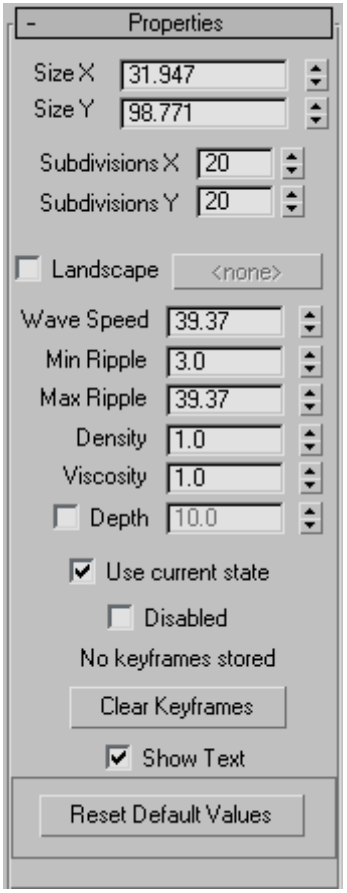
To edit the properties of a Water object:

1. Select the object in the scene.
2. On the Modify panel, open the Properties rollout.

3. Use the rollout to specify properties, as described in the Interface section.

Interface

Water Properties rollout



Size X/Y—The dimensions for the water object.

Subdivisions X/Y—The tessellation for the water's mesh.

Landscape button—Although water is defined using a rectangular area, you can simulate non-rectangular surfaces and obstacles inside the water by defining a *landscape geometry*. Once you designate landscape geometry, any vertex in the Water space warp that are contained in the landscape are fixed during the simulation and waves and ripples will be reflected at those points.

Wave Speed—The speed at which wave crests propagate across the surface of the water.

Min/Max Ripple—Limits on the size of the waves generated in the water.

Density—The relative density of the liquid. This determines which objects will sink into the water, and at what height objects of a lesser density will float. The density is specified relative to the density of water (1.0=1,000 kg/m^3).

Viscosity—Resistance to flow: how difficult it is for objects to move through the liquid. A high value means that the motion of objects through the water is highly damped.

Depth—The depth of the water. Buoyancy is applied only to objects inside the water.

Use Current State—When on, the simulation uses the current animation to calculate the starting state. So if you have, for example, updated the water during the *preview (page 2–815)* using the Update MAX command, the stored state is used as the initial state.

Disabled—Removes the water from the simulation.

 # Stored Keyframes—This indicates if there are any keyframes stored for the water. Keyframes are stored for the water if you create a reactor animation in 3ds Max, or if you use Update MAX in the Preview Window.

Clear Keyframes—Clears any stored keyframes for this water.

Show Text—When off, the label "Water" doesn't appear next to the space warp in the viewports.

Reset Default Values—Restores the default values for the object.

Rendering Water

Water in reactor is defined using a space warp. Space warps don't render, so if you want the water to appear in your final animation, you need to provide reactor with a geometrical representation of the water. To do this, you need to bind a plane or any other planar geometry to the water. The plane will deform according to the *Water space warp (page 2–801)*, and you can then render it as part of your scene.

Binding an object to the water applies a modifier to that object, which lets you set a scale strength for the deformation. For example, if you set this scale value to 2.0, any water deformation is doubled for the vertices in the object.

Procedures

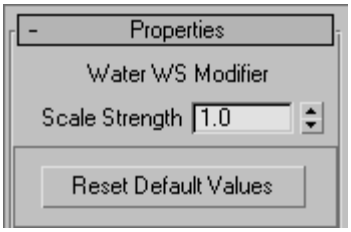
To bind a plane to reactor Water:

1. Draw a plane the same size as your water in the viewport. Put it in a different place so you can easily link the two.
2.  Click the Bind To Space Warp button on the main toolbar.
3. In the viewport, drag from the plane to the water space warp, or vice-versa.

   The reactor Water (WSM) modifier appears in the plane's modifier stack.
4. Move the plane to the same position and orientation as the space warp (or vice-versa).

Interface

Water WS Modifier Properties rollout



Scale Strength—Acts as a multiplier for the geometry's deformation by the water. Increases or decreases the deformation applied by the *Water space warp (page 2–801)*.

Reset Default Values—Returns Scale Strength to its default value: 1.0.

Wind 

Create panel > Helpers > reactor > Wind

menu bar > reactor > Create Object > Wind

reactor toolbar > Create Wind button



The Wind helper object lets you add wind effects to reactor scenes, allowing you, for instance, to make curtains flap in the breeze. After adding the helper to your scene, you can configure various properties for the effect such as its speed, gusts,

and whether objects in your scene can be sheltered from the wind. You can animate most of these parameters. The orientation of the helper icon indicates the direction of the wind: It blows in the direction of the weathervane arrow. You can also animate this direction by animating the icon's orientation.

Procedures

To add wind to a scene:

1. Choose one of the above commands, and then click in any viewport to add the Wind helper.

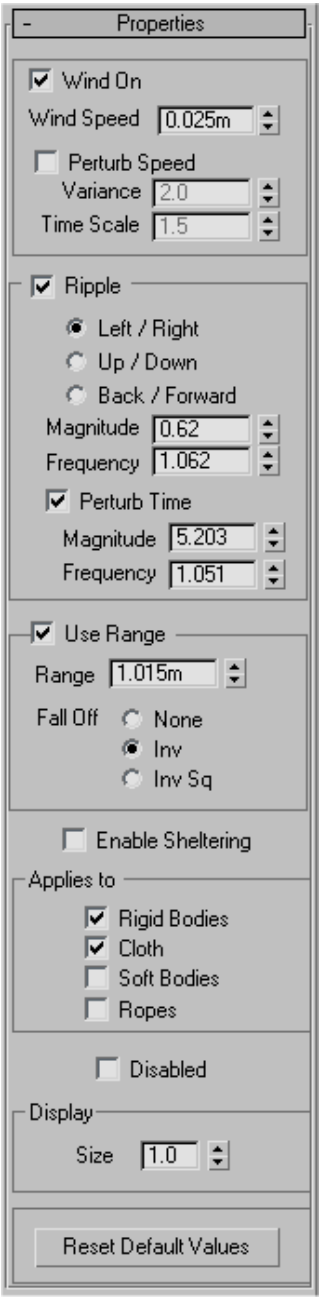
   Note: The icon's position has no effect on the wind's behavior unless you activate the Use Range option.

2. Use the Rotate tool to rotate the icon and set the wind direction.

   The wind blows in the direction indicated by the weathervane arrow.

Interface

Wind Properties rollout



Wind On—[Animatable] Defines whether the wind forces are applied or not.

Wind Speed—[Animatable] The strength of the wind force. The direction is specified by the icon's orientation, which is also animatable.

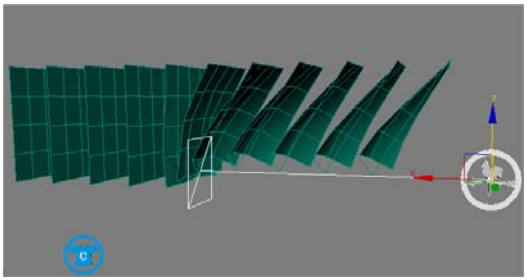
Perturb Speed—When on, the strength of the wind varies over time, using the following parameters.

- **Variance—**The maximum amount of change in the speed.
- **Time Scale—**How quickly speed change occurs, with small values producing slow changes, and large values producing rapid changes.

Ripple—[Animatable] When on, the wind direction becomes a function of space and time, allowing you to add rippling effects to cloth objects affected by the wind. This effect can occur in the Left/Right, Up/Down, or Back/Forward directions, where Forward is the wind direction and Up is the up-axis of the wind icon.

- **Magnitude—**The directional variance for the ripples.
- **Frequency—**The regularity of ripple formation.
- **Perturb Time—**When in, the spatial perturbation is itself perturbed over time, meaning that the ripples move back and forth. The Perturb Time feature has its own Magnitude and Frequency parameters. When off, the rippling cloth should eventually come to rest, unless Perturb Speed is on.

Use Range—[Animatable] When on, the wind effect has the specified range of action, starting from the icon. The range limit is displayed in the viewport, as depicted in the following illustration.



- **Fall Off—**The extent to which the wind effect falls off towards its range limit. You can choose from None, Inv (strength decreases proportionally to distance), and Inv Sq (strength decreases proportionally to the square of the distance).

Enable Sheltering—[Animatable] When on, objects can be sheltered from the wind by other objects.

Applies To—These check boxes allow you to specify the types of objects that are affected by the wind. You can choose from Rigid Bodies, Cloth, Soft Bodies, and Ropes.

Disabled—Removes the wind from the simulation.

Display—The size of the wind icon in the viewports.

Reset Default Values—Resets the default property values for this Wind helper.

The reactor Utility

Utility panel > reactor

You can access much of reactor's functionality through the reactor utility. It lets you preview the simulation, change world and display parameters, and analyze the convexity of objects. It also lets you see and edit the rigid body properties associated with objects in the scene. Perhaps most

important, it lets you switch between the Havok 1 and Havok 3 engines.

The utility is divided into a number of rollouts. For more information, follow the links. More information about the two different Havok versions is available later in this topic, in *Differences between Havok 1 and Havok 3 (page 2–806)*.

- About: Choose Havok 1 or Havok 3 solver
- *Preview and Animation (page 2–806)*
- *Havok 1 World/Havok 3 World (page 2–808)*
- *Collisions (page 2–810)*
- *Display (page 2–812)*
- *Utils (page 2–813)*
- *Properties (page 2–815)*

Differences between Havok 1 and Havok 3

You can use either of two different dynamics solvers in reactor: Havok 1 or Havok 3. To switch between them, choose the desired solver from the drop-down list on the About rollout. The active solver is saved with the scene file.



The difference between the two, in a nutshell, is that Havok 1 provides a broader range of functionality, but Havok 3 is faster and more accurate. If your simulation requires only rigid bodies, you'll generally get better results with Havok 3.

Following is a more detailed list of differences:

Feature	Havok 1	Havok 3
Simulation	Discrete only	Discrete or *Continuous* with extra parameters
Supported bodies	Rigid/Cloth/Soft/ Rope/Deforming Mesh/Water	Rigid only
Rigid Body Properties (page 2–717)	Standard	Standard + *Shell/Penetration/ Quality*
Cooperative Constraints > *Breakable Constraints (page 2–735)*	Linear, Angular	Threshold
Cooperative Constraints > *Strength settings (page 2–735)*	Strength, Tau	Strength only
Preview Window: display	Textures, smoothing, actual geometry	Simulation geometry only
Preview Window: interaction during preview	right-button drag	hold Spacebar and move mouse
Toy Car (page 2–766)	Standard	Additional parameters

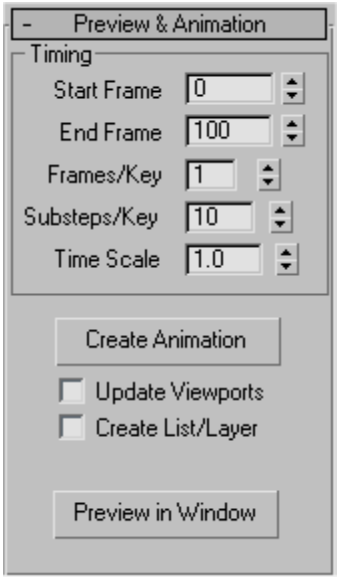
Note: If you preview a simulation or create an animation using Havok 3 and the scene contains non-supported entities such as soft bodies, reactor simply ignores these entities and excludes them from the simulation.

Preview & Animation Rollout

Utility panel > reactor > Preview & Animation rollout

This rollout lets you run and preview your reactor simulation, as well as specifying timing parameters for the simulation.

Interface



Start Frame—When creating a world to be simulated or previewed, reactor needs to access the objects in 3ds Max at a fixed point in time. This parameter defines this point in time (in 3ds Max frames). Initial shapes, positions and velocities are taken from the actual shapes, positions and velocities of the objects in the scene at this particular frame. When creating an animation, keyframes are created starting from this frame to the End Frame time.

End Frame—The last frame to simulate. When reactor creates an animation, it generates keyframes from the Start Frame time to this frame.

Frames/Key—The number of frames for every keyframe that reactor creates (time step). For example, a value of 2 will create a keyframe every other frame. Increasing this value forces reactor to take bigger time steps, potentially reducing the accuracy of the simulation; you might need to then update the Substeps/Key setting accordingly (see following).

Substeps/Key—The number of reactor simulation substeps per keyframe (that is, per time step). The higher this value, the more accurate the simulation will be, although it will also require more computation to simulate. For information about time steps and substeps, see *Time Steps (page 2–709)*.

Time Scale—This parameter maps between time in the simulation and time in 3ds Max. Changing the value lets you slow down or speed up the animation. Values less than 1.0 produce slow-motion animations, while values greater than 1.0 produce sped-up animations.

Create Animation—Runs the simulation and creates keyframes, starting at Start Frame and ending at End Frame.

Update Viewports—When on, updates the scene in the viewports as the animation is created.

Create List/Layer—When on, creates a *List controller (page 2–342)* (unless one is already present) for the position and rotation tracks of those rigid bodies using Position/Rotation/Scale transform controllers. A new subcontroller is added to the List controller to hold the keys created by reactor. This new controller is weighted to 100, while the previous controllers are weighted to 0. In the special case of rigid bodies belonging to a **character studio** *Biped (page 2–843)*, a new biped layer is created every time a reactor animation is generated.

Preview Animation—Previews the simulated scene in the *Preview Window (page 2–815)*.

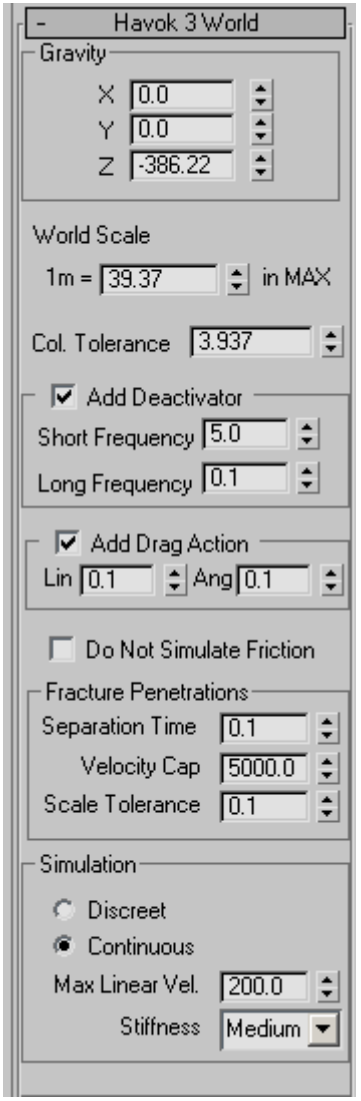
Havok 1 World/Havok 3 World Rollout

Utility panel > reactor > Havok 1 World/Havok 3 World rollout

This rollout lets you set some general parameters for your simulated world, such as the strength and direction of gravity, the scale of the world, and how easily objects can collide with each other. When Choose Solver is set to Havok 1, the rollout name is Havok 1 World, and when Choose Solver is set to Havok 3, the name is Havok 3 World.

Interface

The only difference between the two different versions of the Havok World rollout is that the Havok World 3 rollout contains an extra group of parameters at the end, named Simulation. The following illustration shows the Havok 3 version of the rollout.



Gravity—The acceleration, in world units, the objects in a scene have due to gravity. It is an important value because it affects the overall feeling of scale in a dynamics simulation.

reactor will generate a warning if you specify a non-standard Gravity value compared to World Scale (see following); that is, if you set the value much higher or lower than –9.8 m/s^2, where *m* is World Scale's representation of one meter. The

warnings are provided for guidance and can be safely ignored if you are happy with the behavior of your animation. The default Gravity value reflects "real-world" gravity: –9.8 m/s^2 (–386.22 inches/s^2) on the Z axis. For more information, see *Scale (page 2–711)*.

World Scale—The distance in 3ds Max world units, that represents one meter in the reactor world, and hence determines the size of every object in your simulation.

Note: Changing the World Scale value can drastically alter simulated objects' behavior. For more information, see *Scale (page 2–711)*.

Col. Tolerance—(Collision Tolerance) One of the tasks that reactor performs at each simulation step is detection of whether any objects in the scene are colliding, and then updating the scene accordingly. If objects are closer together than the Collision Tolerance value, reactor considers them to be colliding. A high Collision Tolerance value results in a stable simulation, but it can also cause gaps between "colliding" objects.

The default value is 1/10th of World Scale (10 cm if you are modeling using real-world sizes), and it is a good rule of thumb to always keep it above 1/40th (4 mm), due to floating point precision limits in the processor. For standard scenes and object sizes, the default parameters for world scale and tolerance should be adequate. If you are simulating very small objects where the tolerance is still to visible, try reducing the World Scale value and the Collision Tolerance and Gravity accordingly. For more information, see *Scale (page 2–711)*.

Add Deactivator—When on, reactor adds a deactivator to the simulation.

The deactivator keeps track of the objects in the simulation and deactivates (stops simulating) objects that it determines are at rest. This keeps reactor from wasting system resources simulating objects that aren't doing anything.

When Add Deactivator is on, you have access to two distance parameters for determining whether an object is moving or not. The reason that there are two properties rather than just one is for bodies that vibrate. It's possible that a body can end up vibrating after a collision or some other occurrence. A vibrating body might vibrate a few millimeters in distance, which would keep it active for a typical Short Frequency test. However, the object is not really moving anywhere, so keeping it active wastes CPU time. However, if it vibrates in place, it fails the Long Frequency test and is deactivated.

- **Short Frequency—**The minimum distance, usually in millimeters, that an object must move during each step of a simulation. If an object in a simulation does not move the specified distance in each step, reactor deactivates it.
- **Long Frequency—**Also sets a distance, usually larger than the Short Frequency value. Long Frequency checks every few steps simulation rather than at every step. Any object that does not move the required Long Frequency distance is deactivated.

Add Drag Action—When on, ensures that rigid bodies are subject to constant drag. This damps their linear and angular velocities, so they come to rest sooner. This is useful, for example, where rigid bodies are joined with *constraints (page 2–724)*, as it helps the constrained bodies to come to rest despite the forces applied by the constraints.

- **Lin—**The linear damping applied by Add Drag Action.
- **Ang—**The angular damping applied by Add Drag Action.

Do Not Simulate Friction—When on, reactor ignores all Friction values during the simulation, and objects slide easily across each other.

Fracture Penetrations—These parameters let you adjust how reactor simulates Fracture objects. For

details, see *World rollout > Fracture Penetrations group (page 2–773)*.

Simulation group

These controls are available only with the Havok 3 engine.

The first lets you choose how reactor computes the simulation:

- **Discrete**—reactor checks for collisions only at the beginning and end of each simulation step. This is faster but less accurate.
- **Continuous**—reactor checks for collisions constantly, throughout each step. This is the default choice, and results in slower but highly accurate simulations. This option significantly reduces the chances of missed collisions.

Max Linear Vel.—The assumed maximum linear velocity for all bodies.

This does not set objects' linear velocity. If, later on, a rigid body's velocity significantly exceeds this value, it might tunnel through other objects even if its Quality property is set to Critical.

The default value is 200 meters/second. This value is used during continuous simulations by the collision solver to optimize the simulation, so if it can assume a maximum linear velocity it can discard several cases and get a solution faster.

Stiffness—The hardness of the constraints in the scene.

Use this parameter to configure the dynamics solver. When constraints used in the scene are very hard or stiff, set this to Hard; the Soft or Medium setting works most other cases. Setting Stiffness to Soft can lead to unstable solutions if the constraints are too stiff.

Tip: Another way to reduce the stiffness of the simulation is by lowering the Strength value of constraints, thus relaxing the constraints.

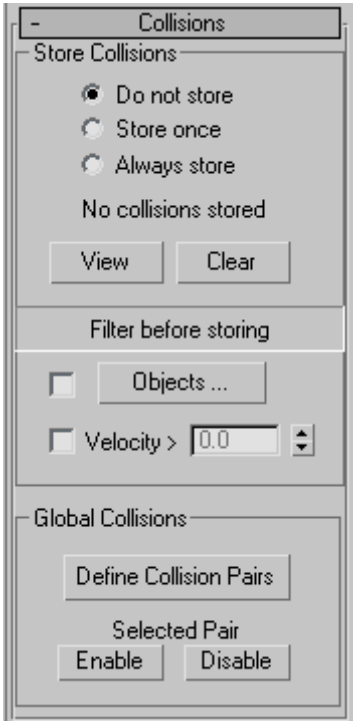
Collisions Rollout

Utility panel > reactor > Collisions rollout

These settings let you store collision details from your scene, and enable and disable collision detection for specific pairs of objects. Pairs of objects with collisions disabled pass through each other during the simulation. Disabling collisions can be useful when you have objects that are attached together, such as the wheels of a car and its chassis. This means reactor doesn't check for every small collision between the objects, which can slow down the simulation.

Interface

Collisions rollout



Store Collisions—See *Storing and Accessing Collisions (page 2–774)*.

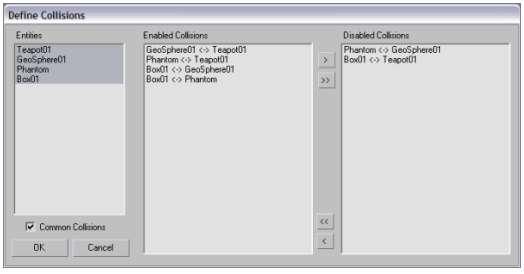
Filter Before Storing—See *Storing and Accessing Collisions (page 2–774)*.

Define Collision Pairs—Opens the *Define Collisions dialog (page 2–811)* . This allows you to toggle collision detection between specified pairs of objects. Pairs of objects with disabled collisions pass through each other during the simulation.

Selected Pair: Enable—Lets you quickly enable collisions between the selected pair of objects. This button is available only when both objects are part of a collection (not necessarily the same collection) and their collisions are currently disabled.

Selected Pair: Disable—Click this button to quickly disable collisions between the selected pair of objects. This button is available only when both objects are part of a collection (not necessarily the same collection) and their collisions are currently enabled (the default condition).

Define Collisions Dialog



Entities—Lists the names of the reactor bodies in the scene (any object added to a collection). Highlighting one object populates the Enabled Collisions and Disabled Collisions lists for that object with respect to the other bodies in the scene. If you select more than one object in this list, how the Enabled/Disabled lists are populated depends on the value of Common Collisions.

Common Collisions—When on, if more than one object is highlighted in the Entities list, the Enabled and Disabled lists are populated from the list of possible pairs that can be made using only the highlighted objects. When off, the lists contain all possible pairs from the scene that contain at least one of your selected objects.

For example, take a scene containing four bodies: Box01, Box02, Box03, and Box04. If Common Collisions is on, and you then highlight Box01 and Box02 in the Entities list, then the only possible pair that just includes your selected objects is Box01<->Box02.

If you highlight Box01, Box02, and Box03, the possible pairs using these objects are Box01<->Box02, Box02<->Box03 and Box01<->Box03. If you highlight the same objects with Common Collisions off, then the lists also include the highlighted objects paired with the remaining objects in the collection. This means that the available pairs would now include Box01<->Box04, Box02<->Box04 and Box03<->Box04.

This option is on by default.

Enabled Collisions—Lists pairs with collisions enabled.

Disabled Collisions—Lists pairs with collisions disabled.

Disable Selection (Right Arrow)—Moves pairs highlighted in the Enabled list to the Disabled list.

Disable All (Double Right Arrow)—Moves all pairs in the Enabled list to the Disabled list.

Enable Selection (Left Arrow)—Moves the pairs highlighted in the Disabled list to the Enabled list.

Enable All (Double Left Arrow)—Moves all pairs in the Disabled list to the Enabled list.

OK—Closes the dialog, saving any changes carried out while the dialog was open.

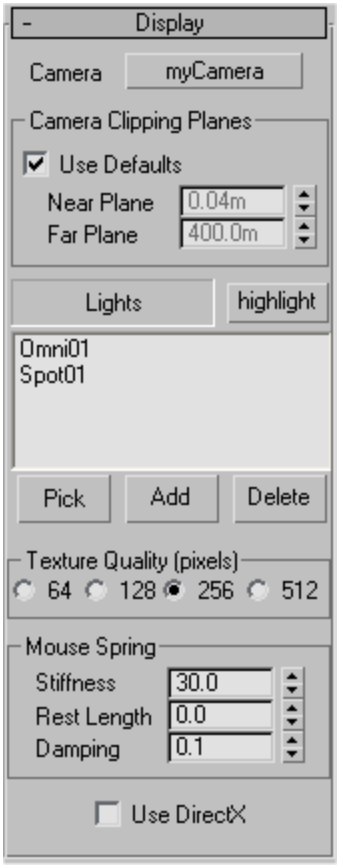
Cancel—Closes the dialog without saving changes.

Display Rollout

Utility panel > reactor > Display rollout

This rollout lets you specify display options for *previewing (page 2–815)* your simulation, including cameras and lighting. These options have no effect on the actual behavior of the final animation; they affect only the appearance of the *preview window (page 2–815)*.

Interface



Camera—Click this button and then pick a camera from the viewports to use as the initial view for the display. Your chosen camera's name appears on the button. If you don't assign a camera, the settings on the current Perspective viewport (if any) are used initially in the Preview Window.

Camera Clipping Planes—If a camera is assigned, the display tries to use the camera to generate clipping planes. If no camera is assigned, reactor uses default values for the clipping planes in the display.

Use Defaults—It is possible that the clipping planes for your specified camera are not sufficient to display everything in the scene. Rather than having to change these values in the window every time the simulation is run, you can turn off Use Defaults and then specify your own clipping planes using the Near Plane and Far Plane values.

- **Near Plane—**The near plane to be used when Use Defaults is off.
- **Far Plane—**The far plane to be used when Use Defaults is off.

Lights—Lists the scene lights reactor uses for the preview. When the list is empty, reactor creates and uses a flashlight at the camera position; you can toggle this light from the Preview Window > Display menu. You can combine up to six omni lights or spotlights to create the preview lighting. Choose lights from the scene one at a time with the Pick button, or add from a list of the available lights in the scene using the Add button. To remove a light from this list, highlight it in the list and then click Delete.

Texture Quality—The size (NxN pixels) of the textures generated for use in the display. A smaller value uses less memory.

Mouse Spring—These options allow you to configure the spring used when you select objects with the right mouse button (Havok 1) or Spacebar (Havok 3) in the *preview window (page 2–815)*.

- **Stiffness—**The stiffness of the mouse spring. If the mouse spring is too strong you might be able to pull objects through each other;

this won't cause errors, but it can create an undesirable state. Default=30.0.

- **Rest Length—**The rest length of the mouse spring. Default=0.0.
- **Damping—**The damping value for the mouse spring. Default=1.0.

Use DirectX—When on, the 3D rendering in the preview window uses the DirectX system. When off, the window uses OpenGL.

reactor Utilities

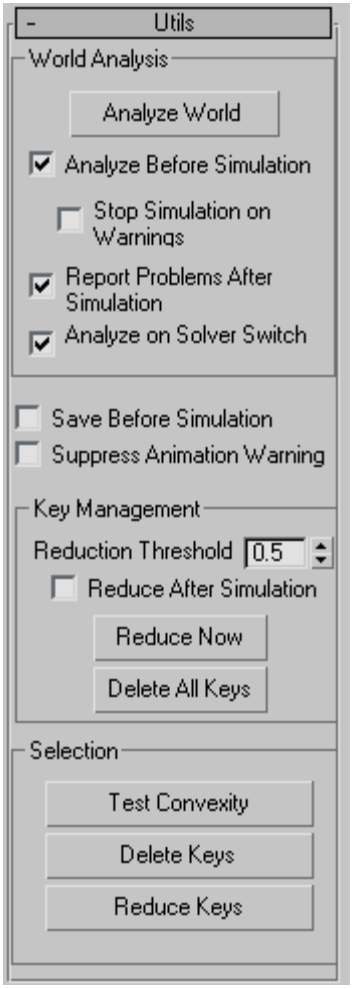
Utility panel > reactor > Utils rollout

menu bar > reactor > Utilities

reactor provides a number of helpful utilities that you can use to analyze and optimize your simulation. For instance, you can check your reactor world for the presence of any unusual physical states that might cause problems with the simulation, such as excessive gravity, and you can remove redundant keyframes from a generated animation.

You can also access these utilities from the reactor menu.

Interface



Analyze World  **—**This utility begins by creating a simulation. If it finds any errors while constructing the simulation, such as invalid numbers of objects in systems or invalid meshes that would stop the simulation from running, it reports these in a dialog. These error checks are always carried out when creating a simulation, and if any of the tests fail the simulation cannot continue.

If your simulation passes the initial error checks, the world analysis begins. This checks for unusual physical states that could cause issues in simulation, such as interpenetrating objects or extreme values for gravity or object density. If any such states are found, the utility reports them in a window and gives you the option to continue or cancel. If you are trying to simulate, choosing Continue runs the simulation; otherwise it closes the window.



Analyze Before Simulation—When on, reactor always calls Analyze World before previewing or running the simulation. Default=on.

Stop Simulation on Warnings—When on, prevents reactor from creating an animation if any warnings occur during analysis. For example, if the analysis finds that a body's property such as Friction exceeds the bounds of real-world values, it halts generation of the animation. Has no effect on the Preview process. Default=off.

Report Problems After Simulation—When on, reactor reports problems detected during the simulation after it's finished. Default=on.

Analyze on Solver Switch—When on, reactor automatically calls Analyze World when you change the solver from the About rollout. Default=on.

Save Before Simulation—When on, reactor saves your scene before simulating.

Suppress Animation Warning—When off, and you click Create Animation, reactor opens an alert warning you that animation creation cannot be undone, and asking you to confirm. When off, reactor simply creates the animation without warning you. Default=off.

Key Management group

After generating an animation with reactor, you often end up with a number of redundant keyframes. For instance, in a scene with a rolling rigid body, creating the initial animation might result in a keyframe for every frame. However, you can actually specify the body's animation with keys at just a few frames (perhaps frames 0, 25, and 90), allowing 3ds Max to interpolate frames between them. This utility allows you to remove any redundant keyframes for specified objects, thus saving memory. This functionality applies only to rigid bodies.

Reduction Threshold—Specifies how aggressive the key reduction is. Increasing this value causes more keys to be deleted, but the animation might lose fidelity. Lowering it retains more keys with greater accuracy, at the cost of memory. Default=0.5.

Reduce After Simulation—When on, reactor automatically applies keyframe reduction each time you simulate.

Reduce Now—Reduces keyframes for all the rigid bodies in the simulation.

Delete All Keys—Deletes all keyframes for all the rigid bodies in the simulation.

Selection group

Test Convexity—Performs a convexity test on the object currently selected in the viewport, allowing you to check whether the object is *convex or concave (page 2–719)* before choosing a *simulation geometry (page 2–719)*.

Delete Keys—Deletes all keyframes for the objects currently selected in the viewport.

Reduce Keys—Reduces keyframes for the objects currently selected in the viewport.

Properties Rollout

Utility panel > reactor > Properties rollout

reactor toolbar > Open Property Editor

For details of this rollout, see *Rigid Body Properties (page 2–717)*.

The Real-Time Preview 

Utility panel > reactor > Preview and Animation rollout > Preview Animation button

Menu bar > reactor > Preview Animation

reactor toolbar > Preview Animation button

It's useful to be able to preview reactor simulations from within 3ds Max. The Preview Window lets you view and interact with a simulation in real time. You can run the simulation, interact with the objects in the scene using the mouse, and even update your objects in 3ds Max with their current state in the preview.

Procedures

To preview a simulation in the Preview Window:

1. Choose one of the above commands.
2. If no errors are found in your current scene setup, the Preview Window opens with your scene.
3. Press P to start the simulation.

   At any time during the simulation, you can orbit around the scene by left-button dragging the mouse within the Preview Window. With Havok 1 and Havok 3, you can pan the window by middle-button dragging; with Havok 3, you can also pan the window by right-button dragging.

To interact with objects in the Preview Window :

1. Position the mouse cursor over the object you want to interact with, and then press and hold the right mouse button (Havok 1) or the Spacebar (Havok 3).

   This creates a reactor spring between the mouse cursor and your chosen object.
2. While holding the right mouse button (Havok 1) or the Spacebar (Havok 3), move the mouse.

   This drags the object, letting you pull it around, unless it's a fixed object, or attached to one, and interact with other objects in the scene.
3. To release the object, release the right mouse button (Havok 1) or the Spacebar (Havok 3).

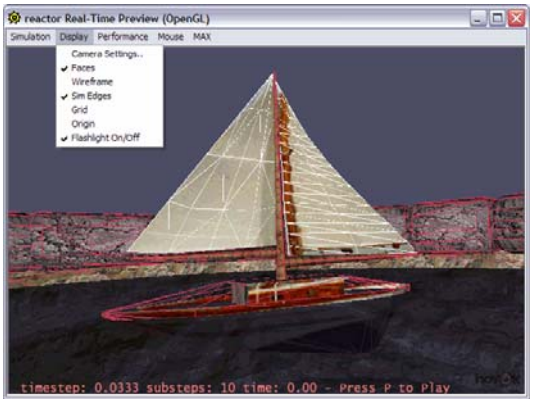
To update 3ds Max from the Preview Window:

1. Preview the simulation as described above.
2. When the scene is in your chosen state (for instance, you might want to drape some clothing over a figure, or let some objects fall and roll on the ground), open the MAX menu and choose Update MAX.

   Note: The scene in 3ds Max will not be refreshed until you close the Preview Window.

Interface

Preview Window Menus



Timing

At the bottom of the window, profiling and timing information is shown :



The information presented is:

- The current *time step (simulation step) (page 2–709)* used for the simulation. This value is initially taken from the *Preview & Animation rollout (page 2–806)* but can also be changed by using the *Performance menu (page 2–817)*.
- The number of *substeps (page 2–709)* used for the simulation. Again, this value is initially taken from the *Preview & Animation rollout (page 2–806)* but can also be changed by using the *Performance menu (page 2–817)*.
- The current simulation time; that is, the time that the current image in the preview represents in the final animation.

During the preview, reactor tries to run the simulation in real time. In other words, if a simulation step of one second takes only 0.7 second of CPU to simulate, reactor "waits" 0.3 second before simulating the next step. That way, the animation is presented in the window at the same speed as the final animation in 3ds Max.

However, for complex scenes or slow CPUs, simulating a particular period of time can take actually more CPU time than the specified period. For example, a simulation step of one second might actually take two seconds to calculate. In that case, it is not possible to present the animation in real time, and the animation is therefore presented during the preview at a slower speed that the final animation will have once it has been created. When this happens, the preview reports it by appending an asterisk (*) to the current time:



Simulation menu

Play/Pause (P)—Starts and pauses the simulation. If the simulation is paused, the display remains active and you can still pan, rotate, and zoom the camera, but the physics world remains still.

Reset (R)—Resets the simulation, returning objects to their initial positions.

Display menu

Important: When About rollout > Choose Solver is set to Havok 3, only Camera Settings is available from this menu.

Camera Settings—Opens a dialog that allows you to specify the near and far clipping planes for the camera, and to change the camera's field of view (F.O.V.). The clipping planes can also be set externally from the reactor utility Display rollout.

Faces—When on, the faces (and not the edges) of the display bodies are rendered.

Wireframe—When on, the edges (and not the faces) of the display bodies are rendered.

Sim Edges—When on, the edges of the physically simulated geometry (the simulation geometry for each body) are rendered in the preview window. This is useful for seeing what's physically happening in a simulation.

Grid—When on, three 2D grids are displayed in the XY, YZ and ZX planes.

Origin—When on, the X (in red), Y (in green) and Z (in blue) axis are displayed at the origin (0,0,0).

Flashlight On/Off—When on, a flashlight located behind the camera lights the scene. When off, lights defined on the Display rollout of the utility are used. If no lights are defined, the preview uses a single fixed light.

Performance menu

Fixed Step (60,50,40,30 fps)—Choosing one of these options changes the frequency of the simulation to the given value; a frequency of 30 fps yields a time step of 1/30 (0.0666) second, 50 fps means a time step of 1/50 (0.02) second, etc.

Substeps (1-100 substeps)—Choosing of these options sets the number of substeps taken in every simulation step.

Mouse menu

This menu is available only when using the Havok 1 engine.

Pick at C.O.M.—During the preview, you can pick and drag objects (see *To interact with objects in the Preview Window : (page 2–815)*. A spring is the attached between the mouse cursor in the screen and the object. When this option is on, the spring is attached to the center of mass of the object. Otherwise the spring is attached to the point on the object where you first right-clicked.

Mouse Help—This option shows some help regarding the use of the mouse to control the camera and mouse picking.

MAX

Update MAX—Takes the position and rotation of the objects in the simulation and uses them to update the objects in 3ds Max.

Use MAX parameters—Resets the substeps and time step (FPS) values to the values set in the reactor utility's *Preview and Animation (page 2–806)* rollout.

Display rollout

The Display rollout of the reactor utility lets you configure various display options for the Preview Window, including cameras and lighting. For more information, see *Display Rollout (page 2–812)*.

Scripts

This section deals with some useful scripts that you can use to speed up your creation of reactor scenes.

If you are interested in creating your own scripts, see the MAXScript Reference, available from the 3ds Max Help menu.

Setup Scripts

These scripts allow you to quickly set up constraints and collections. They are run automatically when you follow the specified procedures.

Constraint Creation

This applies to all the reactor *constraint (page 2–724)* types, except the *Point-to-path constraint (page 2–762)*.

To create a constraint and attach objects in a single step:

This method works with two or more objects only; it doesn't work with a single-body constraint setup such as Spring with a child only.

1. Create the object(s) to constrain, and then ensure they are both selected.

   If you select more than two objects, reactor will create constraints between them. If you link the objects together hierarchically in 3ds Max before creating the constraints, reactor will create constraints between link parents and children. Otherwise, reactor will examine the scene and try to choose appropriate pairs of objects.

2. Create the constraint(s) by using the reactor toolbar, menu or quad menu. The constraint automatically appears in the viewport, and as you will see in the constraint's Properties rollout, your objects are now attached to the constraint.

If the objects are linked in 3ds Max, reactor will automatically make the link's child object the constraint child and the link parent the constraint parent. Otherwise, reactor will examine the scene and try to choose an appropriate parent and child.

In addition, if the constrained objects are already added to a collection, reactor will disable collisions between them (except in the case of *springs (page 2–727)*).

Collection Creation

This applies to all the reactor collection types: *rigid body collections (page 2–723)*, *soft body collections (page 2–788)*, *cloth collections (page 2–781)*, and *rope collections (page 2–792)*. It also applies to the *Fracture (page 2–770)* and *Constraint Solver (page 2–736)* helpers.

To create a collection and add bodies in a single step:

1. Create the object(s) to add to the collection, then ensure that they are all selected.
2. Create the collection by using the reactor toolbar, menu or quad menu. The collection automatically appears in the viewport, and as you will see in the collection's Properties rollout, your objects are now added to the collection.

The Rag Doll Script

This example script lets you quickly set up a simple humanoid character, using a mixture of *Rag Doll (page 2–737)* and *Hinge (page 2–747)* constraints to limit the movement of the character's limbs. You can use the sample humanoid, or create your own character using the provided naming conventions and the script will set up its constraints. Alternatively you can use the script as a basis for creating your own custom rag doll script.

Procedures

To create the sample rag doll:

1. Run this script:
   - menu bar > MAXScript > Run Script > *rctRagDollScript.ms*
   - Utilities panel > MAXScript > Run Script > *rctRagDollScript.ms*

   The Ragdoll dialog opens, with two rollouts: Create Humanoid and Constrain Humanoid.
2. On the Create Humanoid rollout, ensure that your chosen creation options are on, and then click Create Humanoid. A humanoid figure appears in the viewport.



3. To set up the humanoid's constraints, ensure that your new humanoid is highlighted in the Humanoids list on the Constrain Humanoid rollout (by default, it's called Ragdoll) and that your chosen options are active on the rollout.
4. Click Constrain Humanoid.

   The script creates Rag Doll constraints and Hinge joints with appropriate limits for the humanoid. In addition, if Create RBCollection and Create CSolver are on, the script creates these helpers and adds the bodies and constraints to them.

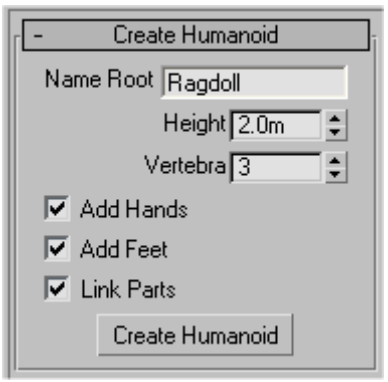
To turn your own humanoid figure into a constrained rag doll:

1. Create your figure using the naming conventions described in *How It Works (page 2–820)*.
2. Run the script as described above.

   The script searches the scene for geometry that follows the naming conventions and adds your figure to the Humanoids list.
3. Ensure your figure is highlighted in the list and then click Constrain Humanoid.

Interface

Create Humanoid rollout



Name Root—This text box stores the root name for the objects to be created. For example, the Name Root "RagDoll" generates objects with names such as RagDoll Pelvis and RagDoll R UpperArm.

Height—The height of humanoid to be created.

Vertebra—The number of vertebra to be created in the character's spine.

Add Hands—When on, reactor creates boxes for the character's hands.

Add Feet—When on, reactor creates boxes for the character's feet.

Link Parts—When on, reactor links the created parts as a 3ds Max hierarchy.

Create Humanoid—Click to create the humanoid objects and add them to the scene.

Constrain Humanoid rollout



Humanoids—Displays a list of potential humanoids found in the current scene that could be used for creating rag dolls. In generating this list only the root node is examined so it is, in effect, a list of valid root nodes in the scene. If the rest of the hierarchy does not exist for a given root, when you use the root to create a rag doll the software outputs errors to the MAXScript listener.

Vertebra—The maximum number of Rag Doll constraints that the script creates for the character's spine. If this value exceeds the number of valid vertebrae in the selected humanoid, the script won't create the excess constraints.

Create RBCollection—When on, the tool creates a *Rigid Body Collection (page 2–723)* and places all the parts of the newly created rag doll character into it.

Create CSolver—When on, the script adds a *Constraint Solver (page 2–736)* to the scene, which is then associated with the freshly created Rigid Body Collection. All the hinges and rag doll constraints that are created for the selected humanoid are added to this. This option is available only if you have chosen to create a Rigid Body Collection.

Add Hands—When on, if the selected humanoid has hands, the script creates constraints to connect the humanoid's hands to the rest of the body.

Add Feet—When on, if the selected humanoid has feet, the script creates constraints to connect the humanoid's feet to the rest of the body.

Snapshot Parts—When on, the script creates a separate geometry object for each object to be attached to a constraint. It is the newly created snapshot that is then used in the constraint and added to the collection (if a collection is being created).

Link To Original—When on, the newly created snapshot geometries are linked to the original geometries. Available only when Snapshot Parts is on.

Change Name—When on, each geometric component of the new rag doll is prefixed with the text in the New Name field. Available only when Snapshot Parts is on.

New Name—When on, each geometric component of the new rag doll is prefixed with the text in the New Name field.

Constrain Humanoid—Click to create the Rag Doll constraints and, depending on the active options, also create a copy of the source objects, a Rigid Body Collection, and a Constraint Solver.

How It Works

The script works by assuming that you've used the following naming conventions for the rag doll parts. This naming convention follows that used by the **character studio** Biped object to name the limbs. However, as long as they are named as following, you can use any geometry created in 3ds Max. The script scans the selection for geometry with names that include one of the following:

- _R_UpperArm
- _R_ForeArm
- _L_UpperArm
- _L_ForeArm
- _R_Thigh
- _R_Calf
- _L_Thigh
- _L_Calf
- _Pelvis
- _Spine
- _Head

The script assumes that there is only one of each type of part, except in the case of the spine, which can have up to five parts.

At this point the adjoining parts are connected, using mostly *Rag Doll constraints (page 2–737)*. However, the script also uses *Hinges (page 2–747)* for the elbow and knee joints, as using Rag Doll constraints would be overkill to achieve the simple constrained motion found in an elbow joint. The script also sets suitable limits for each constraint.

In addition, the script sets the mass of each rigid body in the character to 10.0 kilograms, and the simulation geometry to be a bounding box. It also disables collisions between constrained pairs of bodies.

Frequently Asked Questions



See also: *Troubleshooting (page 2–823)*

Can I assign initial velocities to objects? How?

Yes. At the start of the simulation, reactor assigns to each object the velocity it currently has in 3ds Max. In other words, if an object is moving during the frame range A to B, and you start the reactor animation setting *Start Frame (page 2–806)* in between A and B, the object will start with the velocity it had in your animation.

Warning: Make sure the object has the desired velocity at the exact frame where reactor starts the simulation (Start Frame). Since 3ds Max usually decelerates an object when animating its movement using keys at time A and B, starting the reactor animation at time B won't add any velocity to the reactor simulation. In order to have initial velocity in the reactor animation, start at a frame between A and B.

Is it possible to modify the animation in Cloth/Soft/Rope/Water?

The animation for Cloth, Soft (mesh), and Rope is internally stored in the respective modifier using a memory-optimized format. Unfortunately, there is no direct access to those keyframes. However, you can use the *Point Cache modifier (page 1–758)* to do some manipulations.

There is some MAXScript access to Water keyframes. For details on MAXScript, choose Help > MAXScript Reference.

What happened to Point-Nail constraint?

The Point-Nail constraint from reactor 1 has been superseded by generalizing the *Point-Point constraint (page 2–750)*. The old Point-Nail constraint is equivalent to a-single bodied Point-Point constraint

Can I apply the Cloth/Soft/Rope modifier to part of a mesh?

While the *Soft, Cloth, and Rope (page 2–777)* modifiers apply to all the vertices, you can fix or keyframe vertices of the mesh or spline by using *deformable constraints (page 2–795)*. Fixed or keyframed vertices will not be simulated physically

Can I apply the Cloth/Soft/Rope modifier to more than one object?

In general, no, because the animation is stored explicitly in a single modifier. The only exception is the *FFD soft body (page 2–786)*: If you have an FFD modifier applied to more than one object, you can apply a reactor *Soft Body modifier (page 2–784)* to any (or many) of the bodies. You should still add only one of the objects to the *Soft Body Collection (page 2–788)*.

Which parameters can I animate?

Only *Wind (page 2–803)* parameters are animatable.

Can a keyframed object participate in the simulation?

Yes. For objects that don't change shape, and for which reactor should not create keyframes, use *unyielding* rigid bodies. For objects that change

shape (deform), use *deforming meshes (page 2–793)* (

Note: Unyielding rigid bodies are much faster and simpler to simulate than deforming meshes.

Can reactor store collisions for deformable bodies (Soft, Cloth, Rope)?

No. The *storing collisions (page 2–774)* functionality applies only to rigid bodies.

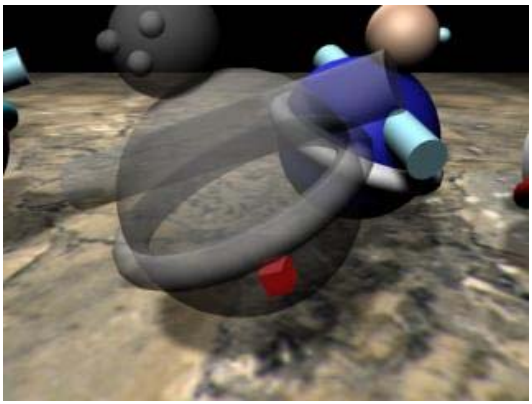
What are the limitations of *Water (page 2–801)* in reactor?

Water in reactor is simulated as a height field (points are displaced vertically only). This can produce effects like ripples and waves, but it cannot reproduce certain other effects:

- Wakes
- Splashes. You can, though, use the information exposed through MAXScript to generate particles or other effects.
- Flowing fluids (like water flowing through a pipe or pouring from a tap).

Can I move the center of mass of my objects? How?

You can't explicitly set the center of mass of an object. However, you can use compound rigid bodies where different pieces have different masses to simulate uneven mass distribution. For example, you can simulate an object with a low center of gravity by dividing the object into two pieces (primitives), top and bottom, where the bottom piece is heavy and the top piece is light; alternatively, you leave the original piece intact, with low mass, and add a small, heavy primitive at the point of the desired center of mass. For more information, see *Compound Rigid Bodies (page 2–722)*.



Can I attach Cloth/Rope/Soft Bodies to other Cloth/Rope/Soft Bodies?

No. However, once a piece of, say, cloth has been simulated, you can take it out of its Cloth Collection and add it to a *Deforming Mesh Collection (page 2–794)*. By doing so, the cloth will still be part of the simulation the next time, but reactor won't recalculate its animation. You can then attach other pieces of Cloth/Soft/Rope to it by using the *Attach To Deforming Mesh (page 2–799)* constraint.

How can I make a rigid body not move until a collision happens?

You can use the *Inactive* property for the rigid body.

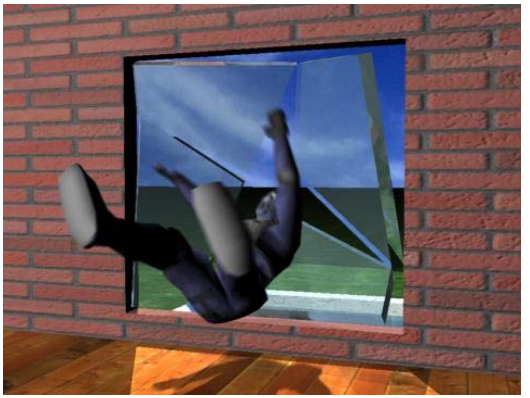
How can I animate a floating object, like a balloon?

There are many alternatives you can try:

- Use *water (page 2–801)*. The buoyancy of a balloon in the air follows the same rules as the buoyancy of a light object inside a fluid. You will need to place the water surface well above the scene so the objects are not visible when they reach it.
- Use very low mass for the objects, and add *wind (page 2–803)* blowing in the up direction.

You might want to also decrease the strength of gravity.

Troubleshooting



Check also : *Frequently Asked Questions (page 2–821)*

General Troubleshooting

Analyzing the World

Whenever you get unexpected behavior from reactor, it is advisable to perform a *world analysis* to get a report of anything detected by reactor as suspicious, unexpected, or prone to error in your scene. For more information, see *reactor Utilities (page 2–813)*.

Discussion Forums

The Autodesk Web Forums provide you with a moderated community to discuss, learn, and share techniques and ideas with other 3D professionals. Need to know how to do something? Having a problem? Just want to show off your work? Post your questions to the appropriate forums.

Simulation takes too long.

General tip:

- Before performing a CPU-intensive operation such as Create Animation or Reduce Keys, make sure you don't have any object selected in the scene. Having an object selected during some calculations can force 3ds Max to evaluate the selected object many times in order to update its UI, slowing down the calculations.

  If the scene is complex but contains independent parts that don't interact with each other, try adding the objects involved to different collections and solvers and create the animations independently for each subset, enabling and disabling the collections and solvers accordingly.

If you are using *Soft Bodies, Cloth, or Rope (page 2–777)*:

- Try to reduce the complexity of the mesh that flows to the reactor modifier. For example, you can apply a MeshSmooth modifier on top of the reactor modifier to render a smooth surface while simulating a simpler one.
- For soft bodies with complex meshes, using *FFD soft bodies (page 2–786)* can offer better results and decrease the computation time.
- Unless necessary, do not use the Avoid Self-Intersections option, because it increases computation time and memory usage significantly.
- Try to reduce the number of *substeps (page 2–806)* used in the simulation. Also try to decrease the number of Internal Steps used by the Cloth/Soft/Rope collection.

  Note: These steps also reduce the accuracy of the simulation, and can therefore cause instability.

If you are using *rigid bodies (page 2–716)*:

- For best results when using rigid bodies only, set Utilities panel > reactor > About rollout > Choose Solver to Havok 3.

- Reduce the complexity of the geometry used for the rigid bodies. Try using bounding boxes or sphere, and avoid using non-fixed concave objects. Use simpler geometry as a *geometry proxy (page 2–719)* for the object.
- If you are using Fracture, slowdowns can be caused by reactor trying to avoid instability. Follow some of *these tips (page 2–773)* to help reduce the chances of instability happening
- Try decreasing the number of *substeps (page 2–806)* used in the simulation. Be aware that this will reduce the accuracy of the solution and therefore may cause instability.

Water doesn't render.

The reactor *Water (page 2–801)* is a world-space modifier and therefore it is not a renderable object. You can bind renderable object (like planes) to the water WSM in order to render a water animation. See *Rendering Water (page 2–803)*.

Cloth/Soft/Rope becomes unstable and explodes.

- Check whether any unyielding rigid body or *DefMesh (page 2–793)* is forcing the Cloth/Soft/Rope into a penetration.
- If you are *attaching to a rigid body (page 2–798)*, make sure that the points of attachment are outside the rigid body; otherwise, make sure Ignore Collision is on.
- If your Cloth/Soft/Rope is *attached to a rigid body (page 2–798)*, use Runge-Kutta as the ODE Solver for the *Rigid Body Collection (page 2–723)* of the attached rigid body
- Try increasing the amount of internal steps in the Cloth/Soft/Rope collection (Advanced rollout).

Rigid bodies become unstable/explode/ fly to infinity.

In rare situations, simulation can become unstable and objects can "explode." This typically happens only when using complex systems of objects, like *Fracture (page 2–770)*, or using *simple constraints (page 2–727)* like springs or dashpots.

- For best results when using rigid bodies only, set Utilities panel > reactor > About rollout > Choose Solver to Havok 3.
- If you are using many *simple constraints (page 2–727)* (spring, dashpots) working together, switching the ODE Solver in the *Rigid Body Collection (page 2–723)* to Runge-Kutta can improve the results.
- Usually, increasing the number of *substeps (page 2–806)* in the simulation improves the stability of the simulation.
- Also for springs and dashpots, avoid attaching objects with very different masses. Using a Strength value similar to the mass of the objects attached, and a Damping value of 1/10th of that Strength usually give good results.
- If you are using *constraints (page 2–724)*, make sure you have aligned the *constraint spaces (page 2–725)* properly.
- If many objects start the simulation too close too each other (inside *collision tolerance*), reactor will try to push them apart at the beginning of the simulation. Try increasing the space between them, or reducing the collision tolerance.
- If you have constrained two rigid bodies so they are in continuous contact/penetration (like an upper arm and a forearm), be sure to *disable collisions (page 2–810)* between those bodies.
- Adding *drag action* can help damp the overall simulation. Try increasing the linear drag in the scene.

- If you are working with Fracture, the *Fracture Tips (page 2–773)* section provides information on how to avoid instability.

Cloth/Soft/Rope object stretches too much

If you have attached (using the *Attach to RB Constraint (page 2–798)*) it to a falling object, be sure to turn off Do Not Affect Rigid Body. If the problem remains, increase the mass of the Cloth/Soft/Rope, or decrease the mass of the attached rigid body. Increasing the strength and damping of the Cloth/Soft/Rope should also improve the situation.

Objects seem to move very slowly

Make sure you are using the proper *World Scale (page 2–808)* in the scene. A 10cm box falling from a 1m height doesn't behave the same as a 1km box falling from a height of 10km: The second one will reach the ground much later). Use real-world units or have a proper mapping between 3ds Max units and simulation units (*World Scale (page 2–808)*).

Why, after simulation, do I get warnings regarding "topology changes"?

The term "topology" refers to the internal connections of the vertices, edges and faces of a mesh. If you change the number of vertices in a mesh, you are changing its topology. While 3ds Max lets you animate the topology of objects (e.g., change the number of vertices or faces over time), reactor can simulate only objects with fixed numbers of vertices. If, during a simulation, reactor detects a topology change, it ignores the object in question and reports the problem as a warning.

Objects interpenetrate.

There can be different reasons for this:

- Check that you haven't disabled collisions for those bodies. You can disable collisions either *globally (page 2–810)* for pairs of objects, or for the object in particular (*Rigid Body Properties (page 2–717)*).
- If the objects are attached using *deformable constraints (page 2–795)*, make sure you haven't disabled collisions in the constraint parameters.
- Make sure *Collision Tolerance* is set to a value that makes sense for the sizes of the objects.
- Try increasing the number of *substeps (page 2–806)* in the simulation.

Cloth/Soft/Rope intersects itself.

Make sure you have turned on the Avoid Self-Intersections option in the *modifier (page 2–777)*.

I can't add an object to a Cloth/Soft/Rope Collection.

Only an object that has the proper reactor modifier applied can be added to one of these collections. Make sure you have applied a reactor *Cloth (page 2–778)*, *Soft Body (page 2–784)*, or *Rope (page 2–789)* modifier to the object.

I can't add an object to a Rigid Body Collection.

If you are using *compound bodies (page 2–722)*, make sure the group is open, as only the group parent can be added to a *Rigid Body Collection (page 2–723)*.

Mass and other properties are disabled.

Make sure you have selected only one object, and that this object is geometrical. The most common cause for this is the use of groups. You can use groups to create *compound rigid bodies (page 2–722)*, but in order to change the properties of

the primitives inside the group, you'll need to open it and selected the individual primitives inside it.

Rigid bodies don't bounce enough.

Increase the *Elasticity property (page 2–717)* of the object. Try also increasing the Elasticity of the other object involved in the collisions, as the elasticity of a collisions is a function of the *elasticity property (page 2–717)* of the two objects involved.

Objects stop moving unexpectedly in the middle of a simulation.

This can be caused by objects being deactivated. There are two ways objects can be deactivated: through the *world (page 2–808)* deactivation or, if the object belongs to a constrained system, through the *Constraint Solver (page 2–736)* deactivation parameters

In both cases, try modifying the deactivation parameters for a less-aggressive deactivation, or disable deactivation in general.

An object doesn't move during the simulation.

Make sure that you've added the object to a *RBCollection (page 2–723)* and that it has a *mass (page 2–717)* other than 0.0 (objects with mass 0.0 are fixed). If you are using *Fracture (page 2–770)*, make sure none of the pieces is fixed (has mass 0.0).

A keyframed object doesn't move during the simulation.

Make sure that you enabled the *unyielding (page 2–717)* property for that rigid body.

My old scripts don't work.

reactor exposes parameters and functionality to MAXScript in a much more consistent and complete way than previously. For the update, some reactor MAXScript interfaces had to be renamed, but all functionality exposed in previous versions is still available, and a lot more functionality has been exposed. While this means that some reactor-specific scripts need to be updated, it also means that a great deal more flexibility and power is now achievable using MAXScript with reactor. For more information, see the MAXScript Reference, available from the Help menu.

I can't I disable collisions between two bodies.

If your objects don't appear in the *Define Collisions dialog (page 2–811)*, make sure you've added the two bodies to a collection; only objects in collections are shown in that dialog.

Rigid bodies don't sink/don't float in the water.

Objects sink or float depending on their density and the density of the water. You define water density explicitly in the *Water space warp (page 2–801)*. The density of a rigid body is defined as its mass divided by its volume.

- Check that your objects have real-world sizes (or adjust the *world scale (page 2–808)*) and *masses (page 2–717)*.

If your bodies still don't sink enough for your liking, try the following:

- Increase the mass of the bodies.
- Decrease the density of the water.
- Decrease the viscosity of the water, this can slow down the sinking.

Follow the opposite steps if you're trying to make your objects float more.

Water doesn't ripple enough

You can either increase the Max Ripple value in the *Water space warp (page 2–801)*, or you can scale the effect of the space warp on a particular object

that you've bound to the space warp by modifying the Scale Strength value in the *Water world-space modifier (page 2–803)*. The first option (increasing the Max Ripple value) is preferable, but you need to re-create the animation in order to see the effect. The second option (increase the Scale Strength parameter in the modifier) affects only the object bound to the Water space warp. It can thus produce inconsistent results if you have objects floating in the water, because they won't be affected by the change.

I can't see any stored collisions.

Make sure you've set reactor to *store collisions (page 2–774)* during the simulation. Collisions are stored only for rigid bodies; not for Cloth, Soft Bodies, or Rope.

Constrained objects snap at the beginning of the simulation.

During a simulation, reactor tries to match the two constraint spaces (sub-objects) defined in the constraint (see *Constraint Concepts (page 2–725)* for more info). When you create a constraint those spaces are aligned by default. After setup, each space moves with the corresponding attached object. If you have moved the objects and wish to realign the constraint spaces, there are tools in the helper to do so (see *Working With Constraint Spaces (page 2–726)* for more info).

File size increases dramatically.

- If you are using rigid bodies, try the *Reduce Keys (page 2–813)* utility to reduce the number of keyframes created for each rigid body.
- *FFD-based soft bodies (page 2–786)* can also create a large number of keys in the FFD modifier (one for each vertex for each frame). Although the *Reduce Keys (page 2–813)* functionality in reactor applies only to rigid bodies, you can reduce keys in the FFD modifier using the 3ds Max Reduce Keys tool, available from *Track View (page 2–501)*.

reactor runs out of memory during simulation.

Reduce the complexity of the objects simulated:

- For rigid bodies, use a simpler *geometry or proxy (page 2–719)* to simulate the object.
- For Cloth and Soft Bodies, try to use a coarse mesh for the simulation, then apply a MeshSmooth modifier *on top* of the reactor modifier to smooth the final result.
- For Soft Bodies, try to use *FFD-based Soft Bodies (page 2–786)* instead of mesh-based soft bodies.

For any deformable body, the Avoid Self-Intersection option is particularly memory-intensive. Avoid using this option if you are having problems regarding memory.

Using FFD-Based Soft Bodies to animate a part of a mesh, the whole object seems to be rotated by the simulation.

If you are using *FFD Soft Bodies (page 2–786)* to animate just a subpart of a mesh, make sure the *Animate Transform* check box is off, as this option works only for FFD animating the whole mesh.

Using FFD-Based Soft Bodies to animate a whole mesh, some weird squashing happens on occasion.

If you are using *FFD Soft Bodies (page 2–786)* to animate the whole mesh, sometimes the FFD modifier deformation can show artifacts when the lattice points are rotated. Try switching on the *Animate Transform* check box in the modifier to avoid those artifacts.

Picking objects during preview makes them go unstable.

Try reducing the strength and damping of the mouse spring on the *Display rollout (page 2–812)*.

Objects barely move when picking during preview.

Try increasing the strength and damping of the mouse spring on the *Display rollout (page 2–812)*.

I can't see anything in the Preview Window.

It might be caused by the lack of a proper camera; create a camera and select it in the *Display rollout (page 2–812)*. If still doesn't work, try changing the far and near clipping planes.

The preview shows everything in black.

If the scene uses texture maps but the software can't find the map files, they're shown as black. Disable textures by pressing X during the *preview (page 2–815)*.

After the simulation I get a warning reporting that objects moved to infinity.

When instability occurs, rigid bodies can reach extreme velocities and therefore move to infinity. If reactor detects this condition when creating an animation, it does not create keyframes, but instead reports the problem once the animation is finished.

Check also the *Rigid bodies become unstable (page 2–824)* section above.

There is a visible gap between rigid bodies in the simulation.

Try reducing the *Collision Tolerance* of the simulation.

reactor complains about a no longer existing node.

This can happen if you rename an object that is used as geometry or display *proxy (page 2–719)* for a rigid body; proxies are stored by name, so if you change the name of the proxy you'll have to manually update the rigid bodies that use it.

Error : All vertices are coplanar, use concave mesh.

By default, geometries in reactor are simulated by using their convex hull. reactor can't create a convex hull from geometry in which all the vertices lie in the same plane, such as a standard 3ds Max plane. Those geometries need to be simulated as concave. Check the *Rigid Body Properties (page 2–717)* section for more information.

Rope simulation modifies the smoothness of the spline.

Although reactor simulates cloth as a series of linked vertices, the *Rope modifier (page 2–789)* tries to produce a smooth curve going through the vertices. In most cases the curve generated by the modifier is suitable, but in some cases (for example, if the vertices in the spline are far apart) you might find that the generated curve is not optimal. In those cases you can:

- Modify (refine) the spline below the Rope modifier so there is a uniform distance, not too great, between vertices. You will need to re-create the animation in order to see the final results.
- Modify the spline after the Rope modifier using an *Edit Spline modifier (page 1–680)*. This is the most flexible option.

Animation works in the preview, but not in 3ds Max after creating the animation.

reactor creates the animation for *rigid bodies (page 2–716)* by sending transforms (new rotation

and new position) to the controllers assigned to those bodies. Some controllers, like procedural controllers (Noise) and some system controllers (Biped) do not handle these explicit transform updates. To overcome this, assign a regular keyframable controller such as Euler XYZ or Bezier to the objects or, if that's not possible or desirable, create a copy or snapshot of the bodies and make sure that a suitable controller is assigned to the snapshot copies. Use these copies for the reactor simulation instead.

Rigid bodies in animation flicker and jump when playing the animation but not when moving the time slider.

This is usually a problem with the subframe interpolation of the Euler (XYZ) rotation controller. While reactor sends valid transformation keys to the controller, the conversion from an arbitrary rotation to X,Y and Z Euler angles may cause non-smooth animations between keyframes (the controller produces noncontinuous X,Y and Z values). This problem can be solved by using a quaternion-based rotation controller, such as the TCB controller, for the rigid bodies.

Preview Animation fails

While the *Preview Window (page 2–815)* should work with most OpenGL and DirectX compliant hardware, some specific hardware and driver configurations might cause problems:

- Update to the latest drivers for your graphics card. Refer to the manufacturer's Web site for details.
- Try to use DirectX instead of OpenGL (or vice versa) for the preview. This setting, independent of what is used for the 3ds Max viewports, can be changed on the *Display rollout (page 2–812)*. If you discover that your card works only with a particular configuration, you may want to create startup script (a script in the *scripts/startup* folder) that calls reactor.displayUsingDirectX = true/false.
- Some problems can be caused by lack of display memory. Avoid having multiple sessions of 3ds Max open at the same time, and reduce the size of the windows and the display resolution.
- Try changing the 3ds Max to different drivers: OpenGL, Direct3D or Software.

Wind doesn't seem to work

- Make sure the Wind Helper icon is pointing towards the right direction.
- If the wind should act upon cloth, rigid bodies, etc., make sure the correct check box is on in the Wind Helper parameters.
- Try to increase the strength of the wind, or decrease the weight of the objects it should affect.

14 character studio

The **character studio** functionality in 3ds Max provides professional tools for animating 3D characters. It is an environment in which you can quickly and easily build skeletons and then animate them, thus creating motion sequences. You use the animated skeletons to drive the movement of geometry, thus creating virtual characters. And you can generate crowds of these characters using **character studio**, and animate crowd movement using a system of delegates and procedural behaviors.

character studio comprises three components: Biped, Physique, and Crowd.

- *Biped (page 3–916)* builds and animates skeletal armatures, ready for character animation. You can combine different animations into sequential or overlapping motion scripts, or layer them together. You can also use Biped to edit motion capture files.

  Note: Biped does not create character mesh objects. You should create your character mesh before using Biped to create a skeleton for it.

- *Physique (page 3–994)* uses the biped armatures to animate the actual character mesh, simulating how the mesh flexes and bulges with the movement of the underlying skeleton.

  Note: You can use Physique with other hierarchies beside the biped skeleton.

- *Crowd (page 3–925)* animates groups of 3D objects and characters using a system of delegates and behaviors. You can create crowds with highly complex behaviors.

If you are new to 3ds Max, read *What You Should Know to Use character studio (page 2–832)* and *Understanding character studio Workflow (page 2–839)*.

If you are already a 3ds Max user, proceed to *Understanding Biped (page 2–833)* to continue.

To find out about Physique, read *Understanding Physique (page 2–834)*.

To learn about the Crowd system, read *Crowd Animation (page 2–1154)*.

What is character studio?

character studio is a set of components that provides a full set of tools for animating characters.

character studio lets you create skeletal hierarchies for two-legged characters (called *bipeds*) that can be animated quickly using a wide variety of methods. If your character walks on two legs, the software provides unique *footstep animation (page 2–856)* that automatically creates movement based on gravity, balance, and other factors.

If you want to animate motions manually, you can use *freeform animation (page 2–886)*. This type of animation is also suitable for characters with more than two legs, or characters that fly or swim. With freeform animation, you can animate the skeleton with traditional inverse kinematic techniques.

You can also animate a biped skeleton by *loading motion capture files (page 2–925)*. You can convert back and forth between these methods to take advantage of their several benefits.

character studio also provides tools for skinning a character using the biped skeleton, or any other type of linked hierarchy.

The software provides a unique ability to separate the motion of the animation from the structure of the character. This means you can animate a giant walking and apply that motion to a tiny elf, or animate a fat character and retarget the same motion to thin one. Using a library of motions, you can animate a character doing hundreds of different actions, as easily as loading a file.

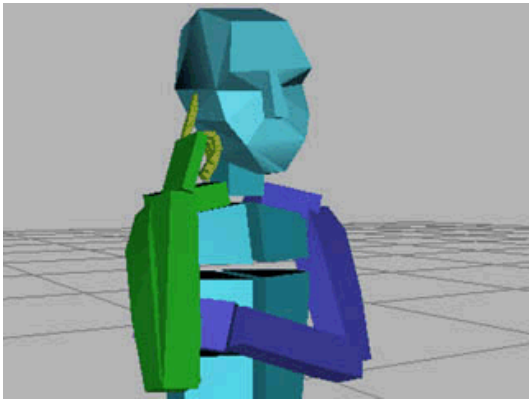
The software provides a comprehensive range of tools for motion editing. You use motion scripting to sequence animations with transitions. Different animations can be overlaid using layers, or blended with the new nonlinear *Motion Mixer (page 2–604)*.

You can use function curve editing, available on the track bar or in Track View, or in the new *Animation Workbench (page 2–1008)*, to alter animations. The workbench also provides specialized tools to analyze and correct errors in motion. Finally, the software provides options for *creating crowds of bipeds (page 2–1154)* or other objects with a procedural animation system that uses forces and behaviors to drive the characters' motions.

character studio consists of three basic components:

- **Biped®** provides tools to create and animate skeletons.
- *The **Physique®** modifier (page 2–1076)* associates a skeleton with the character's mesh objects, so you can control the mesh with the skeleton.
- **Crowd** provides tools to create and animate crowds of animated objects, including bipeds.

What You Should Know to Use character studio



character studio provides you with a broad range of tools for animating virtually any type of character. To expedite your introduction to the product, you should be familiar with the following 3ds Max concepts:

- Creation, transformation, and modification of objects
- Selection of objects through clicking or dragging in viewports, and by using the *Select Objects dialog (page 1–78)*
- Navigation in viewports and changing your viewport configuration
- Use of the Track View - Dope Sheet and Track View - Curve Editor, and the trackbar to view and edit animation tracks and keys

For information on these areas, consult the remainder of this reference. If you don't know how to use 3ds Max, do some of the introductory tutorials found online in Help > Tutorials.

Understanding Biped

Biped (page 2–843) is a 3ds Max component that you access from the Create panel. Once you create a biped, you animate it using the Biped controls on the Motion panel. Biped provides the tools you need to design and animate the figures and motion of characters.

The Biped

The biped skeleton created with the Biped module is a two-legged figure created as a linked hierarchy, and designed for animation. The biped skeleton has special properties that make it instantly ready to animate.

Figure and Keyframe Modes

character studio is designed to interchange motion and characters. In *Figure mode (page 2–982)*, you pose the biped to fit your character model. In *Keyframe mode (page 3–960)*, you animate the skeleton. Motions created for the biped can be saved and loaded onto other biped skeletons with completely different physical characteristics. For example, you could animate a giant ogre, save the animation, and load it onto a small child. Motion files are saved in the proprietary **character studio** *BIP format (page 3–916)*.

These files can be used in a variety of ways with Motion Flow, the Motion Mixer, or the Crowd tools to combine animation or animate multiple characters at once.

Animating the Biped

There are two primary methods used in creating biped animation: *footsteps method (page 3–943)* and *freeform method (page 3–945)*. Each method has advantages. You can convert from one method to the other, or you can use a combination of both techniques in a single animation. For detailed information, see the sections that follow: *Creating Footstep Animation (page 2–856)* and *Creating Freeform Animation (page 2–886)*.

Biped Properties

The biped skeleton has a number of properties designed to help you animate faster and more accurately.

- **A human structure—**Joints on the biped are hinged to follow human anatomy. By default, the biped resembles a human skeleton and has a stable inverse kinematics hierarchy. This means that when you move a hand or foot, the corresponding elbow or knee orients itself accordingly, and produces a natural human posture.
- **Customizable for non-human structures—**The biped skeleton can easily be made to work with a four-legged creature or an animal that naturally leans forward, such as a dinosaur.
- **Natural rotations—**When you rotate the biped spine, the arms maintain their relative angle to the ground, rather than behaving as if they were fused to the shoulders. For example, say the biped is in a standing position, arms hanging at its sides; when you rotate the spine forward, the biped's fingers will touch the ground rather than point behind it. This is a more natural position for the hands, and this speeds the process of keyframing the biped. This feature also applies to the biped head. When you rotate the spine forward, the head maintains a forward-looking orientation.

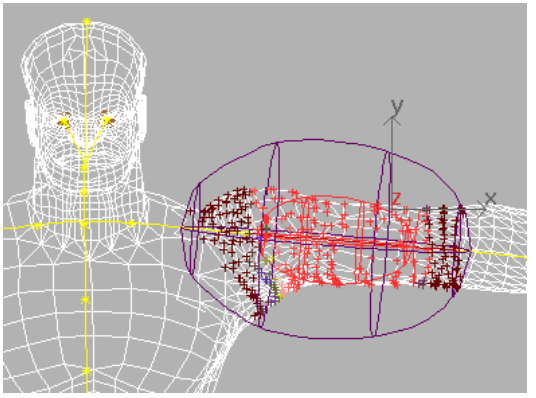
- **Designed for footsteps—** The biped skeleton is specially designed to animate with **character studio** footsteps, which help solve the common animation problem of locking the feet to the ground. Footstep animation also provides an easy way to rough out animation quickly. See the section *Creating Footstep Animation (page 2–856)*.

See also

Biped (page 2–843)

Biped User Interface (page 2–932)

Understanding Physique



Physique is a modifier that, when applied to a mesh, allows the movements of an underlying skeleton to move the mesh seamlessly, like bones and muscle under a human skin. Physique works on any point-based objects including geometric primitives, editable meshes, patch-based objects, NURBS, and even FFD space warps. For NURBS and FFDs, Physique deforms the control points, which in turn deform the model. It will attach to any skeleton structure including a biped, 3ds Max bones, splines, or any 3ds Max hierarchy. When you apply Physique to the skin object(s) and attach the skin to the skeleton, Physique determines how each component of the skeleton influences each vertex of the skin, based on settings you specify.

Physique affects a mesh after you click Attach To Node on the Physique rollout and select a root node in the viewports. During the attach process, Physique works its way through all of the children in a hierarchy, starting at the object you select, and creates its own links with associated *envelopes (page 2–1085)* for each link it finds. The links created by Physique are referred to within this documentation as the Physique *deformation spline (page 3–927)*. Vertices that fall within envelopes are influenced to follow the links and animate the mesh. Splines and 3ds Max bones can also be added using the Add button in the *Floating Bones rollout (page 2–1110)*.

Biped and Physique

When the biped pelvis is selected in the viewports and *Attach To Node (page 2–1106)* is turned on, Physique traces its way from the pelvis down the legs to the toes. From the pelvis it also traces its way up through the spine and branches at the collar to the arms, hands, and fingers, and up the neck to the head. A link and associated envelopes are created for each link found. If any other objects, including 3ds Max bones, are linked to the biped, Physique treats them similarly: it creates a new link and envelopes.

Keep this in mind when you use Physique to attach a mesh to the biped; if your character has additional limbs, link 3ds Max bones to the biped for the extra arms before using Attach To Node to create links and envelopes. When Physique is applied and Attach To Node is used, it creates links and envelopes for all the links in the biped, as well as for the linked bones. Objects that should not deform, like skin, but need to be linked to the biped, like a sword, should be linked *after* Attach To Node is used to link a mesh to the biped. This way, a link and envelopes will not be created for the sword.

Envelopes and Weighted Vertices

Envelopes are the Physique modifier's primary tool for controlling skin deformation; tendons and bulge angles are used to fine tune mesh deformation after envelopes are adjusted. All envelopes have an inner and outer bound (boundary). Vertices falling within the inner bound of a single link receive a full weight of 1.0 from that link. Those falling outside the outer bound receive no weight from that link. Vertices falling between the inner and outer bounds receive a weight between 0 and 1.

Vertices move together with the link that influences them. Where multiple envelopes encompass a vertex, that vertex receives weight from each envelope and follows each link to an average position based on these weights. This weighting from multiple links is considered *blending*. It is possible that weights assigned to some vertices don't reach a total weight of 1.0 or greater. Rather than leaving these vertices behind, Physique by default normalizes them to a value of 1.0.

Adjusting falloff, overlap, scale, and other envelope parameters changes vertex weight distribution across links. This, in turn, changes the way skin behaves as the biped moves. Much of the work you'll do to correct the way skin deforms on a character will be to adjust envelopes.

Deformable and Rigid Envelopes

There are two Envelope types per link, *deformable (page 3–927)* and *rigid (page 2–1085)*. Deformable envelopes follow the Physique deformation spline that runs through the joints in the hierarchy, and can be deformed using bulge angles, tendons, and link parameters. Rigid envelopes determine vertex-link assignment based upon the linear 3ds Max link and move in an immobile relationship to the link. Vertices in a rigid envelope, however, are deformed (blended) in the overlap area of other envelopes.

Typically you use deformable envelopes; however, game developers with game-engine restrictions may want to use rigid envelopes exclusively. Both rigid and deformable envelopes can be turned on for the same link. For example, by scaling both envelopes, you could deform the shoulder with a rigid envelope and the armpit with a deformable envelope.

The Number of Links That Can Affect a Vertex

Any number of overlapping envelopes (*N Links (page 3–978)*) can influence vertices. Normally, N Links are preferred. For special purposes such as games requirements, you can limit the number of links (envelopes) that can affect a vertex. The No Blending parameter is similar to the method used in version 1 of the software; a vertex is assigned to only one link.

Physique Workflow

Before Physique is applied, align the biped to the mesh in *Figure mode (page 2–982)*. Use a pose with the arms outstretched so the hands are away from the torso. Save a figure file, so it's easy to return to this pose whenever you need. Select the mesh and choose Physique in the Modify panel. Turn on Attach to Node and select the root node in the hierarchy (biped Pelvis or root node in a bones hierarchy, not the COM). In the Physique Initialization dialog, click Initialize to create default envelopes based on the links in the hierarchy. The remainder of the work is adjusting envelopes and optionally adding bulge angles and tendons.

Envelope size, overlap, and other parameters are adjusted with the character in an animated position (with Figure mode turned off). By scrubbing the time slider back and forth, you can spot problem areas and adjust the envelopes affecting the problem areas. In Place mode is

useful to keep the character stationary during envelope adjustment.

Link parameters, Bulge angles and tendons are the finishing touches. Skin sliding, the amount of twist, and crease blending as a character moves are controlled using link parameters. Bulge angles are used to bulge areas like the biceps, legs and chest relative to the angle created by a link and its child in the hierarchy. Tendons can span multiple links in the hierarchy to stretch and pull a character's skin.

See also

Physique (page 2–1076)

Physique User Interface (page 2–1106)

Understanding Track Editing

There are several places to view your animation represented as tracks in 3ds Max. These views include:

- Track View - Curve Editor
- Track View - Dope Sheet
- Track Bar

The tracks can be displayed as function curves, or as keys and ranges on a box graph. **character studio** makes special use of function curve editing capabilities with a customized version of the Track View - Curve Editor called the *Animation Workbench (page 2–837)*.

Biped Curve Editing Using the Track Bar

After creating animation of the biped in the viewport using keyframing tools from the Motion panel, you might need to work on the animation tracks, either to adjust key interpolation for smooth motion or to adjust ranges and affect timing. 3ds Max allows for a number of basic ways to do these tasks for all scene objects using direct key manipulation.

The simplest and most direct access to keys and ranges can be found on the track bar. Keys for selected biped objects are immediately visible there. With some limitations, you can create keys by Shift+dragging in the track bar. Right-clicking any key will display a menu; choosing from this menu will give quick access to key properties such as interpolation controls. This Key Properties dialog is also accessible from a list on the Motion panel > Assign Controller rollout.

A right-click on the time slider also displays the Create Key dialog, which is another quick way to set keyframes.

Selecting any two keys displays the range as a bar beneath the keys. You can reposition this range bar or resize it to adjust the timing of the animation. You can combine accelerator keys, like Ctrl and Alt, with left, middle, and right mouse buttons to extend or stretch the range displayed in the track bar.

You can expand the track bar to quickly display the keys on function curves. From there, you can select and manipulate keys.

Biped Curve Editing in Track View

The same function-curve editing controls are found in Track View. Track View has two modes: Curve Editor and Dope Sheet.

character studio displays function curves for biped components in the Track View - Curve Editor.

Dope Sheet mode displays the footstep tracks and other biped keys as boxes. Here, you cut, copy, and paste keys to create repeated motion. You can use Time Editing in Dope Sheet mode, and cut, copy, paste, and insert time segments, complete with keys, into your animation.

Curve display can help pinpoint troublesome spots in your animation. You can see where a curve has problems, usually corresponding to motion problems in the viewport animation. You move the keys to compensate for the problems.

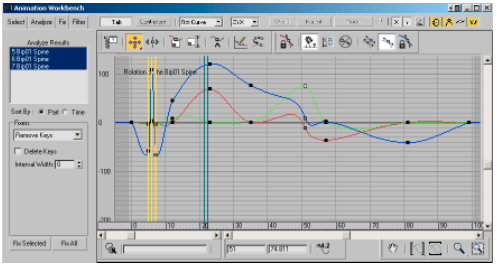
Understanding the Workbench

The Workbench is a curve editor customized for use with **character studio**. It provides specialized tools for selecting and displaying curves, and also for locating and fixing errors and discontinuities in motion. It is a visualization tool that allows you to see and manipulate quaternion function curves represented as euler angles. Also you can see curves for the the position of a biped body part in any coordinate space.

The Workbench automatically limits selections to just biped body parts and scene objects related to that biped. You can analyze these for errors using a variety of detectors, and fix them using the provided fixers. You can also apply filters to perform operations on different biped body parts to generally affect the animation without error identification.

Function Curve editing for Biped is also available in the 3ds Max Track View - Curve Editor and in the expanded track bar, but without the specialized Workbench tools.

Especially when working with imported motion capture data, you might find that curve editing is difficult using the standard 3ds Max tool set. This is because you might have a key on every frame, so curve manipulation becomes cumbersome and awkward. It is difficult to visually pinpoint where the trouble spots are located. The Animation Workbench offers automatic functions to reduce keys or apply filters to the motion curves to smooth animation.



The Animation Workbench

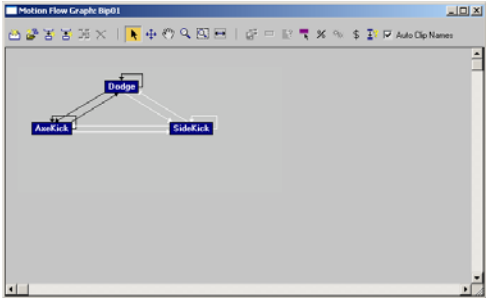
See also

Working with the Workbench (page 2–1008)

Animation Workbench (page 2–1012)

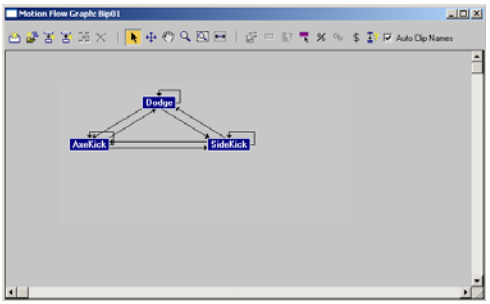
Understanding Motion Flow

Motion Flow is a tool that graphically arranges clips (motion files), flowing from one motion to the next. You can use a Motion Flow graph to set up a series of clips with transitions between them, which will make the biped perform the series of motions in sequence.



Motion Flow graph with transitions between motion files

You can also create a network of clips on the graph, where each clip has a transition to two or more clips. With this type of graph, you can tell **character studio** to generate the actual motion sequence for you based on random selection.



Multiple transitions from each clip

A crowd simulation can also be used to generate a motion sequence from this type of graph. Each biped in the crowd chooses from the motions and transitions in the graph based on where they want to go and how fast they need to get there. These factors, in turn, are determined by the crowd parameters you set.

See also

Working with Motion Flow (page 2–1026)

Motion Flow Mode (page 2–1043)

Motion Flow Rollout (page 2–1045)

Understanding Crowds



The Crowd system in **character studio** lets you create realistic simulations using large groups of characters, humanoid and otherwise, that behave and interact with one another by procedural means. You can use it to easily animate scenes containing hundreds of people and/or creatures, all with similar or widely varying sets of behaviors, which can vary dynamically according to other factors in the scene.

At the heart of the system are the Crowd and Delegate helper objects. A single Crowd object can control any number of delegates, which serve as stand-ins for crowd members. You can group delegates into teams, and assign behaviors such as Seek, Avoid, and Wander to individuals or teams. You can combine behaviors with weighting, so that, for example, a crowd member could seek a goal while wandering slightly.

Crowd simulations can range in sophistication from simple and straightforward to highly complex. Aiding at the latter end of the range is the Cognitive Controller feature, which lets you use the scripts to apply conditional transitions to sequences of behaviors. For example, you could tell a delegate to approach a goal until it gets within a certain distance, and then start moving away. Or you could use a Cognitive Controller to have a delegate move among a series of goals.

Another means of creating complex, dynamic crowd simulations is *motion synthesis*, which can be used in conjunction with the Cognitive Controller. The Crowd system offers two types of motion synthesis.

- When working with bipeds, you use the Motion Flow feature to allow the software to create scripts for the bipeds that match the delegate behaviors.
- When working with non-bipedal characters, such as fish or birds, you use Clip Controllers that let you apply different animation segments to various types of motion. For example, a bird might flap its wings quickly while ascending but slowly while flying level, and stop flapping its wings during descent.

One of the most important requirements of crowd simulations is avoidance; realism suffers if characters pass through each other or other objects in the scene. The Crowd system offers a number of behaviors to help achieve proper avoidance. It also provides the Vector Field, a special space warp, that, when applied to an irregularly shaped object, allows delegates to move around the object without penetrating it.

Used in combination, the Crowd-system tools described here can produce an endless variety of interesting, multi-character simulations. The Crowd topics in this manual provide detailed information about every aspect of crowd simulation.

See also

Crowd Animation (page 2–1154)

Crowd Animation User Interface (page 2–1182)

Understanding character studio Workflow

Biped, Physique, and the Crowd system work together within 3ds Max to provide a complete set of character animation tools. Although these components can be used in a variety of ways, it is helpful to approach **character studio** with a basic understanding of how a typical character animation is created.

The following sections provide a brief summary of the basic workflow and related benefits to creating a character with Biped and Physique. You may not use all the following steps, but you're likely to do them in the following order.

Create Skin Geometry

Create a basic skin shape for your character using any 3ds Max modeling tools and surface types. Be sure to place your character's skin in a neutral pose with arms outstretched and legs spaced slightly apart. You may also want to add sufficient detail to your skin's mesh or control points around joints to facilitate deformation during movement.

Note: Since Physique deformations are based on a volume, you can refine your geometry at a later time with minimal impact on your skin behavior. This means you can create your animation before you've built your model if you wish.

Create a Biped Skeleton

Biped automates the creation of bipedal character skeletons. It also lets you introduce significant changes to the skeleton's structure and sizing at any point during your animation without adversely affecting your character's motion. This means you can animate your character without knowing if it is short or tall, skinny or fat. It also means that if the director changes the character's proportions, the animation will still work.

For more detail on posing a biped skeleton, see *Biped (page 2–843)*.

Attach the Skin

- Position the biped character within its modeled skin. Use *Figure mode (page 2–982)* to scale bone lengths and to orient the skeleton correctly within the skin's volume. Scale bone thickness as desired to achieve a good initial fit. Then save a figure file, so it's easy to return to this pose.
- Use Physique to attach the skin geometry automatically to the biped or a 3ds Max bones hierarchy. The attachment is typically made to the root node in the hierarchy: the pelvis object on the biped or the root node on a bones hierarchy, not the center of mass. The attached skin is deformed as the biped or bones hierarchy moves.
- The links in the bones hierarchy are used to create a system of 3D envelopes that enclose nearby vertices. Envelopes typically overlap

at the parent and child ends of links. Vertices within overlapping envelopes are blended to create smooth skin deformation over joints as the character moves.

Adjust Skin Behavior

Adjust Physique parameters and introduce skin behavior effects to achieve the desired characterization.

- Change default envelope shapes by adding cross sections and control points to isolate a more specific volume of vertices for each bone. Use exclusion lists or per-vertex weighting to apply fine-tuning control over individual vertices.
- Introduce bulge angles to change muscle shape based on the angle of a particular joint. Create tendons to simulate the motion of tendons under the skin, based on link movements.
- Adjust link parameters to change skin twisting, sliding, and creasing as the biped moves. Sliding allows the skin to compress at the biceps and forearm as the elbow is bent. Twisting controls the amount of skin twisting across a joint intersection.
- Create extra links using 3ds Max bones and dummy objects for added control. Links can be added to the abdominal area to control compression, for example, or to create a link to animate the chest rising and falling as the character breathes. If a character has extra appendages, 3ds Max bones can be added to animate them. One common usage is to add a bone to animate the jaw.

For further detail, see *Physique (page 2–1076)*.

Animate the Biped Skeleton

Once the skin is attached to a Biped structure, you can freely animate the Biped character and see then skin behavior update automatically, based on the current pose.

Tip: Since Physique skin deformation can slow visual playback of your Biped animation, you can temporarily hide the skin object or reduce its resolution in the modifier stack to improve performance.

You can also choose to develop Biped animations in a separate scene entirely, and apply them to your final skinned character when you are satisfied with your final motion.

A *biped (page 2–843)* character is essentially an integrated hierarchy of bones that you can position freely using keyframes, IK goals, and footsteps. You can position a biped character using all the rotation and transformation tools found in 3ds Max.

Many of the 3ds Max coordinate systems can be used to position the biped. Local coordinates are useful to move a limb along its axis (the local X axis is always the axis along the biped limb); world coordinates are handy when there is any confusion regarding which way is up. You can use world coordinates as a home base. In 3ds Max, the world Z axis is always up.

Note: Rotations always occur about the local axis.

Use Freeform Techniques

Biped provides a variety of methods for creating character motion easily. You can use a purely traditional approach by manually creating keyframes in freeform mode for different poses, letting the computer interpolate between joint positions and IK goals.

Use a Footstep-Driven Technique

You can also choose a partially assisted approach by using footsteps and Biped dynamics to help you create a default walk, run, or jump cycle, and then adjust the biped keyframes and footsteps individually.

When using footsteps, *biped dynamics (page 3–916)* helps you by simulating gravity and balance.

- Gravity can help in a jumping motion to accelerate a character during the falling period and to bend the legs naturally on landing.
- Balance adjusts a character's position when the spine is rotated and keyframed to retain a character's balance.
- Dynamics can be turned off on a per-key basis or for the entire animation. The animator can override center of mass keyframes created using Dynamics calculation at any time. Simply set the dynamic properties of these keys or choose *Spline Dynamics (page 3–1015)* for keys generated by newly created footsteps.

Convert Between Animation Types

Once you are satisfied with a particular footstep animation and its corresponding dynamic behavior, you can *convert it (page 2–885)* to a freeform animation consisting of a simple combination of keyframes and IK goals. This intelligent conversion gives you control of animation behavior at every frame, for every joint of the character.

Use Layers to Apply Global Changes

Animation layers (page 3–910) offer you a powerful tool for introducing global changes to an existing character animation. For example, by adding a layer on top of an existing run motion, and creating a single keyframe with the biped's spine rotated forward, an upright running motion can be turned into a crouched run. Layered changes can be stacked up, allowing you to refine your motion composition and eventually collapse your layers into a standard non-layered keyframe animation.

Use In Place Mode to Control the View

In Place Mode (page 2–930) is a tool that lets you keep your biped in view during animation playback. It offers a convenient way of adjusting and adding keyframes to a character without constantly changing your view to follow the character's motion.

Import Motion-Capture Files

Motion-capture files can be imported from the BVH or CSM formats, edited, and saved as BIP files. You can import these files with or without footsteps and dynamics and combine them in Motion Flow mode with other animations.

You can use the supplied motion-capture samples as is or adjust them to suit your needs using Biped's collection of animation tools. The ability to import motion-capture marker files directly into **character studio** using the CSM file eliminates much of the cost of post-processing optical motion capture data. You can import motion-capture files with an additional prop bone, to define the motion of an object such as a sword or club.

You can also import *HTR/HTR2 motion-capture files (page 3–576)*, as well as *TRC files (page 3–577)*.

Motion-capture files can be imported with key reduction, making for more manageable tracks for subsequent editing.

Use Track View for Keyframe Editing

Track View (page 2–1002) allows you to edit keys and footsteps relative to the animation time line.

Footstep editing in Track View - Dope Sheet allows you to move footsteps in time. If you need a character to jump higher between footsteps, move the landing footsteps further down in time; dynamics automatically compensate by making the character jump higher to keep it airborne longer.

You can also specify a freeform period in a footstep animation, using Track View - Dope Sheet. This allows you to take advantage of footsteps and dynamics for part of an animation, then switch to manual keyframing during the freeform period. This approach can be particularly useful in animations where there is a mix of animation where the feet are on the ground and then off. Examples of this type of animation include running and diving, or walking and then sitting down.

Keyframe adjustment tools allow you to find the next or previous key for the selected biped body part, use the Time spinner to slide a key back and forth in time, change Tension Continuity and Bias for a key, and to display trajectories.

You can place arms and legs into the *coordinate space (page 3–924)* of another object or the world to simulate interaction with fixed or moving objects. In Freeform mode, for example, putting the character's legs into *world space (page 3–1036)* prevents them from sliding or moving when the animator is keyframing the character's center of mass. Putting a character's hand in the coordinate space of a ball allows the hand to move wherever the ball moves.

Many tools in 3ds Max can be leveraged with **character studio**. For example, the Select and Link tool can be used to attach objects to the biped.

If a character is to pick up and carry an object and then put it down, the Link controller can be used to animate the duration of the attachment. 3ds Max bones can be used to animate character subassemblies, like pistons, and to create extra links for Physique.

Use the Motion Mixer to Mix Animations

You can use the *Motion mixer (page 2–604)* to combine motions on a biped. For example, you could combine a walking motion with a cheering motion, and cause the biped to walk while cheering.

Use Motion Flow to Combine Animations

After you have created and modified various animation sequences, and stored them in biped motion files (BIP format), you can use *Motion Flow mode (page 2–1043)* to combine various motion files into longer animations that can be quickly previewed and edited. Motion Flow mode automatically places the animations end-to-end, allowing you to mix and match both freeform and footstep-driven motion files. Transitions between successive motions are automatically created for you, to provide a first-pass blending between overlapping frames of animation.

The Motion Flow transitions use *velocity interpolation (page 3–1030)* to create seamless transitions between clips. You can use the *Transition Editor (page 2–1051)* to modify a variety of blending parameters, including transition start frame, length, and angle between clips.

Refining Your Character

Great character animation is a result of many refinements that tune the overall personality of your character. You will find the need to refine progressively both the skinning behavior and the animation timing of your **character studio** character. Biped and Physique make this iteration process straightforward by fully using the 3ds Max modifier stack and undo methods.

In addition, Biped's ability to map motions between characters makes it easy to interchange great animations with existing characters, and tune their behaviors to achieve true integration of motion with character motivation and personality.

Use Crowds to Animate Groups of Characters

Once you've created animation sequences for characters or other models (such as a bird flapping its wings), you can replicate the models or characters and apply the motions to these groups using the *Crowd system (page 2–1154)*. You can also combine them with a wide range of supplied behaviors to create lifelike activities in crowds, such as people streaming through a doorway, street traffic, or birds and fish flocking and avoiding obstacles. You can use Motion Flow mode to create motion clip networks so that characters perform animation sequences appropriate to their current movement and transition smoothly between sequences. And you can use Crowd's cognitive controllers to transition between behaviors based on a variety of criteria. For more on crowd behaviors, see *Creating a Crowd System (page 2–1155)*.

Biped

To work most efficiently with bipeds, it is important to follow the general workflow described in this topic.

Create Skin Geometry

Before you create a skeleton for a character, you should already have a character skin to put the skeleton into.

Create a basic skin shape for your character using any of the 3ds Max modeling tools and surface types. Be sure to place your character's skin in a neutral pose with arms outstretched and legs spaced slightly apart. You may also want to add sufficient detail to your skin's mesh or control points around joints to facilitate deformation during movement.



Character mesh in a neutral pose

Tip: Before adding a biped skeleton, freeze your character mesh. When the mesh is frozen, you can still see it, but you can't select or alter it, reducing the chance for error or frustration.

Create a Biped Skeleton

Once you have a character mesh, you can create a biped skeleton to fit inside. Use *Figure mode (page 2–982)* to set up your biped.

Before you position the skeleton, use controls on the *Structure rollout (page 2–984)* to alter the biped to match your mesh, setting the number of links for the spine, arms, neck, or fingers, or adding props to represent weapons or tools.

Tip: You can use ponytails to create animated jaws, ears, or horns.

Note: Certain biped body parts, including fingers, tails, ponytails, props, and clavicles, can be repositioned in Figure mode to suit different characters.

When you position the biped inside your mesh, start with the center of mass (COM), which is the parent of all objects in the biped hierarchy. The COM should be positioned in line with the hips of the mesh character. Scale the pelvis so that the legs fit properly in the mesh, and then use Move and Scale on the 3ds Max toolbar to position your biped skeleton.

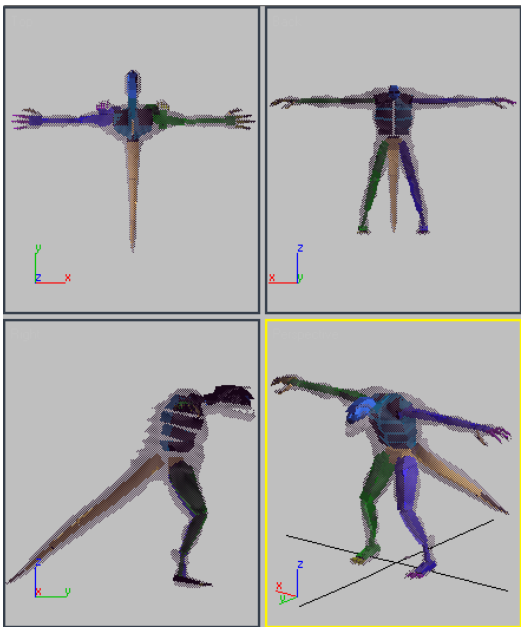
Note: In addition to the standard move, rotate, and scale operations, you can also use modifiers to adjust the parts of the biped.

Note: Biped body parts cannot be removed, however unwanted parts can be hidden. If you delete a part the entire biped will be deleted.

The following list includes some tips for positioning your skeleton:

- Use the **Page Up** and **Page Down** keys to cycle through links.
- Use *Rubber Band mode (page 3–1003)* to move and scale the arm and leg links on your skeleton simultaneously.
- Use tools from the *Bend Links rollout (page 2–952)*, such as *Bend Links Mode (page 3–915)* and *Twist Links Mode (page 3–1028)* to adjust tail, neck, spine, and ponytail links.
- The head, toes, and fingertips should extend slightly beyond the mesh extents to fulfill the requirements of Physique.
- Use the minimum number of fingers and toes.

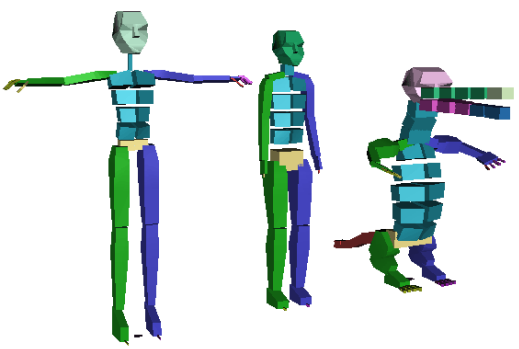
  You need extra fingers or links only if you are planning on complex hand or foot animation. If your character is wearing gloves or shoes, then you probably only need one finger or toe, with one link.
- To create a biped with knees that bend backwards, rotate the biped calves or thighs of both legs 180 degrees about their local X-axis (along the length of the limb). When you exit Figure mode, the biped walks, runs, and jumps with reversed knees.
- When working with a mesh in a symmetrical pose, pose one side of the skeleton, and use controls on the *Copy/Paste rollout (page 2–966)* to paste the posture to the opposite side of the biped.



When you are satisfied with your pose, check the alignment in all viewports to make sure that the skeleton is positioned correctly in the mesh.

Once you have successfully positioned a skeleton inside your character mesh, you are ready to attach the mesh with Physique. For more on this workflow, see *Understanding Physique (page 2–834)*.

Creating a Biped



A biped model is a two-legged figure: human or animal or imaginary. Each biped is an armature designed for animation, created as a linked hierarchy. The biped skeleton has special properties that make it instantly ready to animate. Like humans, bipeds are especially designed to walk upright, although you can use bipeds to create multi-legged creatures. The joints of the biped skeleton are limited to match those of the human body. The biped skeleton is also specially designed to animate with **character studio** footsteps, which help solve the common animation problem of locking the feet to the ground.

The parent object of the biped hierarchy is the biped's center of mass object, which is named *Bip01* by default.

Creating a Biped

A button for creating a Biped appears in the Create panel under Systems. You create a biped by clicking this button and then dragging in the active viewport. You interactively define the height of the biped as you move the cursor.

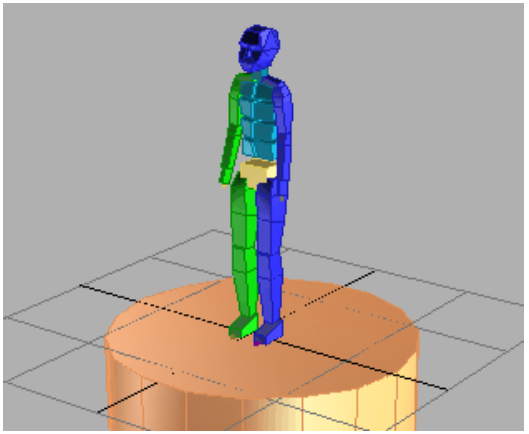
During biped creation, you can change any of the default settings that are used to define its basic structure. The default settings, for Arms, Neck Links, Spine Links, and so on, are for a human figure.

If you turn on Most Recent .fig File, the biped you make will use the parameters stored in the last FIG (figure) file you've loaded.

Changing Biped Parameters

Like other 3ds Max objects, you can change biped parameters on the Create panel at creation time. However, to modify or animate a biped, you use parameters on the Motion panel. For more information on changing biped parameters, see *Structure Rollout (page 2–984)*.

Creating a Biped Using AutoGrid



The AutoGrid feature lets you create objects on the surface of other objects, rather than on the home grid. AutoGrid automatically creates a grid for construction. When you turn AutoGrid on, a tripod cursor is displayed in the viewport. As you move your mouse over the geometry in the scene, the tripod cursor rotates to match the orientation of the face it's touching. A grid is created at that spot and is used to create the biped.

By turning on AutoGrid in the Create Panel > Object Type rollout, biped creation will respect the AutoGrid, letting you build bipeds with their feet on top of geometry.

Procedure

To create a biped on a surface:

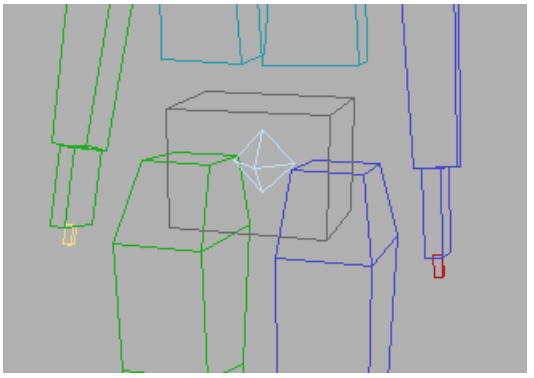
1.  On the Create panel > Systems, click Biped.
2. Turn on AutoGrid.
3. Move your cursor over any geometry in the viewport.

   A Transform gizmo moves with your cursor to indicate the location of the AutoGrid.
4. Drag out a biped.

The biped feet will be in contact with the geometry.

Understanding Biped Anatomy

The geometry of a biped is a linked hierarchy of objects that by default resemble those of a human. The parent or root object of the biped is its *center of mass (COM) (page 3–920)*. This object is displayed as a blue octahedron near the center of the biped's pelvis. Moving the COM positions the entire biped.



Center of mass object

You can select the center of mass by choosing *Bip01* from the *Select Objects dialog (page 1–78)*.

 You can also select the center of mass by clicking Body Horizontal, Body Vertical or Body Rotation in the *Track Selection rollout (page 2–945)*.

Changing the Biped Hierarchy

The biped Hierarchy is a little different than a standard 3ds Max hierarchy in that you can't delete any of the components of the skeleton. If you try to delete any part of the biped skeleton, you delete the entire hierarchy. If you want to create a partial biped, for example a biped with no head, simply hide the objects you don't want to use.

Repositioning Biped Body Parts

Certain Biped body parts can be repositioned in Figure mode to suit different characters. You can move entire arm assemblies by selecting the clavicles and moving them up or down. You can also reposition the fingers, tail and ponytails as you like. See *Posing the Biped (page 2–847)*.

Biped's structure also includes an option to add forearm twist. This uses two to four forearm links to transfer twisting animation into the biped's associated mesh via Physique or Skin.

Props

The biped structure includes an option to add up to three props. Props appear next to the biped's hands and body by default, but can be modified or animated throughout the scene like any 3ds Max object.

Adding Extra Biped Body Parts

To add extra legs, arms, or other body parts you need to create 3ds Max geometry for those parts, then link them to the biped hierarchy. You can use Snapshot to duplicate biped body parts to create these as well. In either case you will need to animate them with standard 3ds Max rotations, because biped IK will not be available on these extra parts.

Changing Initial Biped Anatomy

Use the body parameters to change initial biped anatomy. The Body parameters are in the Create Biped rollout that appears in the Create panel when you create a biped.

Note: You can change body parameters in the Create panel immediately only after creating a biped. Once you leave the Create panel, these settings are still available, but from the Structure rollout in the Motion panel. The body parameters

in the Structure rollout are enabled when the biped is selected and Figure mode is active.

See also

Creating a Biped (page 2–844)

Structure Rollout (page 2–984)

Naming the Biped

If your scene is going to contain more than one biped, it's a good idea to give the biped a unique name. By default, the first biped in a scene is called Bip01. Succeeding bipeds have the same name except that the two-digit number is replaced by another number in sequence: 02, 03, 04, and so on.

Since the parent of the biped hierarchy is the biped's center of mass object, that is also the object selected when you choose Bip01. You can link the biped center of mass to other objects or dummies for additional animation control. For more information on the biped's hierarchy, see *Understanding Biped Anatomy (page 2–846)*.

Procedure

To change the biped name upon creation:

- Immediately after creating the biped and before moving to another panel, enter the new name in the Root Name field in the Create Biped rollout.

  The center of mass is renamed and the entire biped hierarchy inherits the new name.

To change the biped name after creation:

1. Select any part of the biped and go to the Motion panel.
2. In the Biped rollout, click the expander bar to display the Modes and Display groups.
3. Enter a new name in the Name field and press Enter.

Warning: Do not use the usual Name And Color rollout to rename a biped. This changes only the name of the biped's root object (its center of mass) without updating the names of other biped parts in the hierarchy.

To create a Named Selection for the biped center of mass:

1. Select the biped center of mass (diamond shaped object in the pelvis) in the viewports.
2. Type a name in the Named Selection Sets field on the 3ds Max toolbar.
3. Press Enter.

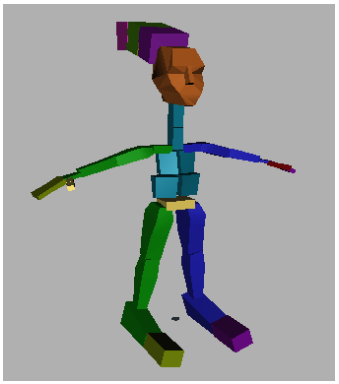
   You can now choose this selection if the biped is hidden and you want to view the biped and edit biped parameters on the Motion panel.

   Tip: Give the center of mass selection the same name as your character in the Named Selection Sets field, "John," for example. With multiple characters, you can quickly select the character you want to edit by choosing the character's name in the Named Selection Sets dropdown.

Posing the Biped

After creating the default biped, you will often need to change the proportions of the skeleton to fit your model. Use Figure mode to change the biped structure in its rest pose.

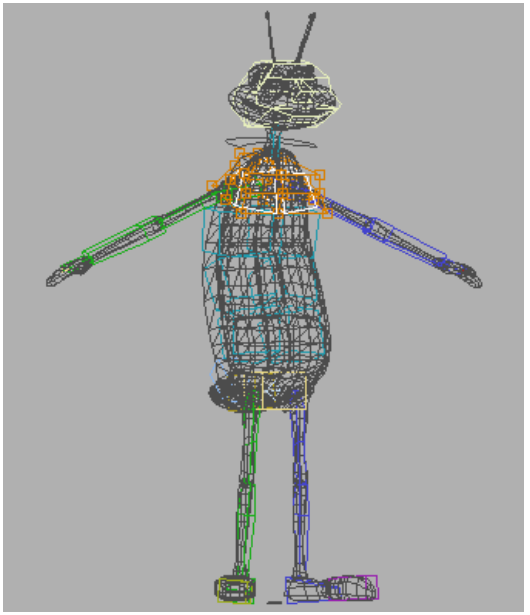
Activating Figure mode returns the biped to its original location and orientation. With the biped in Figure mode, you can use the transform tools to change the proportions and positions of body parts. For instance, you might apply a non-uniform scale to shorten the legs or lengthen the arms.





Biped proportioned to fit inside of Dr. X character geometry

You might rotate the spine objects to create the figure for a hunchback or a dinosaur. Use the move tool to change the position of the thumb or the arms. You can even apply modifiers to the biped skeleton pieces, such as using an FFD on the biped head to adjust its shape.



FFD Modifier used to shape spine and head

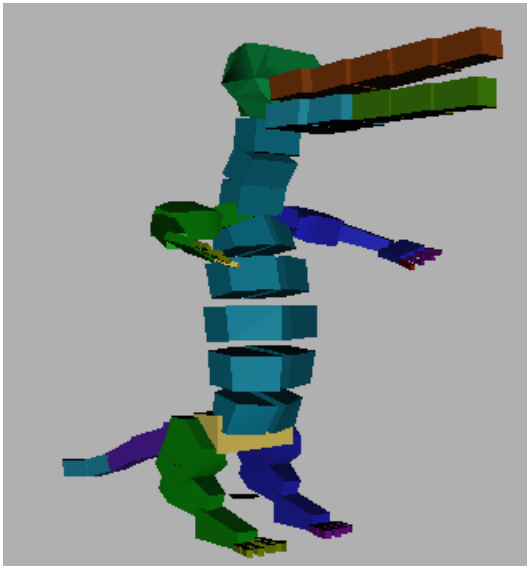
Once you have a default biped, you'll need to match its proportions to the character's geometry. It is quite typical to find the character's geometry with the arms outstretched. The common workflow is to freeze the character and then in Figure mode, reposition the biped, so the center of mass is at the base of the torso. The spine objects, legs, and feet are scaled and rotated to fit within the confines of the mesh, then the arms and hands, neck and head. The tail and ponytail objects can be used for animating wings, fins, jaws, ears, horns, or hair.

Once your biped proportions are correct, you can save them in a .FIG file. Since biped saves the character in the .FIG file, and then animation in the .BIP file, you can change the character's proportions without affecting the animation.

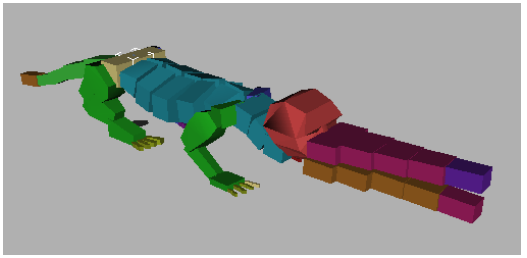
Changing the Biped Structure

Bipeds don't have to appear human. You can change their elements and form to create other kinds of characters. Although you can change some initial aspects of the biped's structure in the

Create panel, you use Figure mode to change all aspects of the biped's structure after its creation.



Bipeds don't have to be human ...



or even walk on two legs.

Any number of bipeds can be in Figure mode at the same time, though you can only work on one biped at a time. When you select a biped, the Figure mode button indicates whether or not the biped is in Figure mode.

In Figure mode you can:

- Specify the number of links in each part of the biped.
- Define the position of the base of the fingers, toes, clavicles, spine, tail, ponytails, and props relative to the body.
- Define the position of the feet relative to the ankles.
- Define the default pose of the biped before animation is applied to it, for example define a hunchback.
- Scale the biped and its various parts.
- Simultaneously scale and position biped parts using *Rubber Band mode (page 3–1003)*. See *Rubber-Banding Arms and Legs (page 2–852)*.
- Create natural links for Physique using Triangle Pelvis.
- Use forearm links to transfer twisting animation to the biped's associated mesh.

For more information on adjusting the biped's proportions see *Posing the Biped (page 2–847)*.

Working with Both Arms or Legs

As you pose your biped, there are two different methods for posing both arms or legs symetrically.

The first method involves using Symmetrical on the *Track Selection rollout (page 2–945)* to select both limbs simultaneously. With both limbs selected, you can rotate and scale the objects, and they will maintain a mirror-image relationship. You can also move the limbs, however they will not retain a symmetrical pose (with respect to the body) if you move them laterally.

The second method involves posing only one side of the biped, then copying the posture to the corresponding limbs on the other side.

Note: You can scale biped limbs only in *Figure mode (page 2–982)*.

Procedures

To work in Figure mode:

1. Select the biped you want to pose, and then go to the Motion panel.

The Motion panel doesn't show Biped controls unless the biped is selected.

2. On the Biped rollout, click Figure mode.

 The button turns blue to indicate a special edit mode.

The biped's pose and location change to the one last specified when the biped was in Figure mode. If the biped was just created and, therefore, was never in Figure mode before, it changes to the pose and location it had when you created it. When you turn off Figure mode, the biped returns to its animated pose and location in the scene.

Note: Biped disables all keyframe editing tools when Figure mode is active. Adjustments made in Figure mode are not animatable.

To fit biped legs to the skin:

1.  Turn on Figure mode on the Biped rollout.
2. Use the Scale Transform gizmo to scale the Z-axis of the biped's pelvis so the biped's leg links are centered in each leg of the skin.
3. Select the biped's two thigh links (LLeg and RLeg) and scale the X-axis of the thigh links so they end at the knees of the skin.
4. Select the biped's two lower leg links (LLeg1 and RLeg1), and scale the X-axis of the lower leg links so the biped's ankles are level with the ankles of the skin.

   Tip: When the two upper legs are selected, you can press Page Down to select the lower legs.
5. In a left or right viewport, scale the biped's feet so their profile roughly matches the profile of the feet of the skin.
6. Scale toes or move them along their local X-axis so each toe is aligned with the corresponding toe in the skin.

   The ends of the final toe links should go through the tips of the skin's toes.

   You might have to change the number of biped toes to match the number of skin toes. A biped must have at least one toe on each foot. If the skin has no toes or the character is wearing footgear, the position and number of biped toes doesn't matter—but they should still extend beyond the skin.

   Tip: During the fitting process, try Freezing the skin object to prevent accidental selection.

To fit the spine to the skin torso:

1.  On the Biped rollout, turn on Figure mode.
2. Select the lowest link of the spine (Bip01 Spine). Choose the Move transform, constrain movement to the spine link's local X-axis, and move it vertically to the waist of the skin, just below the navel.
3. Scale the other spine links in their local X-axis so they fit the upper part of the skin's torso.

   The neck link should begin where the skin's neck begins.

   If the torso of the skin curves, you should also rotate spine links about their local Z-axes, to align the spine with the longitudinal center of the torso.
4. Scale the biped's neck in its local X-axis to match the length of the neck of the skin. The top of the last neck link (also the base of the head link) should be where you want the head to pivot. This is usually just below the ears, centered with the spine.

   Leave the head in its default position relative to the spine and neck links.

To fit both arms using copy/paste:

1.  On the Biped rollout turn on Figure mode.
2. Rotate one upper arm (R Arm1 or L Arm1) in its local Y-axis to center its link in the upper arm of the skin.
3. Scale the upper arm so its link ends at the elbow of the skin.
4. Scale the lower arm (RArm2 or L Arm2) so its link ends at the wrist of the skin.

   If the skin's arms are bent, rotate the lower arm to center its link as well.
5. Scale fingers, or move them along their local X-axis so each finger is aligned with the corresponding finger of the skin.

   The ends of the final finger links should go through the tips of the skin's fingers. You might have to change the number of biped fingers to match the number of skin fingers.
6. When the arm is completely fitted to the skin, select all of it and on the Copy/Paste rollout, click Copy Posture.
7. On the Copy/Paste rollout, click Paste Posture Opposite to pose the opposite arm.

To pose both arms simultaneously:

1.  On the Motion Panel > Biped rollout, turn on Figure mode.
2. Select the biped's left or right hand.
3.  On the Track Selection rollout, click Symmetrical.

   The opposite hand is also selected.
4. Move, rotate, and scale the hands until you have the position and size you want. Use PAGE UP and PAGE DOWN to move the different parts of the arm.

   Note: If you move limbs laterally, they will both move in the same direction, and will no longer be symmetrical about the body.

To create a symmetrical pose by copying one side of the biped to the other:

1.  On the Motion Panel > Biped rollout, turn on Figure mode.
2. Move, rotate, and scale the left arm and leg of the biped until you have the position and size you need.
3. Select all of the bones in the left arm and leg.
4.  On the Copy/Paste Rollout, create a collection, then turn on Posture and click Copy Posture.
5.  On the Copy/Paste Rollout, click Paste Posture Opposite.

   The right arm and leg assume the position and scale of their corresponding bones on the left.

See also

Figure Mode (page 2–982)

Biped Rollout (page 2–936)

Creating a Skin (page 2–1076)

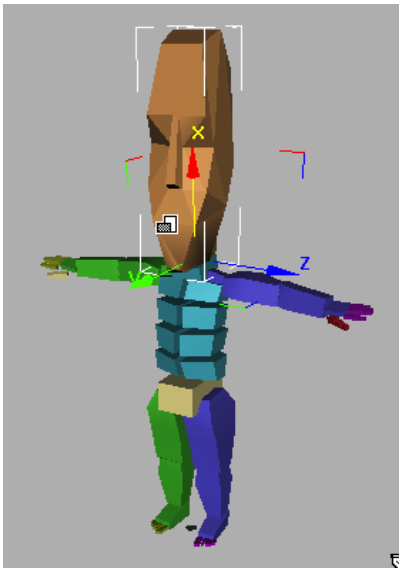
Scaling Links

Use standard 3ds Max scale transforms to adjust a biped's posture by scaling the size of its links. You must be in Figure mode to scale the biped links. If you try to scale a biped without going into Figure mode, nothing happens.

As with rotation, when you scale biped links, Biped constrains the transform to use the link's Local coordinate system. The position of other biped links can change so they remain attached to the resized link. If you shorten the thigh, the

calf and ankle will maintain their size, but change their position.

To scale a link, select any scale icon from the Scale flyout on the Main Toolbar. When you select a body part to scale, use the Transform gizmo to scale along one axis at a time.



Use Scale and the Transform gizmo to scale links.

If your character is symmetrical, select body parts in pairs and scale them at the same time.


Select one body part and then click Symmetrical in the Track Selection rollout. Now both are selected.


You can also scale one and then use Copy Posture and Paste Posture Opposite from the Copy/Paste rollout. This approach ensures symmetry in your character.

Tip: Use the **Page Up** and **Page Down** keys to move through your hierarchy as you work. For example, after you scale the thighs, press **Page Down** to select the calves.

Scaling the biped limbs to fit snugly to the mesh will help when Physique is used to associate the mesh with the biped. The Bounding Box option uses biped limb dimensions to size the envelopes. This saves time when you adjust envelopes in Physique.

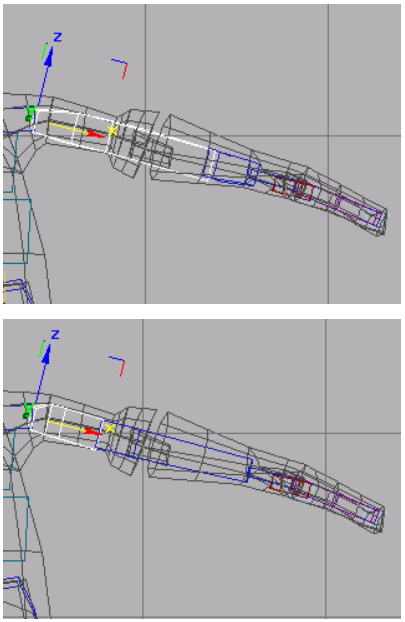
See also

Rubber-Banding Arms and Legs (page 2–852)

Rubber-Banding Arms and Legs

Rubber Band mode provides a way to proportion the arm and leg links simultaneously. Rubber Band mode works with the Move transform rather than the Scale transform. When you move an arm or leg with Rubber Band mode turned on, both the link and its child are scaled in a single step.

Rubber-banding the upper arm rescales the upper and lower arm objects and moves the elbow link without affecting the position of the shoulder or the wrist. If you've spent a lot of time getting the fingers in the right place, you can reposition the elbow by rubber-banding, without affecting the hands.



Using Rubber Band to resize an arm without changing the hand position

Procedures

To rubber band an arm or leg link:

1. On the Motion panel > Biped rollout, turn on Figure mode.
2. Select the arm or leg link you want to rubber band.
3. Turn on the Move transform.
4. On the Biped rollout > Modes group, turn on Rubber Band mode.

   Note: This button is unavailable if you are not in Figure mode, or if a part of the biped other than the upper or lower arm or leg, or the center of mass is selected.

   Note: The Modes group is hidden by default. To display it, click the expansion bar on the Biped rollout.
5. Move the selected arm or leg link.

   As you move the arm or leg link, the hands and feet are stationary as the knees or elbows are positioned.

To adjust the biped center of mass with Rubber Band:

1. On the Motion panel > Biped rollout, turn on Figure mode.
2. Select the center of mass object (diamond shape) on the biped.
3. Turn on the Move transform.
4. On the Biped rollout > Modes group, turn on Rubber Band mode.

   This button is unavailable if you are not in Figure mode or if a part of the biped other than the upper or lower arm or leg or if the center of mass is selected.
5. Move the center of mass.

   While Rubber Band mode is active, the center of mass moves independently from the rest of the biped, which remains in the same position.

   Tip: Another way to adjust the biped's balance is to change the value of Balance Factor. Balance Factor displays on the Key Info rollout when a Body Horizontal key is current. Balance Factor can be keyframed.

Biped Display Options

For greater speed in displaying bipeds, or to make your viewports less cluttered while you edit your scene, Biped lets you turn off the display of some biped elements. These display controls are found in the Motion panel > *Biped rollout (page 2–936)* > Display group, rather than in the Display panel.



These controls allow you to quickly turn on and off the biped bones, objects and footsteps, twist links and leg states, as well as footstep numbers and trajectories. There is also a Display Preferences dialog accessed from here that lets you control which bipeds are visible during Biped Playback.

Note: The Display group is hidden by default. To display the group, click the expansion bar on the Biped rollout.

Procedure

To change biped display:

1. Go to the Motion panel.
2. On the Biped rollout > Display group, click a button corresponding to the display choices you want.

   Note: The Display group is hidden by default. To display the group, click the expansion bar on the Biped rollout.

Deleting a Biped

You can quickly and easily delete an entire biped from a scene.

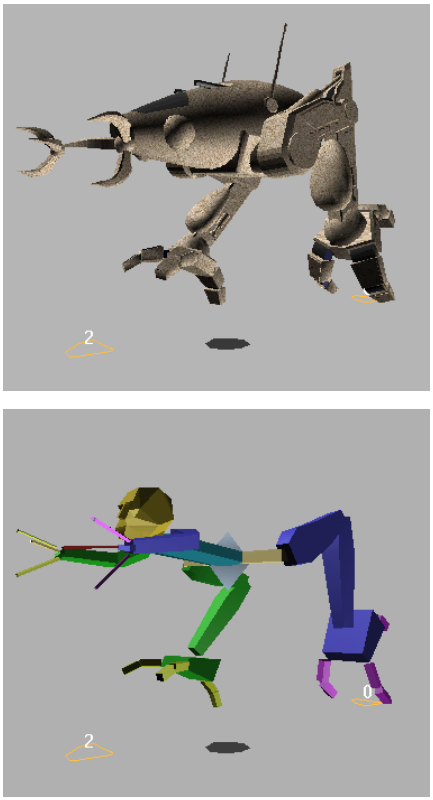
Procedure

To delete a biped:

1. Use any selection tool to select the entire biped or any part of the biped.
2. Press Del or Delete on your keyboard.

Linking Character Body Parts to the Biped

With the Physique component, you can use Biped to animate a deformable skin, usually a mesh object. However, some animations don't require deformation. For example, a knight clad entirely in rigid metal armor doesn't need to deform as skin does. Also, figures seen from a distance don't require the same degree of realism as figures seen close up.



Jointed character linked with biped skeleton

Characters available commercially often come in two varieties: jointless and jointed. Jointless characters have a seamless mesh at limb joints. Jointless characters should be attached to the biped using Physique. A jointed character has separate objects with ball joint geometry for the limbs, and lends itself to the linking technique described in this topic.

Linking objects and other geometry to the biped can also be used in cases such as the following:

- Link a weapon or a flower to the biped hand.
- Link eyeballs or teeth to the biped head.

- Link extra 3ds Max bones or splines to the biped to create extra envelopes when Physique is applied.
- Link 3ds Max bones to the biped to automate mechanical assemblies when the biped is keyframed.
- Link particle emitters to the biped hands or feet to create smoke or dust.

  Note: If you've linked particle emitters or 3ds Max bones (with the IK controller) to the biped, the Animate button must be on when you reposition the biped.

Procedure

To link a mechanical character to the biped (without Physique):

With mechanical or jointed characters, you can simply link objects to the biped without using the Physique modifier.

1. Load or create a mechanical or jointed model. Body parts in this model should be separate 3ds Max objects.
2. Create a biped.
3.  On the Biped rollout turn on Figure mode, then position and fit the biped to your mesh objects.
4. Go through the mesh objects and use Select and Link on the 3ds Max toolbar to link each object to its corresponding part on the biped.

   All of keyframe animation applied to the biped will animate the model.

   Tip: If you plan to use 3ds Max linking and Physique together, link the body parts to the biped after applying Physique. In this way, Physique will not create superfluous links from the biped to the mechanical body parts when it is applied.

See also

Biped Rollout (page 2–936)

Saving and Loading FIG Files

Select a biped. > Motion panel > Biped rollout > Load File (in Figure Mode)

Select a biped. > Motion panel > Biped rollout > Save File (in Figure Mode)

Figure files save all information about a biped's anatomy: links, link positions, twist links, and Figure mode posture, and the scale of geometric elements. Figure (FIG) files have a *.fig* file name extension.

Figure files do not contain the actual biped skeleton, or any animation. They are used strictly for setting or saving a biped's pose in Figure mode.

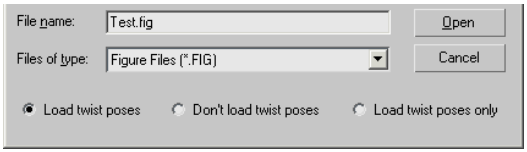
Figure files are useful for quickly setting poses for similar or identical characters.

 Clicking Load File brings the Open dialog, from which you can pick a FIG file and choose from the following *twist poses (page 2–950)* options:

- **Load Twist Poses**: Loads both the *twist links (page 2–984)* and twist poses data. (Default=on).
- **Don't Load Twist Poses**: Loads only the twist links data.

  Note: The pose presets are always available.
- **Load Twist Poses Only**: Loads only the twist poses data without the twist links.

  Note: The Twist and Bias values are not used without twist links. Refer to the *Twist Poses rollout (page 2–950)* for details.



Procedures

To save a biped's figure information to a file:

1. Select the biped to save.
2.  On the Motion panel, activate Figure mode.
3.  On the Motion panel > Biped rollout, click Save File.
4. In the file dialog, enter a name for the figure file, and then click OK.

   Tip: While you work on creating a biped pose, save your work frequently in a figure file.

To load a biped figure:

1. Select the biped to replace with a saved figure.
2.  Activate Figure mode in the Motion panel.
3.  On the Motion panel > Biped rollout, click Load File.
4. In the file dialog, choose the figure file to load.
5. Choose one of the three loading options, depending on whether or not you want to load twist poses, or to load twist poses only.
6. Click OK.

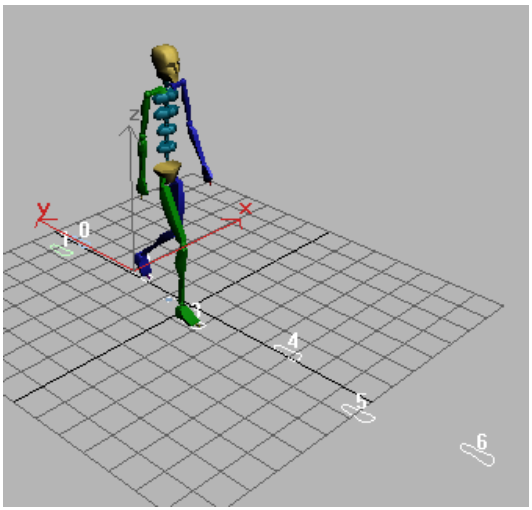
Warning: Loading a figure file replaces the selected biped's pose and base parameters. If you have created a new pose or a new biped structure, save it in a figure file before you load a different biped figure.

See also

Figure Mode (page 2–982)

Footstep Animation

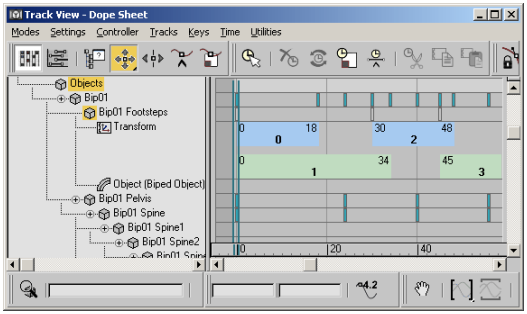
Footstep animation is a central compositional tool in Biped. Footsteps are biped sub-objects, similar to *gizmos (page 3–949)* in 3ds Max. In viewports, footsteps look like the diagrams often used to illustrate ballroom dancing. Each footstep's position and orientation in the scene controls where the biped steps.



Footsteps

The language of footsteps allows you to more directly describe and compose the complex time and space relationships that are found in different forms of locomotion.

While the placement of footsteps appears in viewports, the timing appears in Track View - Dope Sheet Editor. There, each footstep appears as a block of time, with each block representing the time when the foot is planted in a footstep.



Footstep keys in Track View

Footsteps are appropriate for animation where the biped is on the ground or uses the ground a great deal, such as walking, standing, jumping, running, dancing, and athletic motion. For movement that does not require the biped to interact with the ground, such as swimming or flying, *freeform animation (page 2–886)* is more appropriate.

Footsteps are used to lock the foot to the ground. You can create the same result in freeform animation by simply creating *planted keys* for the feet.

Workflow

In general, you compose a footstep pattern with these actions:

1.  Select the biped and access the Motion panel.
2.  Click Biped rollout > Footstep Mode.
3.  In the Footstep Creation rollout, click Walk, Run or Jump to choose the gait (a method or pattern of moving on foot). The gait sets the initial timing for footsteps you create.

   When you first create them, footsteps are inactive. They exist in the scene but don't yet control the biped's motion.

4. Create the footstep pattern in viewports. You can create footsteps in a few different ways, each available in the Footstep Creation rollout:

    Automatically create a number of footsteps with Create Multiple Footsteps, or

    Individually place each footstep with Create Footsteps (append) or Create Footsteps (at current frame). Use Auto Grid to create footsteps on a surface of a mesh.

5. Edit footstep placement by moving or rotating footstep icons in viewports.
6. Edit footstep timing in the Track View - Dope Sheet mode.
7. Activate the footsteps to create keys for the biped. The biped will now move through the footsteps using the placement and timing you set up.
8. Play the motion to test it out, looking only at the leg and foot motion and disregarding the upper body for now. If the footstep motion isn't right, deactivate the footsteps, make changes to footstep placement or timing, reactivate the footsteps and play it again. Do this until the foot and leg motions are correct.

   When footsteps are activated, keys are created for each of the tracks of the figure: the head, spine, pelvis, arms, legs, and, if appropriate, tail and ponytails. These keys form an initial sketch of your animation. The default keys, when *interpolated (page 3–958)*, form the basic, minimal motion required to animate the figure according to the footstep pattern. Most likely you will want to adjust or replace these keys.

9. Edit the animation of the upper body using ordinary 3ds Max animation methods: Turn on Auto Key and move or rotate the biped parts. You can also use the **character studio** Set Key tools on the *Key Info rollout (page 2–954)*.

10. Play the animation again and make any corrections to the upper body motion.

Footstep Method

In the viewports, footsteps represent support periods in space for the biped feet. You can move and rotate footsteps in viewports. In Track View, each footstep appears as a block that represents a support period in time for each of the biped's feet. You can move footsteps in time in Track View. The footstep position and orientation in the viewport controls where the biped will step.

There are three ways to create footsteps for the biped:

- Place footsteps individually.
- Use the footstep tools to automatically create a walking, running or jumping animation.
- Import motion-capture data to footsteps.

A key advantage of the footstep method is the natural adaptation of the biped that occurs when the footsteps are edited in time and space. Editing footsteps in the viewports allows you to reposition all of the footsteps to move the entire animation. In Footstep mode, stride, length, width, and direction can be changed quickly for an entire animation and the biped automatically adapts.

Using the Footsteps Show/Hide button on the Display rollout, all footsteps can become visible. Move the footsteps in the viewports to position them for proper ground collision with the terrain object. For example, if the biped toes are rotated for the Lift key at the last frame of a footstep (to create more toe curl as the character walks) the leg automatically repositions itself to maintain foot contact with the ground (footstep).

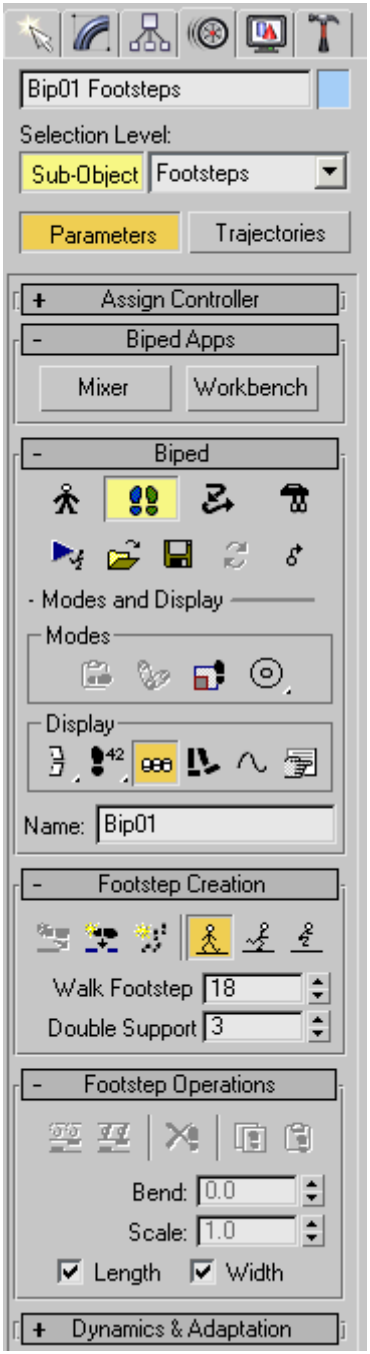
These adaptations speed up the process of creating and editing animation for the biped. If necessary, the animator can prevent biped adaptation by using the *Adapt Locks parameters (page 3–905)* on the *Dynamics & Adaptation rollout (page 2–980)*.

Foot States

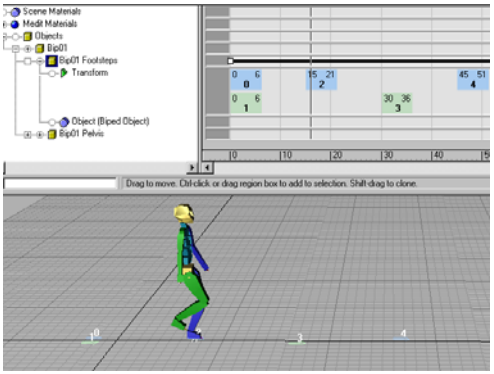
Within a footstep animation, there can be four *foot states (page 3–943)*: *move, touch, plant*, and *lift*. These correspond to the state of the biped feet in relationship to the footsteps. Use the foot state displays in the Biped rollout to determine the state of the biped feet when you are editing the biped foot or leg keys. The foot states can be displayed in the viewport by turning on Leg states in the Display group of the expanded Biped rollout. The Foot states are represented as colored keys in Track View - Dope Sheet.

Footstep Mode

 When Footstep mode is active, footstep creation controls appear in the Motion panel.



Footstep controls in Motion panel



Footsteps displayed in Track View - Dope sheet editor

Using the default keyframes as a starting point, you can interactively insert, replace, or delete keyframes in order to refine the motion of the biped and fill in the details of movement that are unique to your animation.

Note: By default, when *Biped Dynamics (page 3–916)* is turned on in the Dynamics and Adaptation rollout, gravity (Dynamics Blend) and ballistic tension calculate the trajectory of the center of mass for all newly created keys in a footstep animation containing a running or jumping motion. If *Spline Dynamics (page 3–1015)* is turned on in the Dynamics and Adaptation rollout before footsteps are created and activated, the center of mass uses spline dynamics to calculate vertical motion, which does not take gravity into consideration. Using spline dynamics, you must set keys for the top of a jumping motion or the dip when the character lands; this trajectory is automatically calculated with biped dynamics.

Footstep Editing

At any point in the design process, you can choose to interactively edit your footstep's spatial pattern in the scene or the timing of footsteps in Track View. The keyframes adapt to each edit: changes to footstep location retain the details of all your keyframe positions. Keyframe timing remains synchronized with changes to footstep timing,

except in cases where default leg keys must be regenerated to account for timing edits that alter the basic gait pattern, such as creating a hop in the middle of a walk.

Use Footstep mode to create and edit footsteps. Use Keyframe mode (Footstep mode off) to create and edit your character's keys. You can always edit the timing of both footsteps and keyframes in Track View.

While the biped's feet are airborne, you can animate its legs as you do its upper body. Biped does not create keys based on physics while the biped is off the ground, so animating the legs might be necessary to make long leaps realistic. Alternatively, you might want to make the biped appear to be floating in midair, or underwater, or have it ride a bicycle. See *Freeform Editing Between Footsteps (page 2–883)*.

You can make the biped interact with other objects in the scene: throwing or kicking a ball, opening a door, and so on. You do this by attaching a biped limb to an object in the scene.

An animatable IK Blend parameter lets you store the anchored position and combine inverse with forward kinematics. After you've set keys, be sure to remove the anchors.

Tip: You can often get good results by loading an existing biped motion and then varying it. You'll find a set of sample motion files on the program disc in the *\Samples\Motions* directory.

Note: On the time slider or in Track View, you can move one Biped key past another. See *Moving Keys (page 2–1004)*.

See also

Creating Footsteps (page 2–860)

Editing Footstep Motion (page 2–868)

Creating Footsteps

The topics in this section deal with the creation of footsteps for animating bipeds. They are:

Planning for Footsteps (page 2–860)

Choosing a Gait (page 2–861)

Creating Footsteps Automatically (page 2–862)

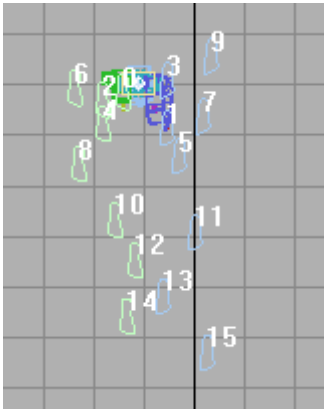
Creating Footsteps Manually (page 2–863)

Activating Footsteps (page 2–865)

Deactivating Footsteps (page 2–865)

Understanding Footstep and Body Keys (page 2–867)

Planning for Footsteps



Spend a little time planning your footstep animation with pencil and paper first. This will make the task of setting up footsteps much easier. Answering the questions below can help orient your planning.

Planning for Footstep Placement

- How are the footsteps going to be positioned in viewports? Draw a sketch of their positions.

- What is the sequence of the footsteps? This will determine the footstep numbers.

  Footsteps always begin with the number 0. Number them in your diagram if you can.

- Is each footstep a left or right footstep?

Planning for Footstep Timing

- What is the support relationship between left and right footsteps? In other words, does the timing of the footsteps resemble a walk, run or jump pattern? A character doing a foxtrot or waltz, for example, is actually moving with "walk" timing – the character shifts weight from one foot to the other.
- How long, in frames, should each footstep be? How long are the biped's airborne periods?
- Are there any periods of time where the biped is standing on two feet for a while?

Choosing a Gait

A *gait* is a method or pattern of moving on foot.

 When you create new footsteps, the timing is determined by the gait you have chosen (Walk, Run, or Jump) and the parameters for that gait. You must choose the gait before parameters can be set.

Choose the gait that most closely approximates the type of movement you want to create. You can alter or adjust the gait after footsteps have been created.

Walk Gait

In a walk, at least one foot is always in contact with the ground. A period where one or both feet are in contact with the ground is called a *support period*. If both feet are on the ground, this is known as a *double support period*.

 Turn on Walk before creating footsteps to create this type of gait.

The gait parameters of a walk are:

Walk Footstep—The number of frames that each footstep remains on the ground.

Double Support—The number of frames in a double support period, that is, when both feet are on the ground.

Run Gait

With running, both feet are never on the ground at the same time. Either one foot is on the ground during the support period, or the body is airborne. While it is airborne, the body moves forward horizontally at a constant speed. In general, the longer the body is in the air, the higher it must fly after lifting off from the supporting foot to stay airborne for the specified period of time.

 Turn on Run before creating footsteps to create this type of gait.

The gait parameters of a run are:

Run Footstep—The number of frames that each footstep remains on the ground.

Airborne—The number of frames that the biped is airborne, that is, when neither foot is on the ground.

Jump Gait

Jumping is a special case of running. Both feet are in contact with the ground at the same time, or airborne at the same time. As with running, forward motion is horizontal and constant, but vertical motion depends on the length of the jump.

 Turn on Jump before creating footsteps to create this type of gait.

The gait parameters of a jump are:

Feet Down—The number of frames in which both feet are on the ground.

Airborne—The number of frames that the biped is airborne, that is, when neither foot is on the ground.

Changing these values changes the biped's jump behavior.

Setting Gait Parameters

Gait parameters can be found in two areas:

-  If creating footsteps automatically with Create Multiple Footsteps, the gait parameters appear on the *Create Multiple Footsteps dialog (page 2–992)*.
-  If creating footsteps manually with Create Footsteps (append) or Create Footsteps (at current frame), the gait parameters on the *Footstep Creation rollout (page 2–988)* are used.

Changing these values changes the timing for any footsteps placed after the values are set.

Changing the Gait After Creating Footsteps

After creating the footsteps, you can change the gait by editing footstep timing in Track View's Dope Sheet mode. See *Editing Footstep Timing (page 2–869)*.

Gait parameters are only one way to define the timing and nature of the biped's gait. For a more complete description of gaits and other parameters that alter the nature of the biped's motion, see *Adjusting Vertical Motion (page 2–878)*.

Creating Footsteps Automatically

Automatic footstep creation is the easiest way to create a walk or run cycle. You can use this method to make the biped climb or descend a flight of stairs, hop repeatedly, and do a variety of motions. Automatic footstep creation places the footsteps for you, generating perfectly timed and spaced footsteps.

Procedure

To create footsteps automatically:

1. Select any part of the biped.
2.  On the Motion panel > Biped rollout, click Footstep Mode.

   You are now in Footstep mode, where you can create, activate, or edit footsteps.
3.  In the Footstep Creation rollout, choose the gait you want to use for the footsteps: Walk, Run, or Jump.

   The selected gait determines the timing pattern of the automatically placed footsteps.
4.  In the Footstep Creation rollout, click Create Multiple Footsteps.

   The *Create Multiple Footsteps dialog (page 2–992)* appears. This dialog determines various aspects of the footstep sequence, such as how far each footstep will be placed from the previous footstep.
5. Set the multiple footstep parameters, and then click OK.

   This places footstep icons in the scene. To make the biped move through the footsteps, you must create keys for the footsteps.
6.  Click Create Keys For Inactive Footsteps on the Footstep Operations rollout.

   Keys have now been created for the footsteps.
7.  Click Play Animation to see the animation.

See also

Creating Footsteps Manually (page 2–863)

Adjusting Vertical Motion (page 2–878)

Creating Footsteps Manually

Creating footsteps manually is useful for complicated footwork, such as dancing. Manual footstep creation allows you to place each new footstep carefully where you want it.

Multiple footsteps can also be generated automatically. To quickly generate a simple walking, running, or jumping footstep patterns, see *Creating Footsteps Automatically (page 2–862)*.

There are two ways to create footsteps manually:

-  You can start creating footsteps at the current frame with Create Footsteps (at current frame) in the *Footstep Creation rollout (page 2–988)*. Any footsteps added subsequently will extend in time from the first footstep. If you attempt to create a footstep at the same time as an existing footstep on the same side, an alert appears and you are not allowed to create the footstep.
-  You can also append footsteps onto the end of the existing footsteps with Create Footsteps (append) in the *Footstep Creation rollout (page 2–988)*. Then Biped computes the frame at which the first footstep should be created based on the chosen gait and the existing footsteps. This option is disabled if there are no existing footsteps.

 As with all footsteps, any footsteps you create with these methods are inactive upon creation. You must activate them with Create Keys for Inactive Footsteps in order to make the biped move through them.

Preparing to Create Footsteps Manually

Before placing footsteps, you must choose the gait, then set parameters for the gait in the *Footstep Creation rollout (page 2–988)*.



Footstep creation and gait buttons

When you select the Walk, Run, or Jump gait, parameters for that gait appear on the Footstep Creation rollout. See *Choosing a Gait (page 2–861)* for information on gaits and parameters.

Switching Between Left and Right Footsteps

By default, manually placed footsteps alternate left, right, left, right. The footstep cursor, which appears during individual footstep creation, shows you whether a left or right footstep will be placed next.

If you want to change the footstep side that will be placed when you next click, press Q to toggle between them. Pressing Q also changes the cursor, showing you which footstep side will be placed next.

Automatic Time Extension

During manual footstep creation, your display remains at the current frame. However, footsteps are created sequentially in time. If the footsteps you create require more frames than are in the active time segment, the footstep creation process extends the active time segment, which can create new frames.

Creating New Footsteps Between Existing Footsteps

If inactive footsteps exist, you can only create new, inactive footsteps within or directly before or after the time of the inactive footsteps. If you need to add more footsteps near a time where there are only active footsteps, deactivate all footsteps first, then add footsteps and activate them again. See *Deactivating Footsteps (page 2–865)*.

Procedures

To prepare for manual footstep creation:

1.  On the Motion panel > Biped rollout, click Footstep Mode.

   You are now in Footstep mode, and can create, activate, or edit footsteps.

2. In the Footstep Creation rollout, choose the gait you want to use: Walk, Run, or Jump. See *Choosing a Gait (page 2–861)* for details on gaits and their parameters.

3. In the Footstep Creation rollout, set parameters for the chosen gait.

To create footsteps beginning at the current frame:

1.  Turn on Footstep Mode.

2.  Click Create Footsteps (at current frame).

3. Click in a viewport to create a footstep. Move the cursor and click again to create another footstep. Repeat until you have created all the footsteps required.

   Tip: Use the Top viewport to create the footsteps.

   Watch the footstep cursor to see whether a left or right footstep will be placed next. If you want to place two left or two right footsteps sequentially, place one footstep, then press Q once before placing the next one.

4. Move and rotate footsteps until you have achieved the desired pattern.

5.  Click Create Keys For Inactive Footsteps.

6. 

   Click Play Animation to see the animation.

To append footsteps onto the existing footsteps:

1.  Make sure Footstep Mode is turned on.

2.  If the existing footsteps are active, deactivate them first. Select all footsteps and click Deactivate Footsteps.

3.  Click Create Footsteps (append).

4. Click in a viewport to create a footstep. Move the cursor and continue clicking to create more footsteps

   Tip: A Top viewport is usually best when you create footsteps individually.

5. Move and rotate footsteps until you have achieved the desired pattern.

6.  Click Create Keys For Inactive Footsteps.

7.  Click Play Animation to see the animation.

To create footsteps using AutoGrid:

The AutoGrid feature creates objects on a surface other than the default construction plane. You can use this feature to place footsteps on an irregular surface.

1. Using any 3ds Max modeling method, create a surface as the terrain for the biped to step on.

2. On Create panel > Systems, click Biped. In the Object Type rollout, turn on AutoGrid.

3. If you already have a biped to use, return to the Motion panel without creating a biped. If you

need a biped, create the biped at this time, then access the Motion panel.

4.  Turn on Footstep Mode.
5.  Click Create Footsteps (at current frame). Move the cursor over the terrain. A transform gizmo moves with your cursor to indicate the location and orientation of the gizmo. Click to place a footstep, then move the cursor and click again to place more footsteps.
6.  When you have finished placing footsteps, click Create Keys For Inactive Footsteps.

   Keys have now been created for the footsteps.
7.  Click Play Animation to see the animation.

   The biped steps on the terrain, following the footsteps.

See also

Choosing a Gait (page 2–861)

Creating Footsteps Automatically (page 2–862)

Editing Footstep Placement (page 2–869)

Activating Footsteps

 When footsteps are created, they are inactive. To make the biped move through the footsteps, you must activate them by clicking Create Keys for Inactive Footsteps on the *Footstep Operations rollout (page 2–990)*.

When you activate footsteps, keys are created for the biped's legs and feet, causing the biped to step through the footsteps. In addition, keys are created for the spine, arms, hips, and tail (if the biped has one) to make the biped move naturally through the footsteps. Keys are not created for the head, ponytails or props.

When keys are created for footsteps, a generic type of motion is generated for the biped's body based on footstep placement and timing. This movement is intended as a starting point for your own animation, not as a final setup.

There are times when it is appropriate to deactivate footsteps, make changes, and activate footsteps again. See *Deactivating Footsteps (page 2–865)*.

Procedure

To activate footsteps:

1.  While in Footstep mode, create footsteps using the automatic or manual method.
2. Move, rotate, delete and edit footsteps as desired.
3.  Click Create Keys for Inactive Footsteps in the Footstep Operations rollout.
4.  Click Play Animation to see the animation.

See also

Creating Footsteps Automatically (page 2–862)

Creating Footsteps Manually (page 2–863)

Deactivating Footsteps

When you change the position or timing of active footsteps, the biped's animation changes accordingly, causing the biped to step into the footsteps in their new positions with the new timing.

Deactivating footsteps temporarily suspends changes to animation when footsteps are changed. When you move deactivated footsteps or change their timing, the biped still continues to do the same motion it did before footsteps were deactivated.

Footsteps should be deactivated when you make substantial changes to footstep timing or placement.

 Footsteps are deactivated by selecting footsteps and clicking Deactivate Footsteps in the *Footstep Operations rollout (page 2–990)*.

 After making changes to the footsteps, click Create Keys for Inactive Footsteps in the *Footstep Operations rollout (page 2–990)*. This will recreate keys for the biped and cause it to follow the footsteps.

Active footsteps are pale in color, while inactive footsteps are brightly colored. This coloring appears both on footstep icons in viewports and on footstep keys in the Dope Sheet.

Note: Deactivating and reactivating footsteps causes all biped keys to be replaced with default animation keys for the current footstep pattern and timing.

Procedure

To deactivate footsteps:

1. Select the footsteps to deactivate.

   The footsteps you select must be in sequence; you cannot deactivate a nonsequential set of footsteps.

2.  Click Deactivate Selected Footsteps in the Footstep Operations rollout.

   The footsteps are deactivated. Viewports and the Dope Sheet display them in bright colors again.

When to Deactivate Footsteps

There are many situations where you will want or need to deactivate footsteps after activating them.

- When you move or rotate active footsteps substantially, you might find that the original keys generated are no longer appropriate for your animation. If this happens, deactivate footsteps, move the footsteps into the correct locations, and activate footsteps again.
- When changing footstep timing in the Curve Editor's Dope Sheet mode, you might receive an error message stating that what you are trying to do violates one of the rules of biped key placement. If this occurs, deactivate footsteps, make timing changes in the Dope Sheet, then activate footsteps again
- If you have a large number of active footsteps in the scene, moving footsteps and making timing changes can cause delays as all keys are adapted and updated. You might find your work is faster if you deactivate footsteps while making your changes.
- If you have activated footsteps then changed the default animation, then find you want to return to the default animation, you can deactivate footsteps and activate them again to recreate the default animation.

Activation/Deactivation Workflow

The fastest workflow with footsteps is to deactivate footsteps and make changes, reactivate footsteps and play the animation to check your work, and repeat this process until the footstep timing and placement are correct. However, deactivating and reactivating footsteps replaces any biped keys with the default animation for the current footstep pattern and timing.

For this reason, it is recommended that you work with footsteps exclusively first, ignoring the upper body animation, and deactivate and reactivate footsteps as needed. When the motion of the feet is perfect or near-perfect, only then should you adjust the biped animation manually. Otherwise, you will lose all your manual animation work every time you deactivate and reactivate footsteps.

Tip: Alternatively, you can use *layers (page 2–974)* to store any upper body animations you want to preserve while you proceed with changes in your biped's footsteps.

Rules for Inactive Footsteps

When you click deactivate footsteps, only selected footsteps are deactivated. If you haven't made any changes to the default animation, it is recommended that you deactivate all footsteps. However, if you've changed or added animation keys to part of the animation, you might want to deactivate only a portion of the footsteps.

When some but not all footsteps are currently deactivated, limitations are imposed on the changes you can make to footsteps:

- You cannot set, delete, move, or edit keys in any way; footstep tracks are disabled in the Dope Sheet.
- You can create footsteps only when the new footsteps are among and adjacent to existing inactive footsteps.
- You can delete only inactive footsteps, and at least one inactive footstep must remain.
- You cannot deactivate a set of nonsequential footsteps.

See also

Activating Footsteps (page 2–865)

Understanding Footstep and Body Keys

When you activate footsteps, keys are created for biped body parts. Before you work with footstep placement and timing, it's essential that you understand the default keys generated and how they relate to footstep editing.

Leg States

A leg moving through a footstep has four states, beginning with the foot on the ground. Then the foot lifts, moves through the air, and returns to the ground again. These *leg states* are labeled by **character studio** to help you work with footsteps and leg/foot keys more easily when editing footsteps.

These states are referred to in these topics to aid in understanding more advanced footstep concepts, so it's important that you become familiar with them.

- **Touch** occurs at the leg keyframe where the foot first touches the footstep, and always corresponds with the start frame of a footstep in Track View. The Touch state is always one frame long.
- **Plant** occurs after touching, and before lifting. It is always in between the start and end frames of a footstep in Track View. The Plant state is usually more than one frame long.
- **Lift** occurs at the keyframe where the foot lifts off the ground, and always corresponds with the end frame of a footstep in Track View. The Lift state is always one frame long.
- **Move** occurs while the foot is in the air, and is always in the intervals between footsteps in Track View. In walking, while one foot moves, the body is supported by the other leg. In running or jumping, there are periods where the body is not supported and moves in midair.

For more information on these keys, see *Editing Footstep Timing (page 2–869)*.


You can display the current state of each leg by turning on Leg States in Biped rollout > Display group. The state name appears in viewports near the foot, corresponding to the state at the current frame.

Note: State names appear only on footsteps that have been previously activated. If active footsteps have been deactivated and edited, the state names that appear will correspond to states the legs were in *before* deactivation. For this reason, leg states displayed between deactivation and reactivation might not be accurate.

Regardless of the footstep pattern, activating footsteps will always cause leg keys to be generated at each lift and touch frame. Although you can alter these keys to some degree by going to a touch or lift keyframe, turning on Auto Key and changing the leg positions or rotations, you cannot delete these keys.

To change the foot/leg animation after footsteps are activated, see*Animating Legs and Feet (page 2–874)* and *Adjusting Body Keys in Track View (page 2–875)*.

Body Keys

When foosteps are activated, keys are also generated for the center of mass (COM) object, the blue tetrahedron at the center of the biped's pelvis. The position and rotation of the COM determines the horizontal and vertical position of the biped body, as well as its rotation.

Unlike other biped parts or 3ds Max objects, the COM has three separate animation tracks, two for motion and one for rotation. You select these tracks by clicking the appropriate button on the *Track Selection rollout (page 2–945)*:

-  Body Horizontal
-  Body Vertical
-  Body Rotation (displays keys for the Body Turning track)

Keys for these tracks are generated automatically when you activate footsteps. You can view the keys for each track on the Track Bar by turning off Footstep Mode and clicking the corresponding button listed above. These tracks are also visible in Track View's Dope Sheet mode.

Note: When you select the COM, only the keys corresponding to the selected track are displayed. If you select the COM and do not select a track, no keys are displayed on the Track Bar, even if the COM is animated. You can find out which COM track is selected by checking the status of the buttons displayed above. When a button is turned on, the track is selected.

You can animate the COM or change existing keys after footsteps are activated. To change COM keys after footsteps are activated, see*Animating the Upper Body (page 2–874)* and *Adjusting Body Keys in Track View (page 2–875)*.

Editing Footstep Motion

The topics in this section deal with the editing of footstep motion when animating bipeds. They are:

Editing Footstep Placement (page 2–869)

Editing Footstep Timing (page 2–869)

Editing Active Footsteps (page 2–871)

Splicing Footsteps (page 2–872)

Animating Legs and Feet (page 2–874)

Animating the Upper Body (page 2–874)

Adjusting Body Keys in Track View (page 2–875)

Shifting the Biped's Balance (page 2–876)

Adjusting Vertical Motion (page 2–878)

Saving Footstep Animation (page 2–882)

Editing Footstep Placement

After you've created footsteps, there are several ways to edit their placement in the scene.

- You can use standard 3ds Max controls to select, move, rotate, or delete footsteps.

  Note: The 3ds Max scale transform will not scale the footstep icons themselves; the size of the icons is set by the size of the biped's feet. The biped foot size is originally set by the biped's height, then later by any scaling you perform on the feet in Figure Mode. See *Posing the Biped (page 2–847)*.

- You can use the Bend and Scale controls in the *Footstep Operations Rollout (page 2–990)* to edit the footsteps. The Bend value curves a selection of footsteps, while Scale changes the distance between selected footsteps.

Selecting Footsteps

To change footstep placement, you must first select the footsteps to change.

 When Footstep Mode is turned on, footsteps can be selected with the same methods used to select 3ds Max objects:

- Click footstep icons while holding down the Ctrl key.
- Click and drag to draw a bounding box around footsteps you want to select.
- Unselect footsteps by holding down the Alt key and clicking footsteps.

You can also select footsteps in Track View's Dope Sheet mode. See *Editing Footstep Timing (page 2–869)*.

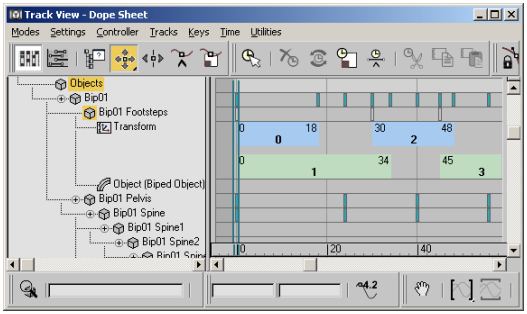
Biped Key Adaptation

When you move, rotate, delete, bend, or scale active footsteps, the software automatically alters biped keys to account for changes in balance and motion. This change is called *adaptation (page 3–905)*. For information on how adaptation occurs, see *Editing Active Footsteps (page 2–871)*.

Editing Footstep Timing

You can edit footstep timing in Track View.

Footstep keys have their own display in Track View's Dope Sheet mode. Footstep keys are displayed as blocks of color, each indicating a period of time when one of the biped's feet is in a footstep.



Footsteps in Track View

The footstep key colors match the colors of footstep icons in viewports. Just as with footstep icons, inactive footsteps are brightly colored, while active footsteps are displayed with pale colors.

The footstep key display relates directly to the different *states (page 2–867)* the leg and foot can be in at any given time. The left edge of each footstep key indicates the foot is in the Touch state, while the right edge indicates the Lift state. The duration of the footstep key between Touch and Lift indicates the time the foot is in the Plant state. The space between footstep keys indicates a time when the foot is completely out of a footstep, a Move state.

You can change footstep timing in Track View by:

- Moving footstep keys to an earlier or later time.

- Moving footstep keys to change the time relationships between footsteps.
- Changing the lengths of footstep keys.
- Deleting footstep keys.

You must select footstep keys before you can change their timing.

Selecting Footsteps for Time Editing

Selecting footstep keys is the same as selecting footstep icons; when one is selected, the other is selected as well. Footstep keys can be selected in a number of ways for editing in Track View.

Note:  To select footsteps, you must have Footstep Mode turned on.

- Select footsteps in viewports as you would any 3ds Max object. Click a footstep to select it, and hold down the Ctrl key and click additional footsteps to select them. You can also draw a selection box around footsteps to select them.

  Selecting footsteps also selects the footsteps in Track View. A selected footstep in Track View has a white border and a dot on both the left and right edges of the footstep key.



Footstep key selected in Track View

- Click the center of a footstep key in Track View to select it. Hold down the Ctrl key and click the center of additional footstep keys to select multiple footsteps.
- To select several footsteps at once, draw a selection box around the footstep keys in Track View.
- You can also select one edge of a footstep key by clicking the left or right edge in Track View. A single white dot appears at the selected edge. You can change the duration of the footstep by increasing or decreasing the length of the key on the selected edge.



Left edge of footstep key selected in Track View

- To select an edge, right-click the footstep track in Track View to display the *Footstep Mode dialog (page 2–1000)*. Select one or more footstep keys, and use the controls in the Footstep Edge Selection group to select an edge.

Editing Footstep Timing

Editing the timing of footsteps is a simple matter of moving selected footsteps or selected edges in time.

Moving an entire footstep changes the time at which it starts and ends, but not its duration. Moving a footstep edge changes the duration of the footstep.

Restrictions on Footstep Time Editing

The timing for both active and inactive footsteps can be changed. Editing timing for active footsteps also changes keys for the center of mass and leg keys, but not the spine, arm or tail keys. This can cause the resulting animation to look different from the way you intended. For this reason, you might prefer to deactivate footsteps before editing timing. See *Deactivating Footsteps (page 2–865)*.

When editing footstep timing in Track View, the following limitations apply to both inactive and active footsteps:

- There must be at least one frame between footstep keys for the same foot.
- Footstep keys cannot start in negative frames.

- Every footstep must be at least two frames long.
- If you want to add footstep keys, use the tools on the Footstep Creation rollout. See *Creating Footsteps Automatically (page 2–862)* and *Creating Footsteps Manually (page 2–863)*.
- Footsteps cannot be cloned.
-  To scale time of footsteps and body keys together, you can use Scale Time on the Track View toolbar.

  Warning: Do not use the Re-scale Time option in the Time Configuration dialog to scale an animation that contains active footsteps. If you do so, biped keys might not adapt appropriately. Instead, use Scale Time on the Track View toolbar.

If you attempt to perform an illegal footstep timing change in Track View, an error message appears and the footstep timing pattern returns to the state it was in before you made edits.

Procedures

To display footstep keys in Track View:

1. Choose Graph Editors menu > Track View - Dope Sheet.
2. If necessary, expand the Objects listing to display footstep keys.

To move selected footsteps in time:

- In Track View, drag any selected footstep key to drag all selected footstep keys in time.

To change the duration of a footstep:

1. In Track View, click the left or right edge of a footstep key.

   In addition to a white border, a small white dot appears to indicate the edge is selected.
2. Drag the selected edge to make the footstep key longer or shorter.

   You can adjust a selection of multiple footsteps this way.

   You can also adjust any combination of left edge, right edge, and center selections. Dragging in Track View adjusts any combination of edge or center selection.

To change the display of numbers of footstep keys:

1. In Track View, right-click the footstep track to display the *Footstep Mode Dialog (page 2–1000)*.
2. Choose the values to display from the Footstep Number Display group.

To scale keys in time:

1. On the Track View toolbar, click Scale Time.
2. In the Track View window, click and drag to indicate the beginning and end of the time period you want to scale.
3. Move the cursor over the selection region until a scale cursor appears. Drag to the left to reduce time, or drag to the right to increase time.

See also

Adjusting Body Keys in Track View (page 2–875)

Editing Active Footsteps

One of the most powerful features of footsteps is the ability to adapt keyframes automatically to edits you make to your footstep pattern. By analogy, the footsteps become a kind of *gizmo (page 3–949)* for manipulating the keyframes of your character's animation. In most cases, edits you make to footsteps act upon your keys in an intuitive fashion.

In order to understand what happens when active footsteps are edited, you must first understand what happens to biped keys when footsteps are activated. See *Understanding Footstep and Body Keys (page 2–867)*.

Key Adaptation for Footstep Placement Edits

When you move or rotate active footsteps as described in *Editing Footstep Placement (page 2–869)*, the biped's keys are immediately influenced by the edit. The following keyframe tracks are influenced in the vicinity of the edited footsteps:

- Body Horizontal keys change to step or hop within range of new footstep locations
- Body Vertical keys change to match possible changes in stride length, since the center of mass must move lower to step longer distances
- Body Turning (Body Rotation) keys change to bank into turns created by changes in path curvature or body speed
- Right or left leg keys in a Move state (when the leg is not in a footstep) are adapted to step between new locations

Other keys are preserved in the adaptation process. For example, if you animate the body to kneel while turned to the right, altering a nearby footstep will maintain the biped's height and turning motion as much as possible.

Key Adaptation for Footstep Timing Edits

To edit active footsteps in time, follow the methods described in *Editing Footstep Timing (page 2–869)*. Keyframes affected by the edit are updated immediately.

A fundamental factor in how your keyframes are adapted is whether the sequence of leg support transitions has changed. Changing the relationship between opposite leg keys effectively changes the gait, which causes entirely new keyframes to be generated at that point.

- If you move footstep keys in such a way that you remove an overlap or create a new overlap between opposite legs, you have changed the gait, and new keys will be generated for legs.
- If you do not change the overlap relationship between opposite legs, existing legs keys are retained.

Warning: Editing a jumping sequence so that the feet hit or leave the ground slightly out of phase rather than at the same time introduces a change in the leg support sequence. Hence, even in this case, the leg keys will be locally replaced with default keys.

Note: If you scale the biped's feet in Figure mode after activating footsteps, then turn off Figure mode, **character studio** will recalculate the footstep display size and the footstep dynamics automatically.

Locking Keys

You can lock specified tracks to prevent them from adapting automatically when you move active footsteps. The controls in the Footstep Adapt Locks group on the *Dynamics & Adaptation rollout (page 2–980)* lock specified tracks.

Procedure

To prevent keys from changing when active footsteps are edited:

1. On the Motion panel, on the Dynamics & Adaptation rollout, in the Footstep Adapt Locks group, select the types of keys you want to lock.
2. Move, rotate, scale, or bend active footsteps.

   The tracks you lock are unaffected by footstep editing.

Splicing Footsteps

You can copy the motion of a biped footstep sequence, and paste it either at the end or into the middle of another footstep sequence. This technique is called *splicing (page 3–1015)*.

You can use splicing to build an extended, or cyclic, motion sequence out of shorter sequences. You can also copy footsteps from one biped and paste them onto another biped with this technique.

Splicing footsteps makes use of a *footstep buffer* to store footstep positions and keys. The buffer is a temporary area in memory where footsteps are stored.

Tip: Use motion splicing when you want to edit segments of a single footstep-based animation. To build longer animations by connecting finished *.bip* files, use *Motion Flow (page 2–1026)* or the *Motion Mixer (page 2–604)*.

Procedures

To copy footsteps to the buffer:

1.  Turn on Footstep mode.
2. Select a series of footsteps.
3.  In the Footsteps Operations rollout, click Copy Footsteps.

To edit the footstep buffer:

You can edit the footstep buffer only when footsteps have been copied to the buffer, as described in the previous procedure.

1.  On the Biped rollout > Modes group, click Buffer Mode. This button is active only when there are footsteps in the buffer.

   The viewports now display the footsteps in the footstep buffer, rather than the footsteps in the currently activated footstep sequence.
2. Edit footsteps or motion keys as you normally would – by moving and rotating footsteps, or turning on Auto Key and moving and rotating parts of the biped skeleton.

   Note: You can't splice (copy and paste) footsteps while in Buffer mode.

   If you load a BIP file by clicking Load File on the Biped rollout while Buffer mode is active, the footstep buffer is replaced with the motion in the BIP file.
3.  Click Buffer Mode to turn it off.

   Now you can splice the edited buffer by pasting footsteps.

To splice the footstep buffer:

1.  Click Paste Footsteps on the Footstep Operations rollout.

   The footsteps from the buffer appear in their saturated colors.
2.  Use Select and Move from the 3ds Max toolbar to position the first pasted footstep at the point where you want to splice. When you move the first pasted footstep over an activated footstep of the same side, right or left, the activated footstep turns red to indicate splicing is possible.

   This step is usually easiest to do in a Top viewport.

   Note: If you deselect the pasted footsteps by clicking elsewhere in a viewport, no paste occurs and the pasting process ends.
3. When you have moved the buffer footsteps to the desired location, release the mouse.

   The contents of the footstep buffer are placed over the active footsteps. The first footstep in the buffer replaces the footstep that turned red, and the rest of the buffer footsteps follow.
4. The footstep that turned red and the remaining original footsteps now appear in the viewports. They will be near the activated footsteps, in their saturated colors. They can now be moved and pasted onto the end of the spliced motion, which will, in effect, perform an insert of the buffer motion. The red footstep and is

following footsteps can also be pasted anywhere in the footstep sequence. If you pasted onto the last footstep, this step does not occur.

If you don't want to append the leftover footsteps, simply click anywhere in the viewport to dismiss them.

Warning: If any footstep in the buffer overlaps in time with a footstep previous to the one onto that you are pasting, a message appears and the paste will not be performed.

Tip: To create a cycle of a motion with alternating footsteps, you must copy and paste at least three footsteps.

See also

Footstep Operations Rollout (page 2–990)

Animating Legs and Feet

When you activate footsteps, keys are created for the legs and feet according to the footstep pattern and timing. Leg and foot keys are set at each Touch and Lift frame, and between footsteps.

Although the parts of the leg (thigh, calf, foot, toes) can move and rotate separately, one key dot appears on the Track Bar and in Track View for each leg structure unless you have turned on separate tracks in the *Keyframing Tools rollout (page 2–962)*.

When the foot is not in a footstep, you can move or delete keyframes, and animate or reposition any part of the leg or foot. Simply turn on Auto Key and move or rotate the leg or foot to its new position.

Animating Legs and Feet in Footsteps

The leg keys created at either end of each footstep are *locked*, meaning they cannot be moved or deleted. You can, however, change the animation at these frames, or add new keyframes within the timespan of the footstep key.

When the foot is in a footstep, you can animate or reposition the leg and foot to some degree:

- You can rotate the foot around its pivot point within the footstep.
- You can animate the pivot point with the Select Pivot option in Key Info rollout > IK bar. See *Animating Pivots (page 2–905)*.
- You cannot move a foot out of a footstep. If you do not want the foot to be in the footstep at a particular frame, change the timespan of the footstep key so it does not include that frame. See *Editing Footstep Timing (page 2–869)*.

Animating the Upper Body

When you use footsteps to animate the legs and feet, animating the spine, neck, head, arms, and other upper body parts is accomplished in the same way as for freeform animation. Simply turn on Auto Key and start moving and rotating body parts. See *Animating by Moving Links (page 2–890)* and *Animating by Rotating Links (page 2–891)*.


You can also click the Set Key button on the *Key Info rollout (page 2–954)* to set keys for selected body parts at the current frame. This option allows you to move body parts as you like, setting keys only when you have decided to do so.

Warning: 

Do not use the 3ds Max Set Key button to set keys for the biped.

Animation on the center of mass (COM) works differently, depending on whether the biped is in a footstep at that particular frame.

- If at least one foot is in a footstep, moving or rotating the COM will animate the entire body along with the COM, except a foot that is in a footstep. A foot in a footstep will stay planted in

the footstep and the leg will bend or straighten as the COM is moved up or down.

- If the biped is airborne (no footsteps at that frame), the COM can be moved horizontally and can be rotated, but cannot be moved vertically. This is because **character studio** automatically calculates how high the biped can go based on the biped's height and the current gravity setting. To make the biped jump higher or lower during airborne periods, see *Freeform Animation Between Footsteps (page 2–883)*.

Working with Existing Keys

Because keys are created for the spine, arms, and tail when footsteps are activated, you will want to take care when creating new keyframes. As with 3ds Max animation, placing keyframes one or two frames apart creates jerking or popping motions.

To avoid placing keyframes on upper body parts too close together, you can do one of the following:

-  Turn on the 3ds Max Key Mode toggle, select the body part, and use Next Key and Previous Key to move in time from one keyframe to another. Animate only on existing keyframes.
- Delete all the body part keys except the key at frame 0, and create entirely new keys throughout the animation.

  Warning: If you delete all keys, including the key at frame 0, be sure to set a key at frame 0 before animating on other frames. **character studio** does not automatically place a key at frame 0 when you animate on frames other than 0, as 3ds Max does.

Interdependencies Between Legs and Upper Body

When footstep animation is used, animation of the upper body and animation of the legs are interdependent. Any time you change the position of the body's center of mass, leg positions are automatically updated to ensure the feet still move from footstep to footstep.

Interdependencies between the upper body and legs are updated after a key has been set, sometimes causing the appearance of a pop in the leg position. However, in playing the animation you will find the motion to be smooth.

When you set a leg key while the foot is on the ground, the location of the foot is automatically recalculated based on the natural roll of the foot from its point of initial contact with the ground. The pivot points are based on collisions between the footstep plane, the two corners of the heel, and the bones that connect the toes' links.

This ability to preserve interdependent relationships allows you easily to experiment and improvise with motion, since you are assured that edits to one body part will not corrupt the integrity of other parts.

In general, changes to the center of mass, spine, and pelvis cause the legs to adjust. Changes to the legs do not affect the center of mass's vertical position.

Adjusting Body Keys in Track View

When you activate footsteps, keys are created for biped body parts. After activating footsteps, you can also create keys by turning on Auto Key and animating body parts. All these keys can be viewed in Track View – Dope Sheet mode.

In general, editing biped body keys in Track View is the same as editing keys of any sort in 3ds Max, with a few restrictions.

Before reading this section, please read *Understanding Footstep and Body Keys (page 2–867)* for terms and definitions used in this section.

Note: The information in this section applies only to biped body keys for the arms, legs, spine, neck, and other parts of the upper body. For rules for editing footstep keys, see *Editing Footstep Timing (page 2–869)*.

Locked Keys

Some biped keys are locked, meaning they cannot be moved in time. Locked keys are colored red in Track View.

Locked keys for the COM Body Vertical track are created at each footstep's lift and touch frames. These are the only keys that are locked. These keys can be moved in time by changing footstep duration in Track View.

Although you cannot move or delete a locked key in Track View, you can sometimes change the animation on the frame by going to the frame, turning on Auto Key and manually animating the key.

Restrictions on Biped Body Key Editing

All options available for editing biped body key timing are available only when Edit Keys on the Track View toolbar is turned on.

The following rules apply to the editing biped body key timing in Track View:

- You cannot move existing keys outside the area between the first and last footsteps, or to negative frames.
- You can use Add Keys to add keys to any biped body track. You cannot add keys outside the area between the first and last footsteps, or to negative frames.
- You can clone keys the same way you clone 3ds Max object keys, by selecting keys and holding down **Shift** as you drag them. You can also clone keys by copying and pasting postures. See *Copying and Pasting Postures and Poses (page 2–910)*.
- The Track View – Curve Editor can be used to view and edit curves showing the interpolation between body keys. You can also use the Workbench to analyze, edit and fix biped motion curves. See *Working with the Workbench (page 2–1008)*.
- You can use Scale Keys to scale selected keys, except for locked keys.

Tip: If you want to animate a biped performing movements with its upper body while standing, extend the length of the standing footsteps to cover the length of the standing animation, then add keys for the upper body.

See also

Animating the Upper Body (page 2–874)

Shifting the Biped's Balance

After footsteps are activated, you can adjust the biped's overall balance in either of two ways:

-  In the *Biped rollout (page 2–936)*, turn on Rubber Band Mode and move the biped's COM, or
- Change or animate the Balance Factor parameter in the *Key Info rollout (page 2–954)*.

Rubber-Banding the COM

When Rubber Band mode is active, you can move the biped's center of mass in relation to the rest of the body, changing the biped's overall balance. Rubber-banding the center of mass defines your character's balance point in at-rest pose and in any resulting motion.

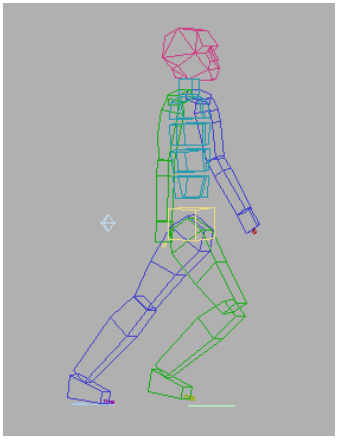
When a character is standing at rest, with feet side by side, the center of mass should be directly above the area where the feet touch the ground. The center of mass determines the center of your character's distribution of weight. For most

characters, the physical center of mass is located near the pelvis.

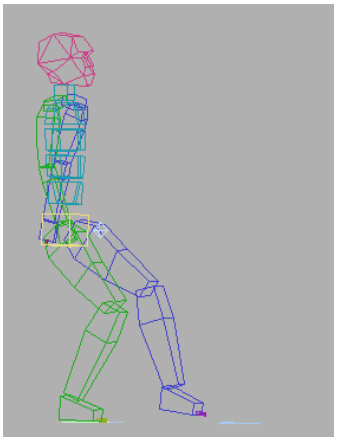
Typical placements for the center of mass are:

- Characters that stand erect, such as humanoids, use the default location of the center of mass within the pelvis.
- Characters that naturally lean forward, such as some dinosaurs and birds, should have the center of mass moved slightly forward of the pelvis. This is also good for robots and droid soldiers.
- Characters that are holding a heavy weight out in front of them, or overweight characters, might need their center of mass moved slightly forward of the pelvis.
- Characters that are pushing or pulling objects might need their center of mass moved slightly behind the pelvis.

  Note: If you place the center of mass too far in front of the pelvis, the character unnaturally compensates for balance. As each step is taken, the legs and body move in an awkward fashion, as if there was an invisible weight attached to the front of the character.



Center of mass moved behind the biped in rubber-band mode



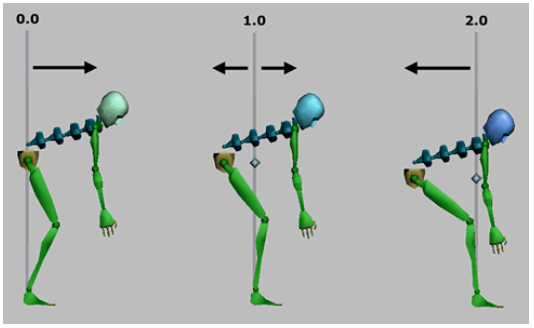
Center of mass moved in front of the biped in rubber-band mode

For a tutorial that uses this technique, see "Creating the Illusion of Weight".

Balance Factor

The *balance factor (page 3–914)* is an animatable parameter that determines how the biped's hips and spine will compensate when the biped bends forward or backward.

The balance factor determines the degree and direction in which the head and/or hips swing out from their original vertical alignments when the biped bends over.



Results when balance factor is set to 0.0, 1.0 and 2.0 before the spine is animated to make the biped bend over.

The Balance Factor parameter can have any value from 0 to 2:

- When Balance Factor is 1 (the default value) and you rotate the spine forward, both the hips and head swing out from their original vertical alignment to a similar degree to compensate for the shift in weight. This value is suitable for times when the biped leans forward while standing.
- When Balance Factor is 0 and you rotate the spine forward, the COM keeps its original vertical alignment, and the head swings forward. There is no movement in the hips to compensate for the shift in body weight. This setting is good for animating a sitting biped who leans forward, to prevent its hips from shifting backward. It can also work well for characters with tails; the tail provides weight in back of the biped, so the hips don't have to shift backward to compensate when the character bends over.
- When Balance Factor is 2 and you rotate the spine forward, the head retains its original vertical alignment while the COM swings backward. The hips compensate strongly for the shift in body weight. This setting can be useful for violent or acrobatic motions.

In order to affect the animation, the balance factor must be set before you animate the biped. The balance factor can be animated to give the biped different reactions to bending motions over the course of the animation.

The Balance Factor parameter appears in the Key Info rollout > Body section.

Procedure

To set the balance factor:

1.  On the Motion panel, turn off Footstep Mode if it is turned on.

   Note: The Key Info rollout is accessible only when Footstep Mode is turned off.

2. Expand the Body bar on the Key Info rollout.
3. Go to the frame where you want to set the Balance Factor.
4.  Click Body Horizontal to select this *track (page 3–1023)*.
5.  If no Body Horizontal key exists at the frame where you want to set the Balance Factor, create one by clicking Set Key in the Key Info rollout.
6. Change the Balance Factor parameter.

   There is no need to click Set Key again after setting the Balance Factor parameter.

7. Rotate the biped's spine.

   The biped's hips and/or head shift position based on the Balance Factor value.

Adjusting Vertical Motion

A *ballistic gait* is any footstep pattern in which there are periods with no feet on the ground, causing the biped to become airborne, or *ballistic*. Running and jumping are ballistic gaits, while walking is not.

When Biped Dynamics is turned on in the *Dynamics & Adaptation rollout (page 2–980)* and you activate footsteps, the Body Vertical keys that are automatically generated during a ballistic gait period will take gravity and landing speed into account. The parameters described in this section affect Body Vertical keys created in this manner.

Airborne Dynamics

With footsteps, each airborne period always begins and ends with Body Horizontal and Body Vertical keys. These keys define the position of the biped at lift-off and touchdown.

When the biped is airborne and Biped Dynamics is turned on, the vertical motion is governed by physically based dynamics. Its airborne trajectory

is based on the current gravity setting, the height of the Body Vertical key at lift-off and touchdown, and the amount of time in the air.

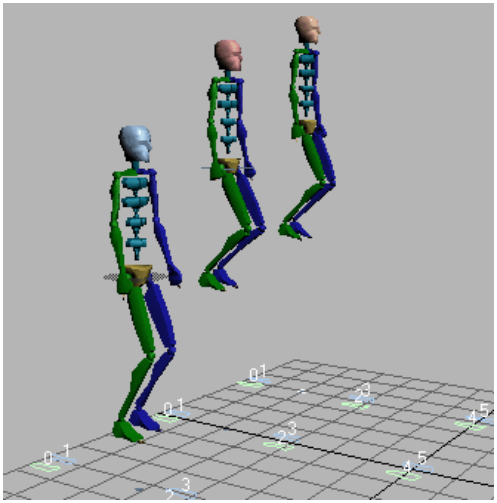
By default, there is no Body Vertical key at the peak of the biped's trajectory; the biped's peak airborne height is calculated and enforced automatically. You cannot, for example, set a Body Vertical key at the peak of an airborne trajectory and move the biped up or down. If you attempt to do so, the biped will snap back to its original airborne height. Although this will cause a *Body Vertical key (page 2–867)*to appear on the Track Bar and in Track View, the key will have no effect on the biped's airborne height.

Gravity and Timing

In reality, the length of time a person, animal or insect stays in the air during a jump is based on two factors:

- How high the creature jumps, which in turn is based on how hard the creature pushes with its legs at the start of the jump. The creature's weight has no bearing on the height of the jump, except to affect its ability to give a good push at the start. A very light character might be taken by the wind and thus stay in the air longer, but that circumstance is not part of the gravitational equation.
- The gravitational pull of the planet from which the creature jumps.

From these two factors, you can calculate how long the creature will stay in the air. You can also perform this calculation backward; if you know how long the creature was in the air and the gravitational equation for the planet, you can figure out how high it jumped.



Jump height increases with time in air

character studio uses the latter method for calculating the height of a biped's jump. It knows from footstep timing how long the biped will be in the air, and it has a method for determining the "gravitational pull" in your scene.

The default gravity setting in **character studio** is based on the standard equation for calculating the Earth's gravitational pull, which is an acceleration of approximately 32 feet/sec/sec. Because this equation depends on an accurate measure of distance, the biped's height is used as a guide to actual distances in the scene. For the purpose of gravitational computation, the biped is considered to be about 5'10" tall, the average height for a male human being.

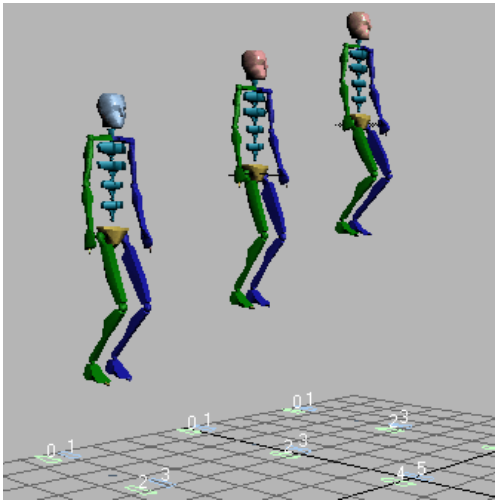
Of course, some of your characters will not be the same height as the average man, so airborne periods between footsteps for these characters will appear inaccurate, with the biped jumping too high or too low for your purposes.

There are two solutions to this problem. You can change the amount of time the character is airborne between footsteps (see *Editing Footstep Timing (page 2–869)*), or you can accelerate gravity with the GravAccel parameter.

Accelerating Gravity

The GravAccel parameter on the *Dynamics & Adaptation rollout (page 2–980)* alters the degree of gravitational pull imposed on the biped during its airborne periods between active footsteps.

The GravAccel default value is based on the standard calculation for the Earth's gravity and its effects on a man of average height. When the GravAccel parameter is increased, the effects of gravity are decreased, and the biped jumps higher. The GravAccel parameter can be set separately for each biped in the scene.



GravAccel values increase height of jump. (Left=500, Middle=1000,Right=1500).

You can change GravAccel at any time during the animation process to affect the biped's vertical airborne motion, both for keys already created and for animation not yet created. This value cannot be animated; the current GravAccel value is used throughout the biped's footstep sequence.

Changes to the GravAccel parameter have no effect whatsoever on the amount of time the biped spends in the air.

Tip: If you want to change GravAccel back to its default value but don't remember what that was, you can find out by creating another biped of the same height and noting its GravAccel value.

Note: If there is not enough gravitational acceleration or time in the air to account for differences in vertical height (such as a biped falling from footsteps on a ledge in slow motion), the biped is placed on the ground at the touchdown frame, causing a discontinuity. You can fix this by either increasing GravAccel or increasing the amount of time between footsteps in Track View. See *Editing Footstep Timing (page 2–869)*.

Touch and Lift Dynamics

When footsteps are activated for a footstep pattern that includes airborne periods, Body Vertical keys are set at each Touch and Lift keyframe.

Based on the time in the air and the biped's height, **character studio** calculates the height to which the biped will jump. From this, the software figures out how hard the biped's legs would have to push at the Lift frame to reach this height, and how much the legs would have to bend to absorb the force of the landing at the Touch frame. Body Vertical keys are set at the Lift and Touch frames accordingly.

You can adjust the timing of touch and lift dynamics by:

- Changing Body Vertical keys at the Touch and Lift frames.
- Setting, changing or deleting Body Vertical keys during a footstep.
- Changing the stiffness or springiness of the biped's legs at Touch and Lift frames with the Ballistic Tension parameter.

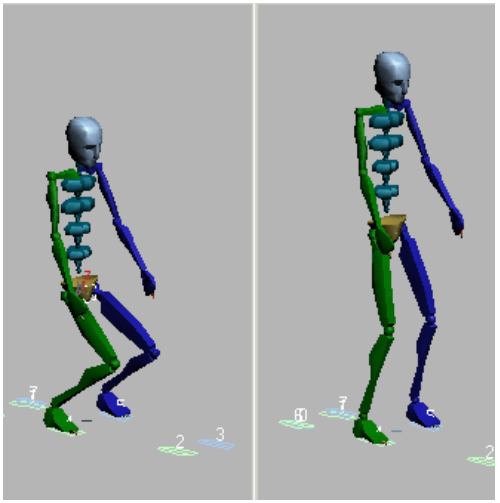
Ballistic Tension

Ballistic tension refers to the stiffness or springiness of the biped's landing at the end of a jump, or takeoff at the start of a jump. The *Ballistic Tension parameter (page 2–954)* changes the biped's motion only when it is set for footsteps

immediately preceding or following an airborne period.

You can edit the Ballistic Tension parameter only at Body Vertical keyframes where the body touches down (a Touch key) or lifts off (a Lift key). On the Motion panel *Track Selection rollout (page 2–945)*, Body Vertical must be selected and a key must be present at the current frame. With footsteps, a Body Vertical key is always set automatically at the landing footstep's Touch frame and the takeoff footstep's Lift frame.

Ballistic Tension can range from 0 to 1, with 0.5 as the default value. Increasing Ballistic Tension to 1.0 makes the legs stiffer at takeoff or landing. Decreasing the tension to 0.0 makes the legs springier and less stiff.



Ballistic tension determines springiness of landing. (Left = 0.0, Right = 1.0)

Dynamics Blend

The Dynamics Blend parameter in the *Dynamics & Adaptation rollout (page 2–980)* determines whether biped dynamics or spline dynamics will be used during airborne periods.

With *biped dynamics (page 3–916)*, the height and trajectory of the biped during airborne periods is determined by gravity calculations and other biped-specific information.

With *spline dynamics (page 3–1015)*, the airborne height and trajectory are not set for you automatically; you must set them manually with keyframes.

When Dynamics Blend is set to 1, biped dynamics are used. At 0, spline dynamics are used. You can animate this value to switch between the two during your animation.

Procedures

To locate vertical center of mass keys:

1.  Select any part of the biped and access the Motion panel.
2.  On the Track Selection rollout, click Body Vertical.
3.  Use the Next Key and Previous Key buttons on the Key Info rollout to jump to the next or previous center of mass Body Vertical keyframes.

   Tip: If the selected vertical COM key is a touchdown key from an airborne period, you can change the Ballistic Tension parameter on the Key Info rollout to control knee bend. Turn on Trajectories on the Display rollout to view how changing parameters affects the trajectory.

To set a Body Vertical key:

1.  On the Track Selection rollout, click Body Vertical.
2. 

Turn on Auto Key.
3.  Click Select and Move on the main toolbar.

4. Click and drag the center of mass up or down in the viewports.

   The biped is repositioned vertically, a key is created in the center of mass Body Vertical track.

To set ballistic tension:

1.  On the Biped rollout, turn off Footstep Mode if necessary.
2.  in the Track Selection rollout, select the biped's Body Vertical track.
3. Move to the Vertical track keyframe you want to adjust, either a Touch or Lift key.

   The Ballistic Tension parameter is available only on Touch and Lift keys.
4. In the Key Info rollout, expand the Body bar to access the Ballistic Tension parameter.
5. Adjust the Ballistic Tension value.

   Adjusting the ballistic tension changes the amount of crouch before a jump, and the amount of dampening that occurs after landing.

To change Dynamics Blend for multiple Body Vertical keys:

1. Choose Graphs Editors menu > Track View – Dope Sheet.
2. Select all the Body Vertical keys you want to change.
3. Right-click over one of the selected keys to display the TCB dialog.
4. Change the value of Dynamics Blend in the TCB dialog. This changes it for all selected keys.

   Note: This only affects the biped during airborne periods in a footstep animation. Changing Dynamics Blend keys in a walk motion does not affect the motion.

Saving Footstep Animation

After you have animated the biped with footsteps, there are two ways you can save the data.

- A BIP file saves both the footsteps and any animation on the upper body. You can save the entire footstep sequence, or just a portion of it.
- An .STP file saves just the footstep data. These files, in ASCII format, are sometimes used by game developers to save just the leg and foot motion data

 To save a either type of file, click Save File in the Biped rollout. Choose the desired file type from the Save as type pulldown.

For more detailed information on these file formats and how to work with them, see *Loading and Saving BIP Animation (page 2–920)* and *Loading and Saving STP Files (page 2–924)*.

Footsteps and Freeform Animation

character studio provides two distinct modes of animation, *footstep* and *freeform*.

Footstep animation (page 3–943) was originally designed for bipedal characters, that is, characters that walk on two legs. It uses a structured system of dynamics and balance to create secondary motion in the upper body, including arm swinging and tail swaying that is automatically applied when footsteps are used. By using default footsteps, you are assured that the biped's foot doesn't slide or pass through the ground, and that the foot easily rolls from heel to the ball of the foot as is common in a typical walk.

Freeform animation (page 3–943) provides a more traditional method of animating to those familiar

with 3ds Max, or other character animation systems. Keys are set by using the **character studio** Set Key tools found in the *Key Info rollout (page 2–954)*, or by turning on Auto Key mode, moving the time slider and transforming the biped parts.

Freeform animation includes preset key types called planted, sliding, and free keys; these make it very quick to lock a biped's foot or hand in space and get results similar to those of footstep animation.

You can add freeform animation to footstep animation during airborne periods, making it possible to use both footstep and freeform animation on the same biped.

Freeform animation is particularly useful for:

- Animating a character with more than two legs.
- Animating a biped switching between foot-based motion and freeform motion, such as a biped doing handstands, jumping into a swimming pool, or leaping up to begin flying through the air.

Freeform animation also provides animated pivot points on the feet and hands. This makes it possible for you to simulate the roll of a foot from heel to toe, or create detailed finger animation.

You can convert back and forth between footstep and freeform animation within some limitations. While it is easy to convert any footstep animation to freeform, there are limitations on the conversion from freeform animation to footsteps.

Because of these limitations, a common workflow is to use footstep animation to rough out the animation, then convert to freeform to finalize the motions. The totally freeform animation is not bound by the rules of the footstep method, so it's easy to do anything you want, such as add keys at the beginning or end of the clip.

See also

Footstep Animation (page 2–856)

Freeform Animation (page 2–886)

Freeform Animation Between Footsteps

For certain types of footstep animation, it is necessary to suspend calculations of gravitational effects on the biped.

For example, if you want the biped to run, dive into a pool then climb out of the pool, you would create footsteps for the running and climbing portions of the animation but not the swimming portion. When you activate footsteps, the biped will become airborne during the swimming period, jumping high in the air between the running and climbing footsteps.

If you don't suspend gravity during the airborne period, you won't be able to change the biped's Body Vertical position to bring it down and place it in the pool. For more information on how the biped behaves during airborne periods between footsteps, see *Adjusting Vertical Motion (page 2–878)*.

To suspend gravity during an airborne period, you can create a freeform period in your animation. During a freeform period, you can animate the biped in freeform mode, which includes positioning the biped anwhere in the scene regardless of the length of the airborne period.

The Track View - Dope Sheet editor is used to insert a freeform period between footsteps. During a freeform period, vertical dynamics (gravitational effects) are suspended, and Body Vertical keys use *spline dynamics (page 3–1015)* instead.

Procedure

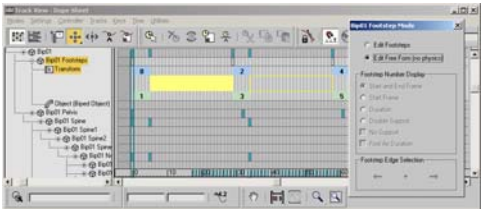
To convert an airborne period to freeform mode:

1. Create a footstep animation with at least one airborne period, and activate the footsteps. For the simplest method of creating a footstep pattern with an airborne period, see *Creating Footsteps Automatically (page 2–862)*.
2. Access Track View by choosing Graph Editors menu > Track View - Dope Sheet from the 3ds Max menu.
3. Expand the footstep key display for the biped.
4. Right-click the footstep key display area. The Footstep Mode dialog appears.

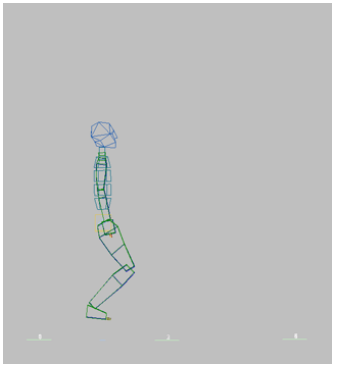


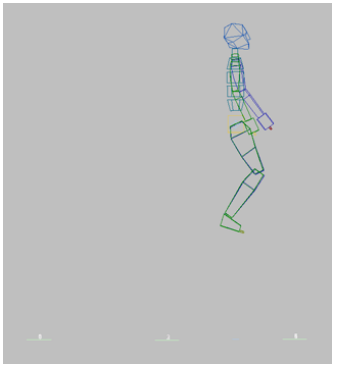
5. In the Footstep Mode dialog, choose the Edit Free Form (no physics) option.

   In each airborne period in Track View, a hollow yellow box appears.
6. In the Track View footstep key display, click inside the yellow box for the airborne period for which you want to suspend gravity. This fills the yellow box, indicating that vertical dynamics are suspended for that period.



7. Play the animation. The biped uses vertical dynamics during all footstep motions except the freeform period.





No vertical dynamics during freeform period, then return to vertical dynamics during next airborne period.

Note: If you activate footsteps after a freeform period has been set between footsteps, the biped might skip footsteps at one end of the freeform period. For this reason, you should activate footsteps before setting up a freeform period.

8. During the freeform period, animate the biped and set keyframes as you like.

   To edit footstep timing after creating a freeform period, you must choose the Edit Footsteps option in the Footstep Mode dialog.

See also

Track View (Biped) (page 2–1002)

Converting Between Footstep and Freeform Animations

There might be times when you want to convert a footstep animation to a freeform animation, or vice versa. This can occur when someone gives you an animation file that's animated in one mode, while you prefer to work in the other.

Animators tend to develop a preference for either freeform or footstep animation. The ability to convert between the two means you can work in either format without concern for compatibility with other artists.

 Using the Convert tool on the Motion panel > *Biped rollout (page 2–936)*, you can convert between footstep and freeform animation.

There are no special considerations when converting from footsteps to freeform; this procedure works for all footstep sequences. However, to convert successfully from freeform to footstep animation, you must observe certain rules in order for footsteps to be generated.

Converting from Freeform to Footsteps

When converting from freeform to footsteps, footsteps are generated for a foot only during periods between two planted keys. In other words, the conversion process looks for times when a foot has two consecutive, identical keys with an IK Blend setting of 1 and Object space selected. In this case, a footstep icon is placed at the foot's location, with the duration of the footstep set by the time between the two identical keys. This is the only condition under which foosteps are generated.

If converting freeform animation to footsteps does not create the expected result, your freeform animation might need adjustment prior to conversion. See *Using IK Keyframe Parameters (page 2–900)* for information on setting freeform planted keys.

If your freeform motion hasn't been prepared with planted keys for converting to footsteps, you can convert it to footsteps by loading it as a motion capture file and specifying footsteps as the extraction method. This technique uses a different method to determine footstep location and duration, and so can sometimes successfully convert to footsteps where the Convert command cannot. See *Importing Motion-Capture Data (page 2–1061)*.

Procedure

To convert between footsteps and freeform animation:

1.  On the Motion panel > Biped rollout, click Convert.

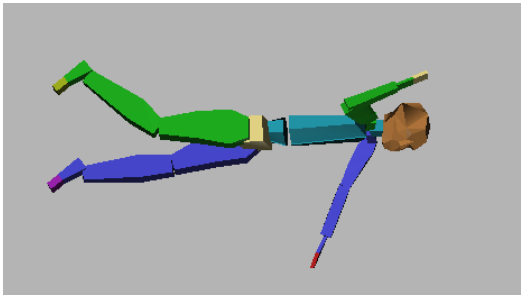
   Depending on which is the current animation method, either the *Convert to Freeform dialog (page 2–999)* or the *Convert to Footsteps dialog (page 2–999)* appears.

2. Change settings in the dialog as necessary, and click OK.

   If you're converting to freeform animation, the footstep patterns are removed from the viewport, gravity and dynamics are removed, and footstep mode is disabled. You are now working with a purely freeform animation.

   If you're converting to footsteps, and the operation is successful, footsteps appear in the viewport, *biped dynamics (page 3–916)* is turned on, and Footstep mode is enabled.

Freeform Animation



Freeform animation of a biped swimming

While **character studio** calculates vertical dynamics and gravity based on its footstep-driven technology, you don't always want your character strictly under these controls. You might want the character to fly, swim, or to do something improbable in a physical world. For these situations, Biped supports a comprehensive set of freeform animation controls that allow you to take total creative control over your character's pose, movement, and timing.

character studio gives you the option to animate characters using freeform mode exclusively, or to create freeform periods during a footstep animation. In freeform mode (without footsteps), you can pose every joint of your character exactly as you like using traditional keyframe methods. Hands and feet can be locked down in space by setting planted keys. You can also animate the pivot points for hands and feet to simulate rolling motion. You can blend dynamically between forward kinematics and inverse kinematics to introduce higher-level control.

To use freeform animation in Footstep mode, you create freeform periods between footsteps using Track View – Dope Sheet. You can convert an existing footstep animation into a freeform animation and back again. This lets you mix the use of footstep and freeform methods.

Warning: You can't add a freeform period at the beginning or end of a footstep animation.

Tip: If you want to do this, stretch out the timing for the first or last set of footsteps, then convert the footstep animation to freeform.

Freeform Method

A freeform animation contains no footsteps; instead it relies on the transforms of the biped body objects and center of mass. Use Freeform for motions like swimming or falling where footsteps are not necessary. If you are a familiar with creating all of your keyframes manually to animate a character, you may want to use the freeform method exclusively.

To start a freeform animation, turn on Auto Key and start positioning the biped. You can also leave Auto Key off and use the red set key buttons on the *Key Info rollout (page 2–954)* to create keyframes.

You can also create freeform animation by importing motion capture data and choosing freeform rather than footstep.

Tip: Take advantage of both methods by combining footsteps and freeform animation. You can create a freeform period for any airborne period between footsteps. A freeform period replaces the ballistic motion calculated using the GravAccel value with a user defined spline motion.

If you are using footstep extraction with motion capture data, you often need a freeform interval to accommodate falling or tumbling motion in the data. The Fit to Existing option on the Motion Capture Conversion Parameters dialog allows for a combination of both methods. Extracting footsteps from motion capture files eliminates sliding feet, a common problem with motion capture data.

Note that while you can add a freeform period to a footstep animation, you *cannot* add a footstep period into a freeform animation. If you want to

add a footstep animation to an existing freeform, you can use the motion flow editor to create a script that sequences the footstep with the freeform.

Inverse Kinematics

Footstep and freeform animations use the same *inverse kinematics (page 3–958)* constraints and extensions. This means that in a footstep animation, you can now edit keys to change footstep duration. By definition, a footstep is the start and end of a sequence of IK constraints in World Space with an IK Blend value greater than 0. Deleting and inserting keys or changing IK space or IK blending alters footstep duration.

In cases involving edits that alter the length of ballistic intervals (when a biped is airborne), the software ensures that there is a vertical key occurring at the lift-off and touchdown frames. This calculates the correct ballistic motion, so vertical keys are automatically inserted if not present.

There are three types of IK keys you can create: *planted*, *sliding* and *free* keys.

- **Planted keys** have an IK Blend of 1. They are joined to the Previous IK Key and are in Object space, rather than Body space. Planted keys lock the foot or hand to the ground, or to any object.
- **Sliding keys** have moving IK constraints. Sliding footsteps are created if there is a *moved* IK constraint occurring in the footstep interval. In a footstep animation this means that the foot can be placed anywhere, even though there is a footstep icon. Footstep icons can be thought of as gizmos rather than the absolute location of a foot. Sliding keys have an IK blend of 1, and are in object space, but are not joined to the previous IK key.
- **Free keys** have an IK blend of 0, and are in body space. They are not joined to the previous IK key. Free keys have no IK constraint.

IK constraints are implemented with a pivot-based system. This allows you to pivot a hand or foot around a selected pivot. For example, in a walking motion you can select a pivot on the heel of a foot and rotate the foot around it,. You can then shift the pivot to the ball of the foot.

Procedures

To create a purely freeform animation:

1. Create a biped.
2. Move the time slider to any given frame then do one of the following:
   - 

Turn on Auto Key, then move or rotate any biped component. This sets a key for that component.
   -  Pose any part of the biped, then in the Key Info rollout, click Set Key .
   -  Pose the hands and feet and then in the Key Info rollout click Set Planted Key . This will create keys that cause the hands or feet to stick in space.

     For more information on setting keys, see *Setting Keys for Feet and Hands (page 2–904)*.

   Note: **character studio** doesn't behave the same as 3ds Max when creating keys. It doesn't automatically create keys at frame 0 when you create your first set of keys at a later frame. For this reason, it's often useful to select all the parts of the biped at frame 0, then click Set Key to creating holding keys for all.

To create a freefrom animation from a footstep animation:

1. Select the biped whose footstep animation you want to convert to freeform.
2.  On the Motion panel Biped rollout, click Convert.

A Convert To Freeform dialog is displayed. Click OK.

Don't turn on Generate A Keyframe Per Frame unless you have a good reason to do so. Having a keyframe on every frame makes animation editing much more complex. Use this only if you are losing motion quality through a conversion process.

Creating Freeform Animation

The topics in this section deal with the creation of freeform animation for bipeds. They are:

Selecting Biped Tracks (page 2–888)

Animating by Moving Links (page 2–890)

Animating by Rotating Links (page 2–891)

Rotating Multiple Biped Links (page 2–895)

Using Controllers (page 2–896)

Using Props (page 2–898)

Selecting Biped Tracks

To animate your character with freeform methods, you need to know how to select the body part you want to animate, as well as the type of movement you want to affect for that part of the body.3ds Max and Biped provide a number of different methods for selecting and moving these animation tracks. Several involve using the 3ds Max Track View, a powerful environment for viewing and managing the geometry and motion data in your scene.

For each biped body part, motion data is viewable in Track View or on the track bar. Once you've selected the biped object, using one of the methods described below, you can see its associated motion data on the track bar or displayed in the Transform branch for that object in Track View.

Biped lets you expand and collapse certain animation tracks to give you more control over your character's movement as you set keyframes. For example, tracks for body parts in the arms can be animated using five separate tracks for maximum control.

You can also collapse these tracks for simplicity, and use a single key to pose the entire arm. Tracks

can be expanded or collapsed in this way for the arms, legs, ponytails, neck, tail, and spine.

In 3ds Max, you have the ability to display the motion data as function curves. You can see and manipulate function curves either in the Track View – Curve Editor, or by using *the Workbench (page 2–1012)*. The Workbench provides specialized tools for analyzing and fixing motion problems found in the curves.

For fast track selection, you can also use the Track Selection rollout. These buttons let you quickly select the motion tracks for the horizontal and vertical movement of the biped center of mass, as well as selection of opposite limbs, or symmetrical limbs. Biped is unique in the way it separates the tracks for the center of mass into three tracks (one each for vertical, horizontal, and rotation).

See also

Dynamics & Adaptation Rollout (page 2–980)

Track Selection Rollout (page 2–945)

Track View (Biped) (page 2–1002)

Procedures

To access Track View, do one of the following:

- On the 3ds Max toolbar click Curve Editor.
- Select any biped object in the viewport, then right-click and choose Curve Editor.

  The hierarchy list automatically displays the track of the selected biped object.
- Choose Graph Editors menu > Track View - Curve Editor.

To select from the screen:

1. Place the cursor over the biped link whose tracks you want to examine. Use zoom controls, as necessary.
2. Pause over a part of the biped, and a tool-tip appears, informing you which object is at the cursor position.
3. Click to select the biped link.

   Tip: Another feedback device is available once you've made your selection. The selected object's name is displayed in the topmost field in the Modify, Hierarchy, Motion, and Display panels.
4.  On the Track View, click Zoom Selected Object.

   The Track View hierarchy window repositions to show the selected object at the top of the window display. Nested below the object's name are the animation, or transform, tracks for that object, if any exist for the current motion.

   Note: By default, all keys for arm objects are in the Clavicle track, and all keys for leg objects are in the Thigh track. If you select a biped foot and click Zoom Selected Object, you may need to scroll up the Track View hierarchy to view the Thigh track transform keys. These keys also store foot transforms.

   Tip: Use controls in the Separate Tracks area of the *Keyframing Tools rollout (page 2–962)* if you prefer transform tracks for individual objects.

To display footstep tracks:

1. Select any biped with footstep animation.
2.  In the Biped rollout of the Motion panel turn on Footstep Mode .
3. Right-click in the viewport and choose Curve Editor.
4. On the Track View menu bar, choose Modes > Dope Sheet.
5. Click any footstep in the viewport and the hierarchy list navigates to the Footstep track.

Selecting Tracks from the Track Selection Rollout

The center of mass has three separate animation tracks, two for motion and one for rotation. These tracks can be selected on the *Track Selection rollout (page 2–945)*:

-  Body Horizontal
-  Body Vertical
-  Body Rotation (selects the Body Turning track)

Selecting Tracks with the Select By Name Tool

You can navigate to and select tracks for a given object using the text-entry field next to the Zoom Selected Object in Track View. Type the object name in the text field, using wildcards as necessary, and then press Enter.

Selecting Tracks from the Track View

You can navigate to a given object simply by traversing the Biped hierarchy in the Track View hierarchy window.

Note that child objects in the hierarchy are nested below parent objects. You may need to open several parent objects to get to the nested object you want. The biped tracks are grouped as follows:

- Center of Mass
- Footsteps
- Pelvis (branch1) Spine, Neck (branch), Head, Ponytail.
- Pelvis (branch2) Thigh, Calf, Foot, Toes.
- Pelvis (branch3) Tail.
- Neck (branch 1), Head
- Neck (branch 2), Clavicle, UpperArm, Forearm, Hand, Fingers.

Each of these tracks can be seen in Track View, expanded trackbar and Workbench.

Note that the hands and fingers are grouped with the arms, and the fingers and toes are grouped with the legs.

The body's translation keys are separated into vertical and horizontal tracks.

Animating by Moving Links

 Use the standard 3ds Max Move transform to move biped links.

As shown in the table below, there are two types of movements you can apply to a biped link:

- **General Move –** When you select and move the center of mass object, it and all of its children move; that is, the entire biped moves. When you move the center of mass and one or more legs are planted on the ground with IK constraints (IK Blend=1.0), the Biped tries to maintain the legs' planted position while the body moves.
- **Inverse Kinematics –** Applies to all parts of the arms, legs, fingers, and toes, except the clavicles. If you attempt to move a hand or foot beyond the biped's kinematic limit, the arm or leg straightens out and moves as far as possible in the direction you drag.

The pelvis, head, neck, non-base spine links, and non-base tail links cannot be moved. They can only be rotated, although they move when their parent moves.

When you select multiple biped parts and move them, the move applies only to selected biped parts that have no selected ancestors. For example, when the entire biped is selected, only its center of mass object is moved. When all of the finger joints are selected, the bases of all the fingers are moved.

Tip: You can select the children of a hierarchy by double-clicking the parent.

To quickly select and move a biped and its entire animation to a new location in space, use *Move All mode (page 2–936)*.

Biped provides controls to help you give both arms or both legs the same pose. See *Copy/Paste Rollout (page 2–966)* for more information.

Biped Link	Link Name	Keyframe Mode	Figure Mode
Center of mass	Bip01 (default)	Positions entire biped	Positions entire biped
Pelvis	Pelvis	Does not move	Does not move
Head	Head	Does not move	Does not move
Neck	Neck	Does not move	Does not move
Spine Base	Spine	Does not move	Positions upper body relative to lower body
Tail Base	Tail	Does not move	Positions tail relative to body
Ponytail	Ponytail1, 11, and so on	Does not move	Positions ponytail relative to body
Clavicles	R Clavicle, L Clavicle	Do not move	Positions arms relative to body
Shoulders (upper arm)	R UpperArm, L UpperArm	IK from shoulder	IK from shoulder
Elbows (lower arm)	R Forearm, L Forearm	IK from hand	IK from hand
Hands	R Hand, L Hand	IK from hand	IK from hand
Fingers	Finger 0, 01, and so on	IK from finger	Positions finger relative to hand (if base selected)
Hips (upper leg)	R Thigh, L Thigh	IK from knee	IK from knee
Knees (lower leg)	R Calf, L Calf	IK from foot	IK from foot
Feet	R Foot, L Foot	IK from foot	IK from foot
Toes	Toe 0, 01, and so on	IK from toe	Positions toe relative to foot (if base selected)

Animating by Rotating Links

 Use the standard 3ds Max Rotate transform to adjust a biped's posture by rotating its links. Use the main toolbar button, or right-click and choose Rotate from the quad menu.

To make biped movement appear natural, certain biped joints are limited in how they can rotate, such as the elbows and knees. When a joint can rotate in all three axes, X, Y, and Z, it is said to have three degrees of freedom (DOF). A joint's DOF can be modified by selecting it and then setting options on the *Locks (page 2–500)* rollout of the Hierarchy panel.

Note: Changing a joint's DOF will not replace its default DOF, meaning a locked axis by default can not be unlocked manually.

The table below shows the degrees of freedom of each link. Biped allows a bit greater freedom than most human bodies are capable of. All rotations are performed in the local coordinate system.

Tip: First select the link to rotate, then click Rotate. While the Motion panel is active, Biped sets the transform managers to reflect the degrees of freedom of that link. For example, if you select the pelvis and then click Rotate, Biped changes the coordinate system to Local and turns on the Y axis constraint. (For joints that have more than one

degree of freedom, you might later need to change the axis constraint setting.)

To rotate the biped pelvis in all three dimensions, refer to *Pelvis as Ball Joint (page 2–892)*.

Biped Motion Constraints

Some biped parts have special-case conditions that govern how you can transform them, as described in the table and sections that follow.

Biped Link	Link Name	Free Axes	Comments
Center of mass	Bip01 (default)	X, Y, Z	Rotates entire biped
Pelvis	Pelvis	X, Y, Z	If feet are planted, adjusts legs to keep feet and toes above ground
Head	Head	X, Y, Z	
Neck	Neck	X, Y, Z	Neck orientation does not affect head orientation
Spine	Spine, Spine1-4	X, Y, Z	Spine rotation adjusts overall balance
Tail	Tail, Tail1-4	X, Y, Z	
Clavicles	R Arm, L Arm	Y, Z	Shoulder orientation does not affect clavicle orientation
Shoulders (upper arm)	R UpperArm, L UpperArm	X, Y, Z	Rotating pivots from shoulder to wrist
Elbows (lower arm)	R Forearm, L Forearm	X, Z	Hinge plus special rotation
Hips (upper leg)	R Thigh, L Thigh	X, Y, Z	Rotating pivots from hip to ankle
Knees (lower leg)	R Calf, L Calf	X, Z	Hinge plus special rotation
Hands	R Hand, L Hand	X, Y, Z	
Feet	R Foot, L Foot	X, Y, Z	If feet are planted, adjusts legs to keep feet and toes above ground
Fingers	Finger0, 01, 02 Finger1, 11, 12, amd so on	X, Y, Z	Finger bases have three free axes; other finger joints have Z only
Toes	Toe0, 01, 02 Toe1, 11, 12, and so on	X, Y, Z	Toe bases have three free axes; other toe joints have Z only

Pelvis as Ball Joint

The pelvis can be rotated in all three axes: X, Y, and Z. This ball joint provides three degrees of freedom, improving the flexibility of animating with a biped, because the pelvis gyrations are isolated from the center of mass (COM) rotation. This causes the COM to assume a smoother trajectory, which also influences the biped's head and arms. As a result, they will look more natural moving in the body's COM rotational space.

It's important to be aware of the following behavioral changes resulting from this:

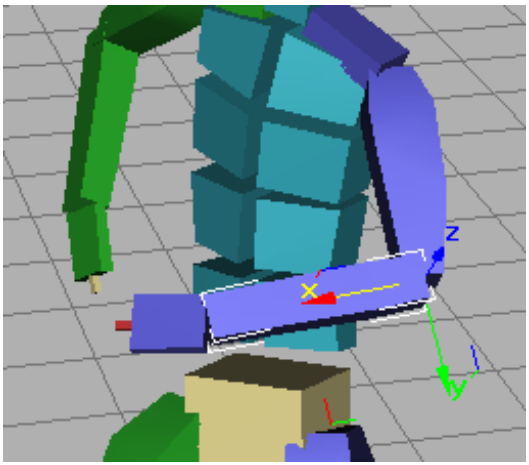
- The three-DOF pelvis motion is displayed as three separate Euler or quaternion (TCB) function curves in the Curve Editor and Workbench.
- Using key reduction on the biped's spine produces improved results.

- Rotations on the new pelvic axes (X and Z) are keyable in layers and supported in MAXScript.
- The default parametric footstep animation is slightly adjusted to use the three-DOF pelvis.
- If a legacy biped asset such as a BIP file is loaded onto the 3ds Max biped, the pelvis-related data will be adapted to the new DOFs and the new coordinate space within which the pelvis rotates.
- Using the Motion Mixer or Motion Flow, assets containing new pelvic rotation tracks will be blended with old assets lacking these tracks.

Special Rotation: Elbows and Knees

Elbows and Knees perform a special rotation when you rotate them about their X axis. They don't actually rotate around their X axis; this does not make sense because they have one degree of freedom. Instead, the upper and lower arm/leg are rotated together about an invisible axis defined by the line stretching from the shoulder to the wrist, and the hip to the ankle. This special rotation can be very useful for positioning the arms and legs.

The special rotation can also be useful for creating characters with reverse knee bends. When the knees are rotated backward, at more than a 90-degree rotation from the front-facing human knee posture, Biped assumes the character has backward knees or bird legs, and uses this as a reference position for all *.bip* motions.



Rotating the forearm along the X axis rotates the arm elements about an invisible axis between the shoulder and wrist.

Balance: Spine

Biped uses only the spine, in conjunction with the center of mass, to maintain the biped's balance. Because of this, rotating all of the spine or any one of its links causes the horizontal position of the body to change relative to its center of mass. These adjustments are performed in the center of mass's local reference coordinate system, ensuring that the figure will rotate naturally about its center of mass; for example, during flips in the air.

On the Motion panel *Key Info rollout (page 2–954)*, you can turn this behavior off by setting Balance Factor to 0.0 for corresponding horizontal center of mass keys. The Balance Factor control is in the Body group on the rollout.

Independent Orientation: Arms, Head, Feet

Changing the orientation of a clavicle (the root Arm object) changes the position, but not the orientation, of its corresponding upper arm. In effect, the clavicles are a support from which the arms are suspended.

Likewise, the orientation of the neck changes the position, but not the orientation, of the head. Although linked to the neck, the head typically rotates independently of the neck, and interpolation of these individually set orientations produces more natural-looking motion.

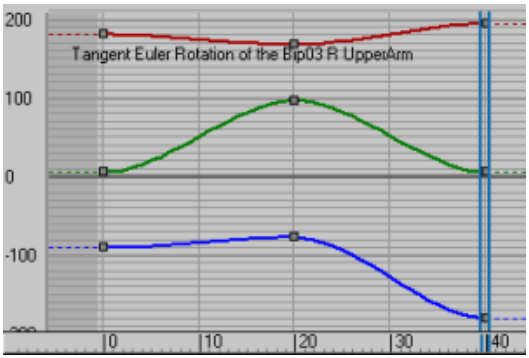
Similar to the head and arms, changing the orientation of the upper or lower leg changes the position, but not the orientation, of the corresponding foot. In this way, the foot orientation remains relative to the ground plane.

Adjusting Keys with TCB Rotation

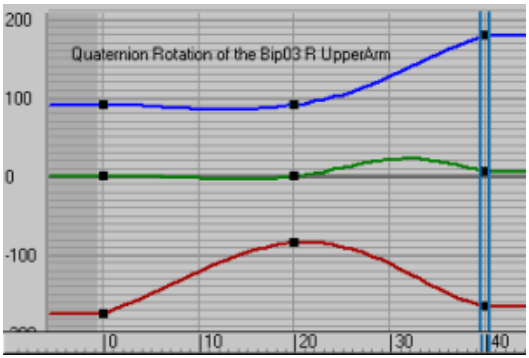
Rather than creating extra keys to fine-tune the motion of the biped limbs, you can use the TCB controls to adjust ease in, ease out, and limb trajectory on keys that already exist.

Visualizing Rotation Animations with Function Curves

Another way to visualize your rotation animation is through the *Curve Editor (page 2–507)*. Each key you add is displayed and connected to other keys, creating a curve that represents your animation. You can use either the *TCB Rotation Controller (page 2–377)* or the *Euler XYZ Controller (page 2–318)* (on the *Quaternion/Euler rollout (page 2–948)*) to display your rotation curve as Quaternion or Euler, respectively. Each controller affects the curve differently based on separate rotation calculations. To learn more about this, refer to *Working with Euler Curves on Biped Animation (page 2–916)*.



Tangent Euler Rotation curve



Quaternion Rotation curve

Procedures

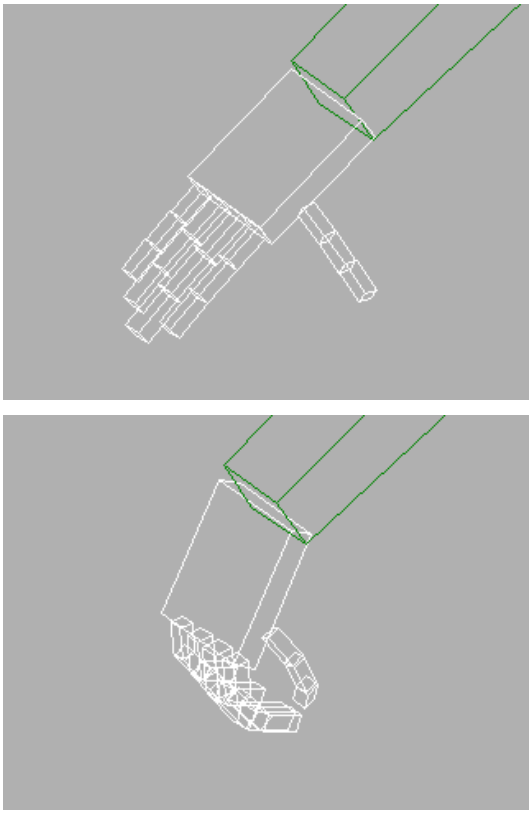
To change TCB for a biped arm:

1. On the Key Info rollout, turn on Trajectories.
2. Select a biped arm.
3. Use Next or Previous key to locate an arm key.
4. If the TCB group is not already displayed, click the TCB divider.
5. Change the Tension, Continuity, and Bias spinners.

   The trajectory changes to reflect the new parameters. Play the animation to see the change.

Tip: You can also access TCB controls by right-clicking on keys in the trackbar, Track View – Curve Editor or the Workbench Curve View.

Rotating Multiple Biped Links



Selecting and rotating a hand and all its fingers causes the fingers to curl.

You rotate multiple biped links to produce curling effects such as fingers curling around a glass or a tail curling up and down. You can choose from two methods to rotate multiple links:

- Select and rotate multiple links manually.
-  Enable *Bend Links Mode (page 2–952)* and then select and rotate any spine, neck, or tail link to use the **character studio** technique of naturally bending the entire spine, neck, or tail.

Selecting and Rotating Multiple Links

When you select and rotate multiple biped links, the rotation is individually applied to each selected link. This is a convenient way to get fingers to curl, for example, or to keyframe a biped's arms, legs, multiple-jointed neck, or tail.

To select and rotate multiple links:

- Select and rotate any number of links.

  Typically, you select an object and all of its children; for example, the hand and all of its fingers.

Using Bend Links Mode and Twist Links Mode



You can use *Bend Links Mode (page 2–952)* to rotate multiple links for the biped's spine, neck, or tail. Bend Links transfers the rotation of one link to the other links in a natural way. When applied to the spine, it is particularly useful for positioning the biped's hips.

You can also use Twist Links Mode on the Bend Links rollout to twist multiple links along the X axis while preserving their relationship with the two other axes. Combining these two modes opens up a wider range of controls over your spine, neck or tail.

To rotate all links in the spine, neck, or tail:

1.  On the Bend Links rollout, turn on Bend Links Mode.

2. Select and rotate a single spine, neck, or tail link. The other links in the spine, neck, or tail rotate to match the single link's rotation.
3.  On the Bend Links rollout, turn on Twist Links Mode.
4. Rotate the link in local X. The other links in the spine, neck, or tail twist properly to match the single link's rotation in X.

Tip: You can use Bend Links mode either to pose the biped or to animate the spine, neck, or tail while in *Keyframe mode (page 3–960)*.

Tip: You can also use the other Twist Links tools in the *Bend Links rollout (page 2–952)* to further control the links' rotation.

Using Controllers

character studio lets you add 3ds Max controllers to the biped object tracks. This allows you to mix the biped animation with standard 3ds Max animation techniques.

Why Add Controllers

Adding controllers to biped body parts lets you create animation that is difficult to do with **character studio** alone. For example, you can create "stretchy" cartoon arm or leg animation by adding a Scale controller to animate the bipeds arms and legs. You could make a character shaking by using a Noise controller for its limb rotations or for a Scale controller on the spine objects to make a character breathing. An Audio controller could be used to synchronize movement with sound.

Biped SubAnims

Different controllers can be added together, blended or sequenced in a list called a Biped SubAnim. These are the same as the weighted list controllers in 3ds Max. By animating the controller weights you can increase or decrease the effects of the different controllers. You can enable, disable or collapse the list controller animation tracks. If you collapse the track, you can then bring them into the Motion Mixer, Motion Flow, or Layer editor or export them to a game engine.

Note: Biped SubAnims exhibits different behaviors based on whether you assign a controller or a constraint as the SubAnim. A SubAnim with a controller assigned will layer the controller animation with the existing biped animation. If you use a constraint, however, this will *replace* the existing Biped Animation.

Tip: SubAnims are best used when combined with freeform animation. You can use them with footstep animations but you may experience unexpected results. If this happens, simply convert the footstep animation to freeform before applying the SubAnims.

Assigning SubAnims

Biped SubAnims can be applied several different ways. You can assign controllers to the Biped SubAnim in the *Assign Controller (page 2–934)* rollout of the Motion panel.

Controllers can also be added in the *Animation Workbench (page 2–1012)* using the Controller button to display the controller list. Right-click and choose Assign Controller from the right-click menu. You can assign controllers to multiple objects by using the SubAnims filter.

Procedures

To add a controller to a biped object:

1. Select the biped object that you want to control, and set a key for that object.
2. In the Motion panel, expand the Assign Controller rollout.

   The Biped SubAnim entry is displayed in the list window.

3. Expand the Biped SubAnim entry in the Controller window.

   List controller entries are displayed for BipScale, BipRotation, and BipPosition.

4. Expand the list controller you want to add to and highlight the first Available entry.

   The Assign controller button becomes active.

5.  Click the Assign Controller button.

   An Assign Controller dialog appears.

6. Choose the controller to want to use, and click OK.

7. Play the animation and observe the effect.

   You can animate the parameters for the controllers using Track View or by right-clicking the list entry in the controller window and then choosing Properties.

   Controllers can also be assigned in the Animation Workbench by displaying the Controllers list, choosing an Available entry, right-clicking and choosing Assign Controller.

To collapse a controller:

You can collapse a controller to use the animation in the Motion Mixer, or with Layers, or to export to a games engine.

1. Select the biped object with the controller, for example an arm with an added noise rotation controller.

2. On the Motion panel, in the Assign Controller rollout, click the Biped SubAnim controller you want to collapse.

3. Right-click the Biped SubAnim and choose Properties.

   A dialog appears with the Collapse choices.

   If a different dialog appears, you don't have the correct selection in the list window. Be sure you've selected the Biped SubAnim entry and not something else.

4. Turn on the tracks you want to collapse in the Enable group.

5. Turn on the options you want to use in the Collapse group.

6. To perform the collapse, click the Collapse button.

   The controller tracks are "baked" into the Biped tracks. The new keys will appear in the track bar, or in the Workbench Curve View if the Workbench is open.

To add controllers to multiple biped parts at the same time:

1. In the viewport select the biped parts to which you want to add list controllers.

2. Open the Workbench.

   The selection list in the Select panel has the biped parts highlighted.

3. 

Click the Controllers button on the Workbench toolbar.

   The controller list appears to the right of the select tab.

4. Select the tracks you want to affect.

5. On the Filter panel, locate the part to Filter group. Set it to Selected Parts if it isn't already.

6. Choose a SubAnims filter from the drop-down list.

7. Adjust any other options you need to in the Filters tab, then click Position, Rotation, or Scale to add the subanims.

8. Expand the controller list to observe the newly added controllers.

To animate the weights of an added controller:

You can animate the weights of an added controller to determine when the effect takes place. A weight of 0 will yield no effect, a weight of 100 will provide the full effect. See the tutorial "Using Controllers

with Biped" for an example of animating the weights.

1. Add a controller to a biped object, using either method outlined above.
2. Immediately after adding the controller select the appropriate BipScaleList, BipRotationList, or BipPositionList entry in the controller list, and expand it to see the Weights.
3. Expand the Weights and select Weight 0.
4.  In the Keyframing Tools rollout on the Motion panel, click the Manipulate SubAnims button.

   Three new rollouts appear in the Motion panel: Position List, Scale List and Rotation List.
5. Navigate to the appropriate List rollout. In the Layer window highlight the Controller you wish to weight.

   The Weight field becomes available
6. 

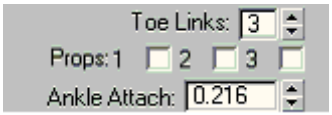
Turn on Auto Key, move the time slider and enter values in that field, or use the spinner to set keys to animate the weights.

   Warning: Don't even think about using the 3ds Max Set Key button to animate the weights. Avoid using this button with **character studio** as a rule. Use Auto Key or use the Key Info rollout set key buttons instead.

Using Props

Props are objects that a biped could be holding or carrying, such as a gun, sword, or briefcase. They are represented as boxes that are additional links in the biped hierarchy, an extension of the biped structure you can access at creation time or in Figure mode. If you are loading motion-capture data that contains prop information, this will transfer to the biped prop.

Each biped can have up to three props. Prop data will work seamlessly with other advanced features of **character studio**: in the Mixer, Motion Flow, Unification, and Layered edits.



Props are defined in the Structure rollout.

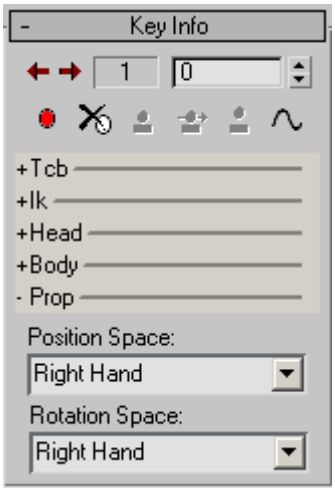
Props appear as boxes in the biped hierarchy. They can be moved, stretched, and scaled in Figure mode as other biped parts can. Props are supported by Physique, and are animated using the biped Transform controller in the same way as other objects. Create your prop and position it in relation to the biped in Figure mode. Props can be made to follow the motion of either hand by keyframing the Position and Rotation coordinate space of the prop.



Bipeds with multiple props as swords

Animating Prop Linkage

Props can change linked parents (similar to the 3ds Max Link controller) at any keyframe. This is done using the Position Space and Coordinate Space lists in the Prop section of the *Key Info rollout (page 2–954)*. Once a key is set on the prop at a given frame, you can set or change the coordinate space in which it is transformed by choosing a new space from the drop-down list. You set rotation and position coordinate space independently.



Position and rotation space controls

Importing Motion-Capture Prop Data

If your motion-capture file includes prop data, **character studio** can read that data and apply it to a prop object.

Upon loading a *CSM file (page 3–925)*, the motion data will be applied to the prop. If the prop switches from one hand to another, the program will automatically determine the switch of the prop's coordinate spaces as the prop is exchanged between the hands. This can be easily changed once the data is in the scene: go to the desired frame, choose a new coordinate space (as described in the previous section), then set a key. When loading a BIP file for retargeting onto other characters, the props move with each of the hands. The software ensures that the hands come together at the precise position on the prop during transfers.

Props can be animated with 3ds Max controllers and that animation can then be collapsed onto the prop's transform controller for editing with motion flow or the mixer.

Procedures

Example: To animate a biped swinging a prop with one hand:

1. Create a biped with a single prop.

   The biped appears with a single prop displayed as a box positioned near the biped's right hand.

2.  On the Motion panel, on the Biped rollout, turn on Figure Mode.
3. Scale the prop so it is the correct size and proportion to be a sword.
4. Rotate and move the prop so it is in the correct orientation in relation to the right hand.
5. Turn off Figure mode.
6. Set a key for the prop at frame 0. To do this, make sure the prop is selected, then use the Set Key button in the Key Info rollout.
7. Now animate the biped as you like, or open a BIP file of a character swinging a sword.
8. Turn on Auto Key, then select the prop.
9. Animate the prop using standard 3ds Max transforms.

Example: To animate the biped switching the prop to the other hand:

1. Follow the steps in the previous procedure to animate the biped swinging the prop.
2. Move to the frame where you want the biped to switch the prop the other hand.
3. Move the biped's other hand so it is in contact with the prop, and set a key for the hand.
4. Select the prop and set a key for it.
5. Move the time slider ahead one frame and set another key for the prop.
6. On the Key Info rollout, expand the Prop section.

   The lists for Position and Rotation coordinate spaces should be available. Change the Position

from the right hand to the left hand, and click Set Key again.

The prop will now follow the movement of the left hand in the frames that follow.

Note: By default, the prop's rotation remains in Body space. Depending on your animation, you might also want to set keys so the Rotation space is Right Hand before the switch, and Left Hand after the switch.

To collapse the transforms of a prop:

1. Animate the prop however you like. You can add list controllers to the prop as you can to any other part of the biped.
2. When the animation is correct, select the prop.
3. In the Motion panel hierarchy list, select the prop transform list you want to collapse, and then right-click.

   A quad menu appears.
4. Choose Properties from the quad menu.
5. In the dialog that is displayed, make the appropriate choices, and then click Collapse.

   After a short delay, the track bar displays the keys that have been added.

   Tip: If you see a different Properties dialog that doesn't offer you the Collapse button, you are too low in the hierarchy. Select the next level up in the hierarchy and try again.

Freeform and IK

Using IK Keyframe Parameters

Biped's inverse kinematics solution has three parameters set at each key of the arm and leg keyframe tracks. As the limb moves through each key:

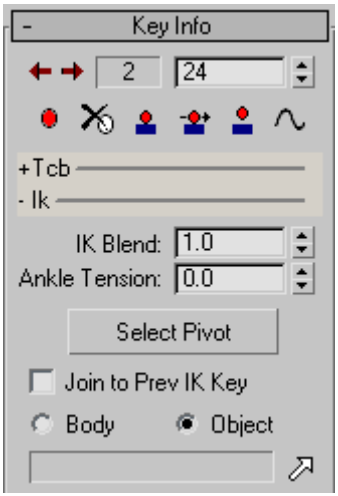
- **IK Blend** sets the motion interpolation to be a blend of forward and inverse kinematics. This allows you to blend swinging motions with hand-or-foot directed motions. Rotating an arm to move a hand is an example of forward kinematics. Using the position of the hand to move the arm is an example of inverse kinematics. The default blend is 0.0, or full forward kinematics. An IK Blend of 1.0 is full inverse kinematics.
- Body or Object determines the reference coordinate space of the IK path. This allows you to move the IK path with your character's body or temporarily attach the hands or feet to follow other objects. The default is Body.
- **Join To Prev IK Key** determines if the key should be part of the previous key (and have the same reference position as the previous key).

Biped has three automatic ways to set these parameters. Setting planted, sliding, or free keys creates keys with different combinations of these three parameters.

- Setting a planted key creates a key with IK Blend set to 1.0, Object and Join To Prev IK Key turned on.
- Setting a sliding key creates a key with IK Blend set to 1.0, Object turned on, and Join To Prev IK Key turned off.
- Setting a free key creates a key with IK Blend set to 0, and Body turned on. Join To Prev IK Key is turned off.

IK Blend

The IK Blend control is in the *Key Info rollout (page 2–954)*; visible when you expand the IK divider bar. You set IK Blend while in Keyframe mode .



IK group parameters

IK Blend can be set per key for any arm or leg track. The IK Blend setting determines whether, at a particular key, an arm or leg is moving through it using inverse kinematics, forward kinematics, or a blending of the two kinematic solutions.

An IK Blend value of 0.0 means full forward kinematics. The arm (or leg) is moved by interpolating the rotations of the joints at the keys. The hand (or foot) tends to move along sweeping circular arcs in this case, and the motion appears to be motivated by the apparent swinging at the joints.

An IK Blend value of 1.0 means full inverse kinematics with the hand (or foot) being used as an end-effector. A spline path is computed through the keys of the hand, and the hand moves along that spline. Joint angles for the rest of the arm are computed to allow the hand to follow the spline. The motion, in this case, appears to be directed by the hand (or foot).

An IK Blend value between 0.0 and 1.0 means a combination of inverse and forward kinematics; when IK Blend is closer to 0.0, forward kinematics are more heavily weighted in the solution, and when IK Blend is closer to 1.0, inverse kinematics are more heavily weighted.

It is best to use forward kinematics when you want the arms to swing, such as when a biped is walking. In the case of a boxer, however, since the hand should follow a directed path when punching, inverse kinematics should be used.

To set the IK Blend value of a key:

1. Select a single arm or leg track by selecting one or more parts of a biped's arm or leg.

   The IK Blend spinner and the other controls in the IK area are enabled only when a single arm or leg track is selected.
2. Set a key if one doesn't already exist.
3. Set the desired value of IK Blend.

Body and Object Options

By default, Biped calculates the kinematics solution using the coordinate system of the biped figure's center-of-mass, or the Body coordinate system. This means that the IK path of the hand (or foot) translates and rotates with your character as it moves. For example, the boxer's hand trajectory always moves relative to the weaving, bobbing, and turning of the boxer's body.

The Object option is used for animating dynamic links between the limbs and other objects in the scene.

The IK Blend control activates when a biped arm or leg (hand and foot) key is current.

- 0 with Body turned on is *forward kinematics (page 3–944)*, or normal biped space.
- 1 with Body turned on is *inverse kinematics (page 3–958)*, creates more straight-line motion between biped keys.
- 1 with Object turned on, but no IK object specified, puts the limb fully into world space.

Use this to control foot sliding in a freeform animation.

- 1 with Object turned on and an IK object specified puts the biped limb into the coordinate space of the selected object; the biped limb follows the specified object.

Join To Prev IK Key

This IK constraint is used to specify if a footstep is sliding or planted. If Join To Prev IK Key is on, then the biped foot maintains a reference position to the previous key, keeping the foot planted. If Join To Prev IK Key is off during a footstep, then the foot can be moved to a new position creating a sliding footstep.

Join To Prev IK Key also functions to lock the biped hands in space. Use Set Planted Key on the biped hands to lock their position in space.

Tip: If you are having trouble with a foot or hand popping back to a previous 'keyed' position, check to make sure that the Join To Prev IK Key is off, especially when you're using pivot points.

See also

Key Info Rollout (page 2–954)

Understanding Walk Cycle Constraints

Walking motions in both freeform and footstep animations should follow the rules of certain IK constraints. In both freeform and footstep animation, a footstep interval is the start and end of a sequence of IK constraints in world space, with IK Blend set to be greater than 0.0. A biped foot in the Move state should have body space turned on with an IK Blend of 0.0. By using these IK constraints, you can convert between the two animation methods seamlessly.

In a freeform walking animation, typically you need two key types for the legs:

- If the foot is planted on the ground, the key must have IK Blend=1.0 with Object Space and Join To Prev Key turned *on*.
- If the foot is in the air between footsteps, then IK Blend=0.0 with Body Space turned on, and Join To Prev Key turned *off*.

In the course of creating a walk or run cycle, you need to alternate these IK constraints for the feet. If the foot is sliding on the ground, then the IK constraints are IK Blend=1.0 in Object Space with Join To Prev Key tuned off. All of these IK parameters can be found on the Key Info rollout.

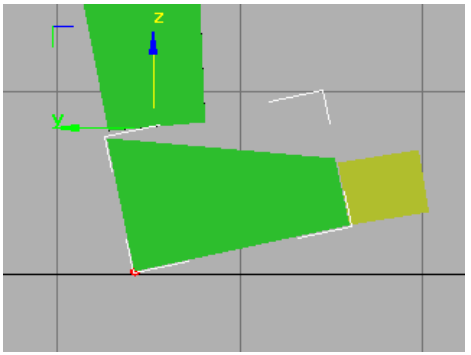
To speed up the process of applying these IK constraints, the Key Info rollout has three additional set key buttons: Set Planted Key, Set Sliding Key, and Set Free Key. By clicking one of these buttons, all the necessary IK constraints are applied automatically. For example, by clicking Set Planted Key, all of the IK constraints necessary are applied at once; IK Blend=1.0 with object space and Join To Prev Key turned on.

See also

Key Info Rollout (page 2–954)

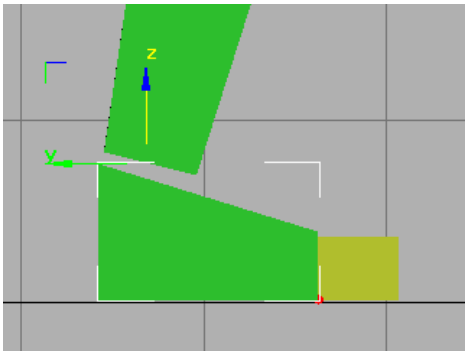
Walking Keys

Let's examine the IK constraints for one footstep. The footstep rolls on the heel of the foot, then rotates down flat on the ground, then raises up on the ball of the foot, rotates at the end of the toes, and finally lifts off of the ground.



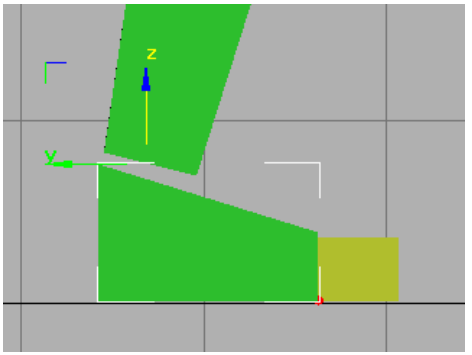
Touch state: pivot planted at heel at frame 14


The foot is touching the ground at the heel. In the Key Info rollout, Set Planted Key is clicked to set IK Blend=1.0 with object space and Join To Prev Key turned on. A pivot is selected on the heel of the foot.



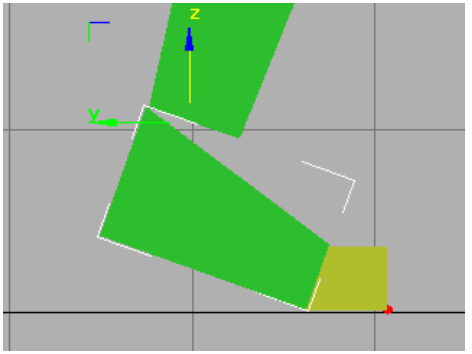
Pivot point shifts to ball of foot at frame 16.

The next keyframe is also a planted key, as the foot is flat on the ground. In the Key Info rollout, Set Planted Key is clicked. The pivot on the ball of the foot is selected.



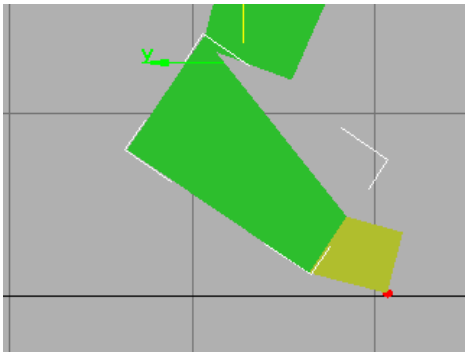
Planted key set to lock pivot at frame 17

This key has the pivot on the ball of the foot as well. Click Set Planted Key. Two consecutive keys with the pivot at the ball of the foot are necessary to rotate the foot about the ball of the foot.



Heel lifts, toe remains flat at frame 18

After rotating the foot about the pivot at the ball of the foot, you set another planted key with the pivot on the toe.



Foot rotates off pivot at end of toe at frame 19.

Here is another planted key with a pivot at the end of the toe; the foot rotates about the tip of the toe.



Free Key allows foot to travel off the ground at frame 20.


Here, Click Set Free Key to set a free key, the foot is off of the ground.

The cycle is repeated to create a walk or run cycle in a freeform animation. The pivot points are selected by turning on Select Pivot on the IK Key Info rollout, selecting a pivot in the viewports, turning off Select Pivot and then rotating the foot in the viewports.

When Join To Prev is turned on, as it is with all planted footstep keys, the foot pivots at the previous key's pivot point. If you find that the foot is not rotating around the visible pivot, remember that it is using the previous key's pivot. In such a case, you must set two consecutive keys with the pivot at the same location: then the foot will rotate around the displayed pivot.


If you adhere to these rules in creating footsteps in a freeform animation, then you can use Convert in the Biped rollout to easily change from a freeform animation to a footstep animation. Note that IK constraints in a footstep animation are applied exactly the same way as in a freeform animation. If you examine a footstep animation, you will notice that the foot IK constraints follow the same rules as for a freeform animation. Footsteps are simply gizmos that define a foot's coordinate system. A foot can slide and move relative to the footstep. Also, if you delete or add a key in a footstep animation, the footstep duration is changed.

Setting Keys for Feet and Hands

You can animate biped's body parts as you do other 3ds Max objects: by setting keys for postures at keyframes.

character studio provides three different ways to set keys at the current frame:

- 
Pose the biped body part, and then click one of the set key buttons in the Key Info rollout.
- 
Turn on the Auto Key button, and then pose the biped part. You can transform the parts using transform tools, or you can copy and paste poses and postures to keyframe the biped.
- Right-click the time slider and use the Create Key dialog to create keys for selected biped parts.

Warning: *Do not* use the large Set Key button, found below the viewports, when animating a biped. Doing so will cause unpredictable results. Instead, always use the Set Key buttons found on the Key Info rollout.


If you are animating a walk cycle or an intricate hand animation, then you should make use of the three different types of set key buttons in the Key Info rollout. Each set key button applies different IK constraints, depending on whether a foot or finger is in a planted state, a move state, or a sliding state. For more information on the three set key buttons, see *Key Info Rollout (page 2–954)*.

You can also set keys to make hands and feet follow a stationary or animated object. See *Animating IK Attachments (page 2–908)*.

If a Bones system using the IK controller or a particle emitter is linked to the biped, the Auto Key button must be on in order to position the biped's limbs. This is also true for biped objects if you turn on 3ds Max Trajectories at the top of the Motion panel, or if you choose Views > Show Ghosting.


The Set Key options on the Key Info rollout have the advantage that you can easily experiment with different poses for your character without unintentionally setting keys as a side-effect. Use the Set Key buttons to commit the changes you make to the pose. The Auto Key mode approach is especially useful when you make adjustments to keyframes that have already been set, or if you are used to working with Auto Key in3ds Max and tend to forget to click the Set Key buttons.

Tip: Use Set Key to insert a key on a frame where a key doesn't exist. You'll often want to refine the motion using controls in the Key Info rollout, without selecting and moving an object in the viewport.

Keyboard shortcut: Pressing **0** (zero) is equivalent to clicking Set Key on the Key Info rollout.


To use **character studio**-specific shortcuts, make sure that the *Keyboard Shortcut Override Toggle (page 3–872)* is active.

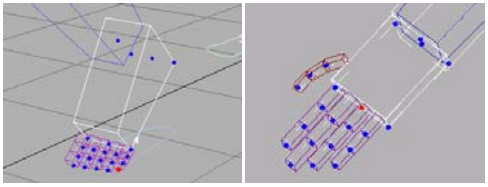
Ground Plane Collision Detection for the Pelvis and Feet

When one or two feet are planted, and the pelvis or a planted foot is rotated, **character studio** detects collisions of the foot and its toes with the ground plane defined by the planted foot's footstep or pivot point. As the pelvis or foot is rotated, **character studio** uses inverse kinematics to alter the rotations of the leg joints so that the foot and its toes do not go below the plane defined by the footstep or pivot point.

Note: A biped foot on a sliding footstep is a special case. The biped foot can be moved to any position and orientation for the duration of the footstep.

Note: Changing rotation values for the legs and toes also maintains collision detection, but the position adjustment is not made until you set a key.

Animating Pivots



In both freeform and footstep animation, pivots allow you to rotate the biped's hands and feet around various points. For example, by activating a pivot on the ball of a foot, you can rotate the foot around the ball of the foot. The biped's hands and feet have the same number of pivots, and pivot location is similar. Pivots are active only if the biped hand or foot is in world or object coordinate space. In a walking motion, you can pivot on the heel first, then on the ball of the foot, and finally

onthe toes. Pivots are essentially extensions of the IK chain.

By setting a planted key for the hand, the hand is anchored in world space: you can move the biped or the collarbone and the hand remains planted. Pivots on the hands make it easy to animate hands and fingers.

While pivots are used in both freeform and footstep animation, the pivots are only visible and accessible when in Freeform mode.

Biped IK

To understand interactive IK limb manipulation using **character studio**, it is useful to distinguish two types of limb joints: primary and secondary.

The primary joints are the shoulder, elbow, hip, and knees. Animals use primary joints to coordinate the positioning of hands and feet, as these joints have the most influence and flexibility in positioning tasks. Even the placement of fingers and toes over specific spots is the task of the primary joints.

The secondary joints are the wrists, ankles, toes, and fingers. These joints are typically used for grasping and support rather than positioning, so they most frequently exercise independent joint angle control suited to a particular task, such as rotating the foot to raise the heels during walking, or curling the fingers around an object. Because these joints have little influence on end effector position, they are rarely engaged for positioning tasks. For example, if you want to place your finger on your nose, you will most naturally rotate only your shoulder and elbow, keeping your wrist stationary.

character studio mimics the IK behavior of "natural" systems. The primary joints are used for IK positioning, while the secondary joints are independently and precisely controlled by the animator. The secondary joints will not rotate unless the animator explicitly selects and rotates them. Therefore, if an IK pivot is set, and the limb is interactively manipulated, (say, the finger is moved onto the nose), only the primary joints adjust to satisfy the IK constraint. The secondary joints remain in place and will never "drift" in an IK solution; they simply obey the values set by the animator. Because there is no drift for the secondary joints, the animator does not have to tediously assign "tension" values to hold them in place.

If an IK pivot point constraint has been set, some by-products of **character studio** "natural IK" are that:

- Interactively rotating a secondary joint will always adjust the primary joints to hold the IK pivot in place.
- **character studio** "special rotations," the forearm/calf X-axis rotations, rotate the elbow about the axis from the shoulder to the wrist and rotate the knee about the axis from the hip to the ankle. These rotations give the animator a direct way to interactively rotate the primary joints while holding the IK pivot in place.
- Interactively translating a limb part to move the IK pivot will alter only the primary joints.
- Interactively rotating a primary joint will move the IK pivot with it.

After an IK pivot has been placed, applying 1) and 2) gives the animator exact control over all possible IK solutions. The IK pivots can be easily moved by applying 3) and 4). In general, you don't need to remember these "rules," because the system works in an intuitive way. You merely need to position the pivots as desired, set the IK constraint, and then fine-tune the posture of the limb, if desired, by rotating the secondary joints and using the special rotation to adjust the primary joints.

Note: During animation playback, the Ankle Tension parameter is used to set the relative

importance of the interpolated ankle joint over the interpolated knee joint for intervals in between keyframes. This is relevant only to interpolation on bipeds.



Using IK constraints and pivots on the biped hands and feet allows you to animate quadrupeds.

See also

Key Info Rollout (page 2–954)

Procedure

To use pivots:

1. Select a biped hand or foot.
2.  On the Key Info rollout, click Set key, and then choose Object.

   You can also select an object if you like at this point, by selecting a non-biped object in the viewport.
3. Turn on Select Pivot.
4. Select a pivot in the viewports.

   Pivots are shown as red dots on the hands and feet. Use wireframe viewport shading if you have problems seeing the choices.
5. Turn off Select Pivot.
6. Rotate the hand or foot around the selected pivot.

   Tip: If you find that you are not rotating around the selected pivot, set two consecutive keys at adjacent keys with the same pivot.

Animating a Quadruped

While **character studio** was initially designed as an specialized program for animating two-legged, or bipedal characters, it works quite well for creatures that walk on four or more legs.

Here are some general rules to follow when animating characters that don't walk upright on two legs:

- Use Freeform animation, not Footstep animation, if your character walks on four legs all the time.
- Pose the biped to match the mesh in Figure mode, scaling and rotating body parts so the spine is horizontal, and the arms stretch to reach the ground.
- If appropriate for your mesh, set Leg Links to 4 and rotate the legs so the knees point backwards. You can set this value at creation time in the Create Biped rollout on the Create panel, or later using the *Structure rollout (page 2–984)* on the Motion panel.
- Animate pivot points with planted keys to mimic the rolling of the feet from heel to toe.

Treat the hands the same as the feet. Set planted keys on the hands and feet, then move the center of mass object to bend the knees and elbows. See *Key Info Rollout (page 2–954)* for more information on the three set key buttons.

- You can save3ds Max objects as part of the BIP file. If you need additional legs (for a centipede) or extra arms or wings you can use standard bone objects with IK chains, and save all of it with the BIP file.
- The Balance Factor in the Key Info rollout Body group, is designed to synchronize upper body and lower body movement. Set Balance Factor to 0 if you find the movement of spine links is affecting the biped's hip movements in an undesirable way.

Animating IK Attachments

A biped can interact with other objects in the 3ds Max scene. Links between objects are usually static, unless you're using the Link controller. In **character studio**, such attachments are "animatable" as well; during the course of an animation, the links between the hands, feet, and objects in the scene can change as your character interacts with them.

This capability is useful for:

- Creating freeform motions (without footsteps and gravity) that require that feet or hands be planted with IK, and then released. Examples are climbing a ladder, riding a bike, or rowing a boat.
- Motions that involve the temporary manipulation of objects, such as bouncing or kicking a ball, opening a door, or touching another biped.

An object to which a hand or foot can be attached is called an *IK object*. See the tutorial "Interacting with Objects" for lessons that show you how to make a biped iron clothing, ride a skateboard, bounce a basketball, climb a ladder, and pick up a briefcase.

Procedures

To make a hand or foot follow an object:

1. Create the object for the hand or foot to follow.
2. Select the biped's hand or foot.

   Tip: You can select any part of the biped's arm to set keys for the hand, or any part of the leg to set keys for the foot.
3. In the *Key Info rollout (page 2–954)*, expand the IK divider to see the IK section of the rollout.
4.  Click Set Key.

   Once a key is set, the IK Blend value and other IK parameters become available.
5.  Click Select IK Object, and choose the object for the hand or foot to follow.

   The object name appears in the field below the Body and Object radio buttons.
6.  On frames where you want the hand or foot to start following the object, set a key, set IK Blend to **1.0** and choose the Object option on the Key Info rollout. You can also click Set Sliding Key to set these parameters all at once.

   Note: If Object space is turned on with no IK object specified, the IK constraints will be in World space. In effect, the hand or foot will be attached to the world, and will not move during the time such a key is set.

   Set one key when you want the attachment to begin, and a second key when you want the attachment to stop. This defines an Object space interval, the duration of a temporary attachment.
7.  On the frame that defines the end of an attachment interval, set IK Blend to **1.0**, choose

the Object option, and turn on Join to Prev IK Key. You can also click Set Planted Key to set these parameters all at once.

8.  On frames where you don't want the hand or foot to follow the object, set a key, set IK Blend to **0.0** and choose the Body option. You can also click Set Free Key to set these parameters all at once.

To anchor a hand or foot:

You can use anchoring to temporarily look a hand or foot to a specific point in space while you are animating IK attachments.

1. Move and rotate the hand or foot into the desired posture.
2.  On the Keyframing Tools rollout, click the button for the limb you want to anchor: Anchor Right Arm, Anchor Left Arm, Anchor Right Leg, or Anchor Left Leg.

   Note: The arm or leg you select beforehand does not actually have to be the same as the arm or leg you are anchoring.
3. Set keys for the arm or leg as described in the previous procedure.

   Note: You can also use anchors to hold a foot or hand in position while you render the animation. However, be aware that anchors are turned off when you close the MAX file. If you want to use anchors from one session to the next, you must turn them on again the next time you open the file.

Editing Freeform Animation

The topics in this section deal with the editing of freeform animation for bipeds. They are:

Editing Biped Keys (page 2–909)

Copying and Pasting Postures and Poses (page 2–910)

Mirroring Motion (page 2–913)

Using Layers (page 2–913)

Editing Trajectory Keys (page 2–914)

Editing Biped Keys

 You can move back and forth between keys by clicking the Previous Key and Next Key buttons on the Key Info rollout. The fields to the right of these buttons indicate the key number and frame number.

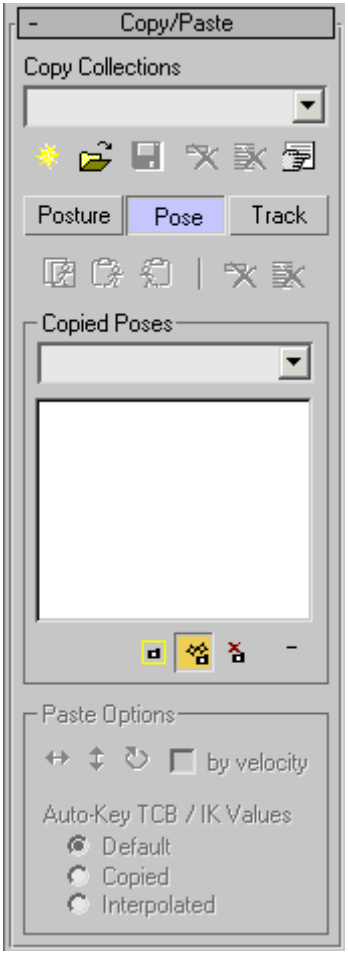
 As with other 3ds Max objects, you can also move back and forth between keys by turning on the 3ds Max Key Mode Toggle, selecting an element associated with a given track, and then using the 3ds Max buttons Next Frame and Previous Frame. For example, you can view keys, move between them, and set keys for a right arm track if any of the right arm's objects (clavicle, upper arm, lower arm, hand, fingers) are selected. Or use the **<** and **>** keys on the keyboard to move back and forth between keys in Key Mode.

 You can delete a key by clicking Delete Key in the Key Info rollout. The biped part that has the key must be selected, and the current frame must be the keyframe you want to delete. Keys that are locked (appearing in red in Track View) cannot be deleted.

Tip:  When you've selected the biped part you want to transform, click Lock Selection Set on the 3ds Max prompt line. Now you can transform the part without accidentally selecting a different part of the biped. The default keyboard shortcut for Lock Selection Set is the **spacebar**.

Copying and Pasting Postures and Poses

The Copy/Paste rollout on the Motion panel provides controls to help you copy and paste biped postures, poses, and tracks. A *posture (page 3–996)* is the rotation and position of any selection of biped objects. A *pose (page 3–996)*is the rotation and position of all the objects in a particular biped. A *track (page 3–1023)* is the animation for any selection of biped parts. For information on working with tracks, see *Copying and Pasting Tracks (page 2–926)*.



You can create and store postures, poses, and tracks in multiple buffers accessed by a drop-down list. You can choose from that list and see a thumbnail image associated with the selection, then paste that selection onto the same biped or any other biped in the scene. You can choose between posture, pose, or track for these operations.

Using these tools lets you work in a traditional pose-to-pose method of animating. You can create a variety of poses and postures, store them in the list, then in Auto Key mode, copy them to any biped at any frame to create animation.

Use the Copy Tracks function to copy the animation of biped parts onto other parts and other bipeds. You can further manipulate these tracks with the Curve Editor and Dope Sheet Editor, using all the standard tools found there. Copy Tracks works with both footstep and freeform animation.

Tip: The *copy collections (page 2–966)* feature acts as an additional layer of organization, allowing your copied postures, poses, and tracks to be grouped together so you can transfer them between files within a single session. After saving a collection, you can load it into another scene either by replacing the current collection or by appending yours to it. If you choose to append, the collection maintains the copy buffer structure.

Note: You must create a copy collection before you can copy a biped's posture, pose or track.

Using Paste Posture

The Paste Posture command is useful for copying a posture in one frame of an animation to another frame of the animation. Copy the posture, then turn on Auto Key and move to another frame, then paste the posture.

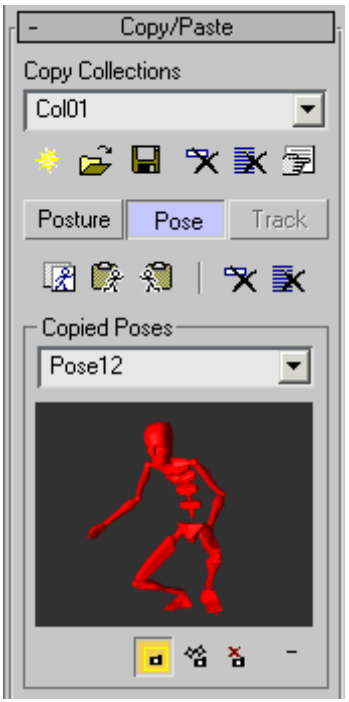
The Paste Posture and Paste Pose commands are also useful for copying a pose from one biped to another. Copy the pose or posture, select the other biped, and then paste.

Copying a posture then using Paste Posture Opposite is particularly useful for copying the posture of an arm or leg or a part of an arm or leg onto the opposite arm or leg of the same biped.

Both Paste Posture and Paste Posture Opposite work differently in and out of Figure mode. Out of Figure mode, only the orientation of the copied links is pasted. In Figure mode, both the orientation and the scale of the copied links are pasted. Also, when you paste the finger base, toe base, spine base, tail base, or clavicles in Figure mode, the position of that link relative to the biped's body is pasted.

Copying the Entire Biped Pose

The Copy Pose functionality allows you to copy the rotation and position of all the parts of the entire biped at once. It doesn't contain the animation track data, only the individual keys at that frame, and the buttons to copy tracks are unavailable when Copy Pose is in use.



Pose shows the whole biped

You can create animation using Copy Pose by simply copying different poses to the same biped at different frames, and setting keys for those poses. Copy Pose works with the **character studio** Set Key buttons found on the *Key Info rollout (page 2–954)*, and with Auto Key mode.

Warning: Don't use the Set Key Mode toggle or the large Set Keys button next to Auto Key. These will not produce the correct results.

Procedures

To copy a posture:

1. Select the set of biped parts that defines the part of the biped's posture you want to copy.
2.  On the Copy/Paste rollout, click Create Collection and rename the new collection **Biped Posture**.
3.  Click Posture to toggle Posture mode and then click Copy Posture. The thumbnail image appears in the image window below the Copied Postures drop-down list.
4. Rename the posture by editing the name in the Copied Postures name field. The posture of the set of selected objects in the biped is copied to the list. You can retrieve this posture and use it to paste (or paste opposite) to any biped at any frame.

For hands-on experience using the Copy/Paste rollout functions, see the lesson called "Creating a Simple Freeform Animation" found in the *Animating with Freeform* tutorial.

To paste the posture:

Note: Turning on the Auto Key button before pasting posture automatically sets a key for all objects whose posture has changed.

1. Choose the Posture you want from the copied postures drop-down list.
2. Select the biped you want to paste to, or move to another frame for the same biped.
3.  Click Paste Posture.

   The copied posture is pasted onto the selected biped. It doesn't matter which part of the biped is selected; all biped parts whose posture was saved by the most recent Copy are pasted.

   Warning: If you paste a posture while the Auto Key button is off, no key is set. If you want the pasted posture to set a key, you must click the **character studio** Set Key button in the Key Info rollout while objects are still selected.

To paste a limb posture on the opposite limb:

1. On the Copy/Paste rollout, click Posture.
2. Select all of the limb, arm, or leg whose posture you want to copy.
3.  Click Copy Posture.

   The posture thumbnail image is displayed in the image window, and the new posture name appears in the named posture field.
4. Rename the posture.

   The default names combine the chosen body track with the selected bodyparts. It's a good idea to replace this with a descriptive name you can easily identify with the action it represents.
5.  Click Paste Posture Opposite.

   The opposite limb now has the posture of the limb you copied.

To copy the entire pose of a biped:

1. Select any part of the biped whose pose you want to copy.
2.  On the Copy/Paste rollout, click Create Collection and rename the new collection **Biped Pose**.
3.  Click Copy Pose.
4. Advance to a new frame or select another biped.
5. 

Turn on Auto Key.
6.  Click Paste Pose or Paste Pose Opposite.

Tip: You can use Paste Pose Opposite to create symmetrical runs and walk cycles.

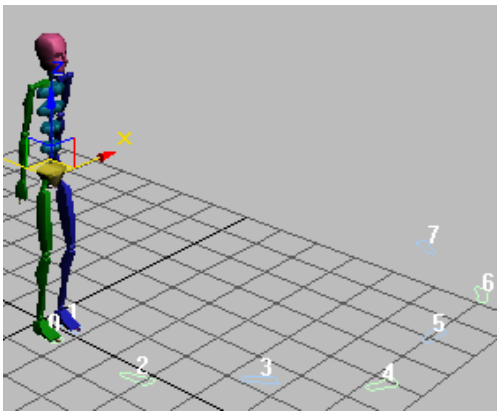
Keyboard Shortcuts

The following are Biped keyboard shortcuts to the copy and paste posture commands. Make sure the *Keyboard Shortcut Override Toggle (page 3–872)* is active.

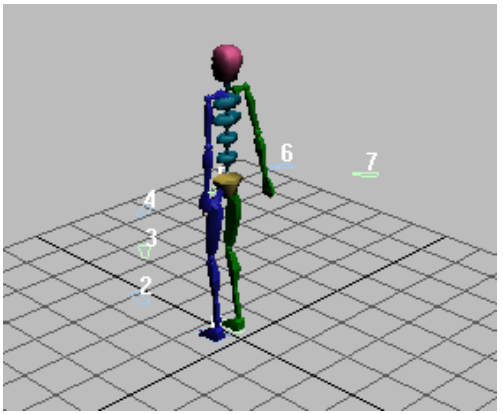
Shortcut	Meaning
Alt+C	Copy Posture
Alt+V	Paste Posture
Alt+B	Paste Posture Opposite

Mirroring Motion 

 Mirror, on the *Keyframing Tools rollout (page 2–962)*, mirrors the motion of the biped through both the X and Y axes of the World coordinate system. The entire biped animation, including all footsteps and keys, is mirrored symmetrically through an axis that joins the center for the biped to the world origin. For example, if a biped is walking from the center of the world grid *toward* the user, the mirror will change the motion so the biped is now walking *away from* the user.



Original motion



Mirrored motion

Use Mirror if you want to create the opposite of a motion you've already created. For example, if a biped walks to the left and swings its right arm, mirroring the motion results in a motion where the biped walks to the right and swings its left arm.

The mirrored motion replaces the existing motion for the biped. If you want to keep the existing motion and add the mirrored motion, use Paste Opposite in the Copy/Paste rollout.

Note: This feature mirrors only those biped tracks that have at least one key. The position of a keyless object is not mirrored.

Using Layers

Layers allow you to add successive layers of animation above the original biped animation. This is a powerful way of making global changes to your character animation. For example, simply add a layer, and rotate the spine forward at any frame, and a run cycle becomes a crouched run. The original biped motion is kept intact, and can be viewed by switching back to the original layer. Layers can be viewed individually, or as a composite of all the Layers. Layers behave like a *freeform animation (page 3–945)*: the biped can adopt any position.

Layers allow you to easily adjust raw motion capture data containing keys at every frame. Simply add a layer, and keyframe the biped. You can also use layers to change the global position of the biped in a freeform or footsteps animation by adding a layer and moving the center of mass.

Procedures

To increment all keys using layers (global offset):

1.  On the Motion panel Layers rollout, click Create Layer to create a new layer. Enter a name for the layer in the Layer name field.
2. Create an offset (or "increment") by setting one key. You can either use Auto Key mode, or click Set key on the *Key Info rollout (page 2–954)* (this can be done at any frame.)
3.  Click Collapse to collapse the layer.

To increment an interval of keys with an envelope (blended offset):

1. Create a new named layer.
2. Create an offset by setting one key. You can use either Auto Key mode, or click Set Key on the Key Info rollout at the frame where the increment is to occur.
3.  To create the envelope, use Snap Set Key to create keys before and after the incremented range.
4.  Click Collapse to collapse the layer.

To increment an interval of keys without an envelope (nonblended offset):

1. Create a new named layer.
2. Create an offset by setting two keys. You can use either Auto Key mode, or click Set Key on the Key Info rollout. Set one key before the frame where the increment starts, and another after the frame where the increment ends.
3.  Use Snap Set Key to create keys before and after the incremented range to clamp the motion to the offset only occurs in the range between the two keys.
4.  Click Collapse to collapse the layer.

See also

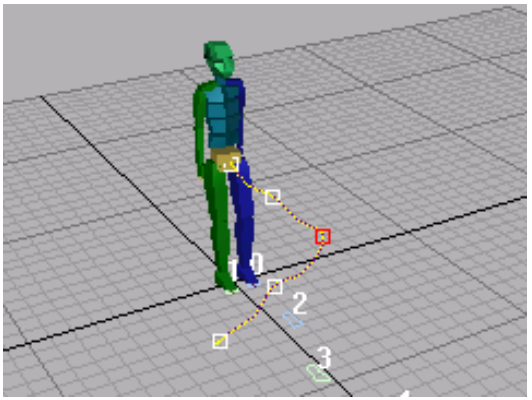
Layers (page 3–961)

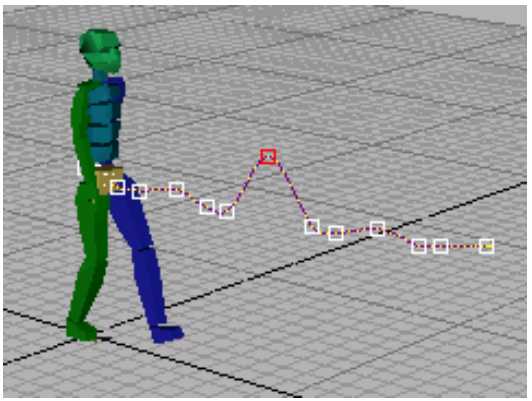
Layers Rollout (page 2–974)

Editing Trajectory Keys

By turning on biped trajectories, selecting the center-of-mass track, and turning on Sub-Object: Trajectories, you can select a key on the trajectory. Then either use Bend Horizontal in the *Keyframing Tools rollout (page 2–962)* to bend the trajectory about the selected key, or simply move the key. You can move center of mass trajectory keys either horizontally or vertically. This allows you to view the entire trajectory of the center of mass and bend or edit the trajectory.

Tip: You can select the center of mass's horizontal or vertical tracks by clicking the appropriate button in the *Track Selection rollout (page 2–945)*.





Above: Bending the trajectory

Below: Moving a key on the trajectory

Warning: Do not use 3ds Max trajectories (Motion panel > Trajectories text button) to edit biped trajectories.

Procedures

To bend a biped trajectory:

1. Select an animated biped.
2. On the Biped rollout of the Motion panel, expand the horizontal bar (click the plus sign at the left) to display the Modes and Display groups.
3.  In the Display group, turn on Trajectories.

   The trajectory is displayed in the viewport for whatever biped part is selected.

   Warning: Don't click the Trajectories button next to Parameters. That is for other scene objects, not bipeds.
4.  In the Track Selection rollout, click Body Horizontal.

   The center of mass object trajectory is displayed.
5. At the top of the panel, below Selection Level, turn on Sub-Object, and choose Trajectories on the drop-down list if it isn't already displayed.
6. In a viewport, select as many keys as you want on the trajectory.

   The selected keys turn red in the viewport.
7. If you want to bend the trajectory open the Keyframing Tools rollout and locate the Bend Horizontal spinner. Then adjust the Bend Horizontal values as desired.

To edit keys on the Center of Mass trajectory:

1. Load a Freeform animation.
2.  On the Key Info rollout turn on Trajectories.
3. On the Track Selection rollout click Body Horizontal or Body Vertical.
4. Turn on Sub Object Trajectories at the top of the Motion panel.
5. In the viewports select and move keys on the Center of Mass trajectory.

Working with Biped Animation

Working with Euler Curves on Biped Animation

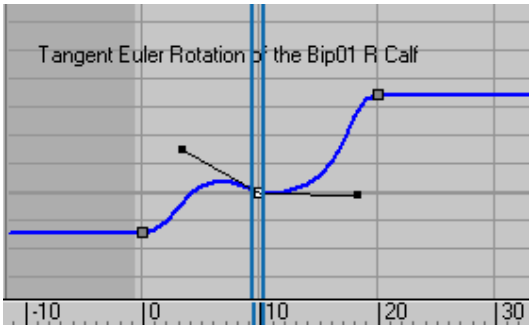
You can control a biped's position and orientation using Euler curves in addition to quaternion curves in the Workbench and Curve Editor. Using the Euler XYZ controller is an efficient way to animate your biped because you can use Bezier tangents to change the interpolation of your XYZ curves (quaternion curves do not have tangents). To learn more about how the Euler XYZ and TCB Rotation controllers differ from each other, refer to *Euler XYZ Rotation Controller (page 2–318)*.

You can switch between Euler XYZ and TCB Rotation controllers via the *Quaternion/Euler rollout (page 2–948)*. The *Curve Editor (page 2–1012)* displays the animation curves based on the chosen controller. Each curve is labeled starting with one of the following:

- "Quaternion Rotation of the ..."
- "Tangent Euler Rotation of the ..."
- "TCB Euler Rotation of the ..."

You can animate most biped parts (center of mass, pelvis, spine, head, neck, arms, legs, and tail) in Euler. However, fingers and toes are considered differently, as all first base links are controlled as quaternion and any subsequent links as TCB/Euler. Fingers and toes do not have tangents.

Biped limbs with only one degree of freedom (DOF), such as forearms and lower legs, are controlled with a single TCB/Euler curve.



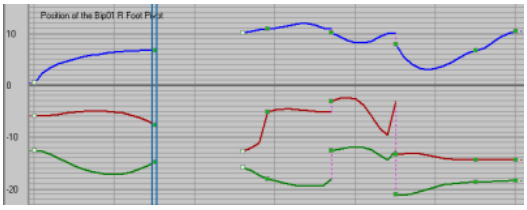
Tangent Euler Rotation curve

Note: Props are not supported with the Euler controller.

Rotation curves on a biped (including its center of mass) are always set in local parent space, whether they are controlled in Euler or quaternion.

The XYZ function curves of an Euler rotation track are locked together. This means that creating a new key on one axis automatically does so for all axes. Also, moving a key in time drags all three axes with it.

Displaying Position Curves



Bezier position curves are available only for the biped's hands, feets, and center of mass (COM). Hands and feets' positions are set in world space while the COM position is set local to the world.

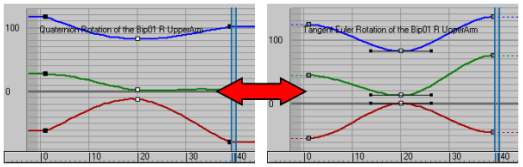
FK/IK key blending is illustrated as follows:

- Full lines represent IK periods.
- Gaps between lines represent FK periods.
- Vertical dotted lines represent a change in pivot points.

Note: Keys set to *Join to Prev IK Key (page 2–900)* are locked in value until the next un-joined key.

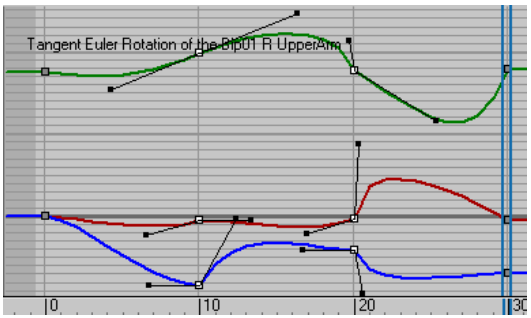
Tip: The Workbench and Curve Editor are synchronized so changing the curve display from the *Workbench toolbar (page 2–1012)* also updates the display from the Curve Editor, as well as activate the corresponding tool from the *Biped Toolbar (page 2–535)*.

Curve Conversion Between Quaternion and Euler



When an animation converts from one controller to the other, its integrity is maintained as much as possible. This means that despite variations in key interpolation results, the keyframed poses are preserved.

If you convert a quaternion rotation curve to Euler, the resulting tangent keys are set to Smooth. If, on the other hand, you switch an Euler rotation curve to quaternion, the bezier interpolation from the tangents is maintained as much as possible. However, if that interpolation is too extreme and can not be matched by TCB values, it is then reduced to fit under quaternion limits.



Extreme tangents on a Euler rotation curve

Using Euler animations with Layers, Motion Mixer, and Motion Flow

Both Euler tangents and quaternion TCB data are stored in each keyframe. Therefore, even if you are using the Euler XYZ controller, you can still change the profile of your curves using quaternion/TCB control values (on the *Key Info rollout (page 2–954)*).

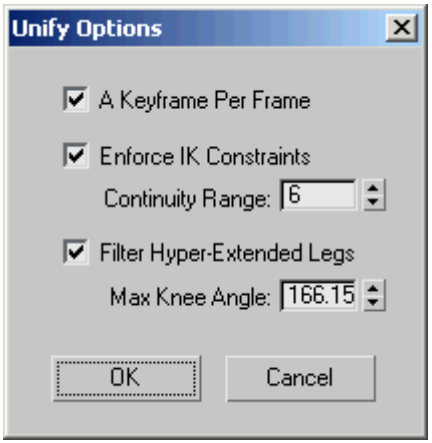
If you add a new layer to an Euler controlled biped part, that layer is converted into quaternion until it is collapsed back with the original layer. On collapse, the curve changes back into Euler following the same conversion rules stated earlier.

When clips containing Euler animations are brought into the Motion Mixer, Euler tangents are preserved in the blending process. However, once they are mixed together in a *mixdown (page 3–973)*, the resulting animation becomes quaternion.

Tip: When creating a mixdown, you can preserve your Euler tangents by turning on A Keyframe Per Frame in the *Mixer Preferences dialog (page 2–624)*.

Motion Flow follows the same rules as the Motion Mixer. If clips brought into Motion Flow contain Euler animations, their tangents are preserved. However, creating an unified motion results in a quaternion animation.

Tip: Euler tangents can be preserved when creating a unified motion by turning on A Keyframe Per Frame in the Unify Options dialog (by clicking Create Unified Motion in the *Motion Flow Scripts Group (page 2–1048)*).



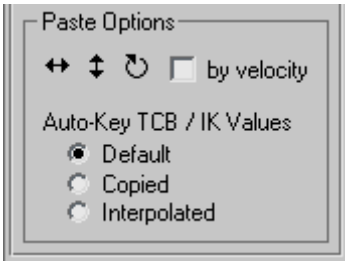
Turn on A Keyframe Per Frame to preserve Euler tangents.

Copy-Pasting Data between Euler and Quaternion Curves

When you copy a track and paste it onto another, its type (Euler or quaternion) is pasted as well. This sometimes results in animations switching from quaternion to Euler and vice-versa.

Euler tangents of copied poses or postures are never copied. If you copy an Euler track and paste it onto any type of track with Auto Key turned on, the current TCB or tangent values (depending on the type of track onto which you pasting) are ignored and instead set to Default .

Following the same rule, pasting a pose or posture on a Euler track with Auto Key turned on doesn't transfer the TCB values of the copied track. If a new key is created before pasting the track, the resulting tangent is set to Default.



Loading and Saving Biped Assets with Euler Curves

When you load an external biped asset file (BIP) onto a biped, they each might contain different animation controllers (for example, a Euler asset loaded onto a quaternion biped). When this happens, the biped animations are converted to match the controller from the biped asset file.

When you save a biped asset file, it stores both Euler tangents and quaternion TCB control values.

Note: While this doesn't affect overall performance, it may slightly impact the biped asset file size.

Working with XRef Bipeds

Using the *XRef Objects dialog (page 3–397)*, you can externally reference bipeds in your scene. This lets you work efficiently in a collaborative environment because the changes that the artist applies to his source XRef biped updates your master scene content. See *XRef Objects (page 3–394)* for more information.

Externally Referencing a Biped

When you create an XRef biped, certain modes on the Motion Panel's *Biped rollout (page 2–936)* become unavailable as XRef data is read-only:

- *Figure Mode (page 2–982)*
- *Footstep Mode (page 2–988)*
- *Motion Flow Mode (page 2–1043)*
- *Mixer Mode (page 2–652)*

As a result, you need to go back to the source scene in order to make desired changes to your biped.

Example: Changing the Structure of an Xref Biped

An integrator externally references a biped into his master scene. He discovers that the biped does not have the right number of spine and neck links, and

that the arm links are too short to reach a specific goal target. Because Figure Mode is unavailable in the master scene, the artist currently in the source scene inputs the changes to the biped's structure, and then saves his file. This update now reflects in the master scene.

Example: Modifying the Footsteps of an Xref Biped

An integrator externally references a biped into his master scene. He sees the footstep animation but cannot edit it because Footstep Mode is unavailable. The artist currently in the source scene goes into Footstep Mode and alters the animation. This update now reflects in the master scene.

Externally referencing a biped automatically creates a new "XREF LAYER" in the *Layers rollout (page 2–974)*, whether or not the source biped contains any animation. As a result, the original layer becomes a level one "Local Layer" which you can use to add animation on top of it.

Note: If the XRef biped contains multiple animation layers in the source scene, 3ds Max collapses them into one layer before creating the XRef.

Note: *Saving animation (page 2–920)* from your XRef biped breaks the XRef link.

Loading, Saving, and Displaying Biped Motion

Working with Biped Motion Files

character studio uses a variety of file formats to save, load, and edit motion.

- BIP file (*.bip)*

  The native **character studio** file format for saving biped motion. A BIP file saves *all* information about biped motion: footsteps, keyframe settings including limb rotation, the scale of the biped, the active gravity (GravAccel) value, and prop animation. IK Blend values for the keys and Object Space settings are also saved.

  See *Loading and Saving BIP Animation (page 2–920)*.

- BVH file (*.bvh)*

  The BioVision motion capture format. Typically, BVH files are obtained from motion capture hardware that records the movements of a human performer. Data saved in a BVH file includes both skeletal data, and information about limb and joint rotations.

  Note: For the BVH file specification, see the *BVH.rtf* document on the program disc.

- CSM file (*.csm)*

  The Character Studio Marker format. This is an alternative format for recording motion capture data. It stores positional data for various markers on a biped figure.

  See *Motion Capture Rollout (page 2–1065)*.

  Note: For the CSM file specification, see the *CSM.rtf* document on the program disc.

- MIX file (*.mix)*

  A MIX file contains data from the Motion Mixer, such as trackgroups, tracks, and transitions.

  See *Using the Motion Mixer (page 2–604)*.

- MFE file (*.mfe)*

  The Motion Flow Editor format saves scripts you create in the Motion Flow Editor, including clip references and transitions.

  See *Motion Flow Mode (page 2–1043)*.

- MNM file (*.mnm)*

The Marker Name File format matches custom marker names in a CSM or BVH file with the standard, preset marker names used by a biped.

Note: For BVH and CSM file specifications, see the *BVH.rtf* and *CSM.rtf* documents on the program disc.

- STP file (*.stp)*

  The Step file format saves footstep data. Unlike a BIP file, it does not save keys for the feet or upper body. STP files are ASCII files. The main purpose for this format is to enable developers to write programs that generate step files. Biped can later be used to create keys for the software-generated footsteps.

  See *Loading and Saving STP Files (page 2–924)*.

Motion File History Lists

The file dialogs for Biped motion files include a history list that lets you quickly choose directories where you have recently saved or opened files. 3ds Max maintains a separate history list for each type of motion file.

Note: When you load or save motion flow snippets, or mixer clips through the reservoir, if a file is already specified for that snippet or clip, its directory will appear at the top of the history list, even if that directory was not originally in the history list. If you cancel the load or save of the snippet or clip, that directory will not be added to the history. If you go ahead and load or save the snippet or clip, the directory will be added to the history.

Loading and Saving BIP Animation

Biped motion (BIP) files have a *.bip* file name extension. They save all information about biped motion: footsteps, keyframe settings including limb rotation, the scale of the biped, and the active gravity (GravAccel) value. IK Blend values for the keys and Object Space settings are also saved.

IK objects, props and other objects in the scene can also be saved with the *.bip* file.

There are several ways to create or acquire *.bip* files:

- By loading one of the sample animation files that come with 3ds Max. Refer to the *installation and support guide* for more details.
-  By creating your own animation with footsteps or freeform methods, and saving the animation with the Save button in the Biped rollout. See the procedures that follow.
-  By loading and filtering a motion capture file, then saving the animation with the Save button in the Biped rollout. See *Filtering Motion-Capture and Marker Data (page 2–1061)*.

When you save a *Motion Mixer (page 2–604)* (MIX) file or a *Motion Flow Editor (page 2–1026)* (MFE) file, the paths to the referenced BIP files are also saved. If these paths are broken (for example, if the BIP files have been moved or renamed), a Missing Motion Mixer Files dialog or Missing Motion Flow Files dialog, respectively, opens listing the missing files. To rectify this, open the *Asset Tracking dialog (page 3–487)* and set a new path for the files.

Procedures

To load biped motion from an existing BIP file:

1.  Select the biped you want to animate, and go to the Motion panel.
2.  On the Biped rollout, make sure you are not in Figure mode, and click Load File.

   Note: When *Figure mode (page 2–982)* is active, the Load File option loads figure (FIG) files.

Anything done in Figure mode changes the basic shape and structure of the biped. When Figure mode is turned off, the Load File tool loads BIP files, which animate the figure.

3. In the file dialog, choose the BIP motion file to load, and then click OK.
4. The biped repositions itself in the scene, as it assumes the initial position of the animation file. You might need to use Zoom Extents to see the biped after it is repositioned.

   Tip: The Biped keyboard shortcut Alt+R sets the animation range to that of the currently selected biped. Since the length of the animation can change after loading a BIP file, this keystroke can be useful. To use it, make sure that the Keyboard Shortcut Override Toggle is active.

To save animation you have created or imported to a BIP file:

1. Animate the biped with footsteps and/or freeform animation.
2.  Select the biped that has the animation you want to save, and go to the Motion panel.
3.  On the Biped rollout, make sure you are not in Figure mode, and click Save File.
4. In the file dialog, enter a name for the motion file, and then click OK.

Importing and Exporting Animation Data

You can use **character studio** with 3ds Max FBX import/export support to simplify the animation of 3ds Max bones. To transform an animated biped into a set of animated bones, export it to *FBX (page 3–558)* and then import your file.

Click the Help button on the FBX dialog to view the latest *3ds Max FBX Plug-in Guide*.

Motion Mapping: Retargeting Biped Motion

One of the more powerful features of Biped is the ability to retarget or map the motion of any biped onto any other biped. If you map the motion of a biped without a tail onto a biped with a tail, default tail motion for the biped with a tail will be computed. Default motion will also be computed when mapping the motion of a biped with fewer links in the legs, spine, or neck onto a biped with more links in the legs, spine, or neck. There are a few ways to perform motion mapping. You can:

- Go into Figure mode and change the structure of your biped. When you exit Figure mode, the new structure of the biped will adapt to the existing animation.
- Save a BIP or STP file from one biped, and load it onto another biped of a different structure and size.
- Copy footsteps from one biped, and paste them onto the footsteps of a differently structured and scaled biped.

Scale Stride Mode

Scale Stride mode gives you control over whether or not certain aspects of motion mapping occur.

If Scale Stride mode is active (the default):

- When you exit Figure mode after loading a FIG file or changing the biped's leg length, pelvis width, or height, the locations of the footsteps will be scaled to match the leg length and pelvis width, and gravity will be changed to match the new height of the biped.
- When you load a BIP or STP file, the locations of the footsteps in the file will be scaled to match the leg length and pelvis width of the

existing biped. Gravity will be adjusted to be proportional to the gravity stored in the file. (A motion stored in a BIP or STP file has a gravity value associated with it.)

- When you paste footsteps copied from one biped onto another biped, the locations of the footsteps in the buffer will be scaled to match the leg length and pelvis width of the existing biped. Gravity will be adjusted to be proportional to the gravity stored in the file. (A motion stored in a BIP or STP file has a gravity value associated with it.)

If Scale Stride mode is not active, no computations will be performed in any of the above cases. Then you might see your biped moving over footsteps that are spaced inappropriately far apart or close together for the size of your biped. Typically, you should leave Scale Stride mode active, unless you want to maintain the spatial relationship between the biped and other objects in your scene.

Procedure

To turn off Scale Stride mode:

1.  Select the biped, and go to the Motion panel.
2.  On the Biped rollout, click the bar at the bottom to expand the rollout.

   (If the rollout is already expanded, the bar shows a minus sign at the left, instead of a plus.)
3.  In the Modes group, click Scale Stride Mode to turn it off.

    The button changes to indicate stride scaling has been turned off.
4.  Use Figure mode to edit the figure.

   When you return to Keyframe mode or Footstep mode, the biped's stride length is unchanged, regardless of the biped figure's new proportions.

Merging and Cloning a Character

At some point you might need to use the File > Merge command to merge a character into your scene. Also, you can use Merge to clone a biped that has a mesh with the Physique modifier applied, as described in this topic.

Cloning a Character

To clone a character (that is, a biped with a Physique mesh), save a scene that contains the character, change the original biped's name, and then use File > Merge to merge the copy from the saved file. You can merge from a saved version of the scene that is currently open. See the procedures for details.

Procedures

To merge a skinned biped:

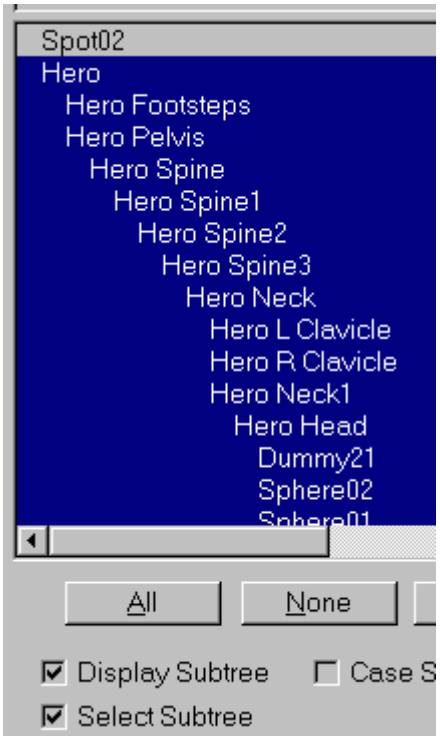
1.  Before you merge, make sure the root name of the biped you want to import is different from that of bipeds already in your scene. If there is a name conflict, select the conflicting biped, go to the Motion panel, and use the Structure rollout's Root Name field to change the biped's root name.

   If you skip this step, you will get a Duplicate Name dialog for every object in the merged biped's hierarchy!
2. Choose File > Merge, then choose the MAX file that contains the biped to import.

   The Merge dialog displays.
3. Click to turn on Select Subtree in the Merge dialog.
4. Locate and click the biped's center of mass name in the Merge dialog list.

The default name for the biped center of mass is *Bip01*. If the biped was renamed on the Structure rollout, find the renamed center of mass in the list.

Note: The center of mass is the root object in the biped hierarchy, if this is selected with Select Subtree turned on, then all the child links are selected, including extra bones and the mesh skin (the mesh skin is linked when Attach To Node is used in Physique).



Selecting a biped while merging a MAX scene

In the illustration, the center of mass object is named *Hero*. With Select Subtree active, all the children are also selected when you click *Hero*, including the Physique mesh.

5. Click OK.

   Probably a Duplicate Name dialog is displayed. If it isn't, the merge is done; if it is, see the following step.

6. On the Duplicate Name dialog, turn on Apply To All Duplicates, and then click Merge.

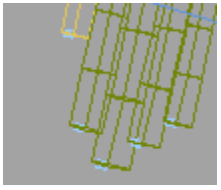
   Note: Even when the bipeds in your scene have different root names than the biped you are merging, the biped finger/toe/head dummies in the file to be merged can have the same names as other biped dummies in your scene. This step merges these dummy objects. Dummies with duplicate names can appear in the Select Objects dialog, so the duplicate names are not a problem.

   The biped and its animation is merged with the scene.

To hide the finger, toe, and head dummy objects:

The dummy objects for fingers, toes, and the head are visible on the newly merged biped. Usually these dummies are hidden from sight. They are used by Physique to create envelopes for all the finger tips, toe tips, and head; these dummies display when a character is merged. The quickest way to hide them again is simply to toggle the Objects button in the Biped rollout's Display group, as described in these steps.

1.  With the merged biped still selected, go to the Motion panel.

   Fingertip dummies displayed as blue cubes

   The number of biped dummies varies according to how many fingers and toes the character has.

2.  At the bottom of the Biped rollout, click the bar by the plus sign to expand the rollout.

(If the rollout is already expanded, the bar shows a minus sign at the left, instead of a plus.)

3.  In the Display group, turn Objects off and then back on again to hide the biped finger/toe/head dummy objects.

To clone a skinned biped:

1. Save a copy of the scene that contains the biped you want to clone.
2. On the Structure rollout, change the original biped's root name, as described above.
3. Merge the saved biped, as described in the previous procedure, "To merge a skinned biped."

   The original biped, with its Physique modifier, is cloned. It appears in the same location as the original biped.
4.  On the Motion panel, on the Biped rollout, turn on Move All Mode.
5.  Use Select And Move to move the clone to a new location in the scene.

Combining BIP Motions

character studio provides two main ways of combining BIP files to build more complex character animations.

- *Motion flow (page 2–1026)* uses BIP files as *clips* in a *script*. The motion-flow script joins clips together using *transitions*. Transitions can be unconditional, they can be chosen at random, or they can be governed by rules such as collision detection. You can control when a transition begins and ends. You can use motion flow to animate a single biped, or a crowd of bipeds.

  Motion flow scripts are saved as Motion Flow Editor (MFE) files.
- The *Motion Mixer (page 2–604)* also uses BIP files and MFE files as clips. In addition to creating transitions from one BIP animation to another, over time, Mixer scripts can combine the upper-body movement of one biped with the lower-body movement of another, adjust the timing of movements, and provide a number of other effects. The Mixer is especially useful when you work with *motion-capture files (page 2–1059)*.

  Mixer scripts are saved as BIP files or MIX files.

Loading and Saving STP Files

Step (STP) files save footsteps, but don't save body keyframes.

The STP file format is an ASCII format that enables developers to write programs that generate step files for biped motion. The online document *stp.rtf*, provided with **character studio** in the *..\cstudio\docs* directory, describes the STP format.

Procedures

To load footstep data:

1.  Select the biped to animate via saved footsteps, go to the Motion panel.
2.  On the Biped rollout, make sure you are not in Figure mode, then click Load File.
3. In the file dialog, choose Step Files (.STP) as the file type to load.
4. Choose the footstep file to load, and then click OK.

The footsteps are loaded onto the biped, and new default keys are generated to match the footsteps.

To save footstep data:

1.  Select the biped whose footsteps you want to save, and go to the Motion panel.
2.  On the Biped rollout, make sure you are not in Figure mode, then, click Save File.
3. In the file dialog, choose Step Files (.STP) as the file type to save.
4. Enter a name for the footstep file, and then click OK.

Using Motion-Capture Data

Besides animating a biped with footsteps or with keyframing (freeform animation), you can import a motion-capture file. The overall workflow for motion capture is straightforward:

- Import the motion-capture data

  See *Importing Motion-Capture Data (page 2–1061)*.

- Filter the data

  Motion-capture files are typically dense, with keys on every frame. Filtering the data reduces the number of keys, resulting in an animation with better performance. You filter the data with the *Motion Capture Conversion Parameters dialog (page 2–1070)*. This dialog can reduce large numbers of keys at once. It is displayed when you use the Motion Capture rollout to load a BIP, BVH, or CSM file. Once you have loaded a motion-capture file, you can filter the data further by clicking Convert From Buffer (also on the Motion Capture rollout), which also displays this dialog.

  See *Filtering Motion-Capture and Marker Data (page 2–1061)*.

- (Optional) Edit the data

  Once you are happy with the converted animation, you can use the freeform animation tools to make specific changes.

  Tip: The *Motion Mixer (page 2–604)* provides another way to edit motion-capture data.

Motion-capture files can be one of three file types: the native **character studio** BIP format, the BioVision (BVH) format, or the Character Studio Marker (CSM) format.

For an introduction to using motion capture, see the tutorial "Working with Motion-Capture Data."

Correcting Posture

A particular motion file might position a biped body part inappropriately. For example, the collarbones might be rotated down too far, affecting your mesh deformation. All you need to do is go into Figure mode, rotate the biped collarbones up, and then exit Figure mode: this corrects the collarbone position for the entire animation. The motion references the Figure mode position, if the biped is adjusted, and this adjustment is reflected in your animation when you exit Figure mode.

Tip: Save a FIG file for the biped pose you used when you applied Physique. Then you can use the FIG file to reload this position if you need to reapply the Physique modifier or reinitialize the Physique settings.

How Biped Uses Figure Mode

When you animate the biped, the Biped plug-in maintains the at-rest pose you have created for these elements of the biped body:

- Spine
- Neck
- Clavicles
- Tail

- Ponytails
- Center of mass position, relative to the body.

When Biped adapts the keyframed motions stored in BIP files to different characters, the keyframes of the above elements are recreated as an offset from the at-rest posture associated with each character's figure. The at-rest posture associated with the arms and legs is always assumed to be a standing posture, with straight legs.

Restructuring a Biped to Match a BIP File

When you load a BIP file, there is an option to restructure the biped to match the file. If you turn this on when you load the file, the biped's structure will change to match the figure of the biped in the BIP file.

Talent Figure Mode and Adjust Talent Pose

Talent Figure Mode and *Adjust Talent Pose* on the Motion Capture rollout have a purpose similar to that of Figure mode. They are used to size and position biped body parts to better fit raw motion capture data. After importing motion capture data, you might discover that certain biped limbs or the biped scale need a global adjustment in order to provide a closer match to the figure of the talent who performed the motion.

See also

Motion Capture Rollout (page 2–1065)

Copying and Pasting Tracks

Biped's Copy/Paste rollout on the Motion panel provides controls for copying and pasting biped tracks from one part of the biped to another, or from one biped to a different biped. At the top of this rollout are three buttons: Posture, Pose, and Track.



This section covers the copying and pasting of tracks. For information on the other options on this rollout, see *Copying and Pasting Postures and Poses (page 2–910)*.

To copy and paste biped tracks, turn on Tracks. The remaining controls in the rollout change to reflect that you are working with tracks:

-  The copy button changes to a Copy Tracks button.
-  The two Paste buttons become Paste Tracks and Paste Tracks Opposite, respectively.

-  The Paste Horizontal, Paste Vertical, and Paste Rotation buttons in the *Paste Options group (page 2–973)* become active.

For hands-on experience in using Copy and Paste Tracks, see the lesson called Creating a Simple Freeform Animation in the tutorial *Animating with Freeform*.

Procedure

To copy a track from one biped to another:

1. Select any part of the biped and access the Motion panel.
2.  On the Copy/Paste rollout, click Create Collection and rename the new collection **Biped Track**.
3. Click Track to use Track mode. Then, select the biped parts whose tracks you want to copy.
4.  Click Copy Track.

   Biped creates a new track buffer, and adds it to the list as the active buffer. The buffer's name indicates which biped parts you selected.
5. Select any part of the other biped.
6.  In the Paste Options group, enable all three Paste buttons.
7.  Click Paste Track or Paste Track Opposite.

   The tracks for all biped parts are applied to the other biped regardless of which parts are selected on the biped.

Repositioning the Biped

The method you should use to reposition a biped depends on which tools were used to animate the biped:

- If the biped uses freeform or footstep animation, use Move All mode.

  This is also the easiest way to move a biped that is not animated.
- If the biped uses freeform animation *and* IK attachment (that is, there are keys where IK Blend is greater than zero), then you need to use a dummy object.
- If the biped is animated with motion flow, use the Start Position spinners.

Procedures

Note: The procedures in this topic assume that you have already selected the biped and gone to the Motion panel.

To reposition a freeform animation (with no IK attachment) or a footstep animation:

1.  On the Biped rollout, turn on Move All Mode.

   The biped's center of mass is selected and displayed in a larger-than-usual size. On the main toolbar, Select And Move is turned on.

   Move All mode is a convenient way to select the biped's center of mass, and to reposition the biped. Use it when the biped is not yet animated, when the biped has freeform animation with no IK attachment, or when it has footstep animation.

   Note: The alternative way to move a biped with footstep animation is to select all of its footsteps, rather than its center of mass.
2.  Move the biped.

   The Collapse button on the Move All dialog box allows you to reset the position and rotation values in the Move All dialog to zero, but does not change the position of the biped.

3.  Turn off Move All Mode.

To reposition a biped with limbs attached to an Object Space object (IK attachment):

1. Create a dummy object and position it near the biped's center of mass.
2.  On the main toolbar, use Select And Link to link both the Object Space object and the biped's center of mass to the dummy object.

   Important: Make sure you link both at the same frame.
3.  Move the dummy object to reposition the biped.

To reposition a biped with limbs attached to world space (IK attachment):

A biped with limbs in world space has an IK Blend setting of 1.0 for each limb, with no Object Space object specified. For example, when you plant the feet of a biped so its knees will bend when you lower its center of mass, or you plant the hands of a biped doing pushups, typically you do not specify an Object Space object.

1. Create a dummy object near the biped's center of mass.
2.  On the main toolbar, use Select And Link to link both the biped's center of mass and the biped part with planted keys to the dummy object.

   Tip: To be on the safe side, you can double-click the upper limb and link all of the limb to the dummy.

   Important: Make sure you create all links at the same frame.
3.  Move the dummy object to reposition the biped.

   Note: Depending on the animation, at some of attached object's keys you might need to set IK Blend back to **0.0**. Otherwise, the limb can get "stuck" in a posture.

To reposition a biped animated with a motion flow script:

By default, the biped's position at the beginning of a motion flow script is the position specified by the first BIP or STP file in the script.

1.  Turn on Motion Flow mode.
2. On the Motion Flow rollout, use Start Position X, Start Position Y, and Start Position Z to change the XYZ position of the biped.
3.  Turn off Motion Flow mode.

To use layers to reposition a biped with freeform animation:

Using layers and Auto Key when you reposition a freeform biped gives you some editing choices, as described below.

1.  On the Layers rollout, click Create Layer.

   Note: If the biped uses footstep animation, the Move All Mode button will become unavailable.
2. 

Turn on Auto Key.
3. Reposition the biped as described in the previous procedure, "To reposition a freeform animation (with no IK attachment) or a footstep animation."
4. 

Turn off Auto Key.

   You now have these aids to animation editing:

   - Lower layers show the biped's original position, as does the red stick figure.
   -  Clicking Snap Set Key on the Layers rollout snaps the biped to its position on the previous layer, and sets a position key. The

animation is blended between the biped's new position and its previous position at frames where the key is set.

Previewing Biped Motion

There are two types of animation playback available within **character studio:**

-  You can use the Play Animation button in the viewport playback controls to play biped animation, the same as you can any other 3ds Max animation.
-  Use the *Biped Playback button (page 3–916)* on the Biped rollout to play back the animation with the biped displayed as a stick figure, while hiding everything else in the scene.

Note: It is possible to turn on both buttons at once. This is not recommended. Depending on your system configuration, turning on Play Animation and Biped Playback at the same time can considerably slow down performanc.

Biped Playback

Biped playback previews the motion of all existing, visible bipeds. If you hide a part or all of biped, the hidden biped or biped part does not appear in the Biped playback. Objects in the animation that are not bipeds do not appear in Biped playback. If the 3ds Max home grid is visible, a grid appears at Z=0 in the Biped playback viewport.

While Biped controls are visible in the Motion panel, you can also press V to start or stop Biped Playback. For this keyboard shortcut to work, the Keyboard Shortcut Override Toggle must be turned on.

Biped always previews the existing animation range. Biped Playback responds to the parameters in the 3ds Max Time Configuration dialog. If Real Time playback is chosen in this dialog, Biped plays back at the current frame rate, sometimes skipping frames if necessary. If Real Time is turned off, Biped plays back as fast as it can, depending on the capacity of the graphics card installed on your system.

Note: Hardware acceleration has no effect on Biped playback. If you are using a hardware-accelerated display card, you might find 3ds Max playback to be faster under certain circumstances.

Procedures

To preview biped motion using biped playback stick figure:

1. Activate the viewport with the view you want to see.
2. Hide or show the bipeds you want to appear in the playback.
3.  Select one of the visible bipeds, and go to the Motion panel.
4.  On the Biped rollout, click Biped Playback.

   Stick-figure biped animation plays back in the active viewport. Animation plays back for all visible bipeds.

To preview biped motion using the full biped model:

1. Drag the time slider.
2.  Click Play Animation.

Biped animation plays back in the active viewport.

Depending on your system, biped animation might not play back in real time using 3ds Max viewport playback.

Warning: You might miss critical frames of your animation if Real Time is selected in the Time Configuration dialog.

In Place Mode 

Create or select a biped. > Motion panel > Biped rollout > Expand rollout > Modes group > In Place button on In Place flyout

In Place mode (page 3–956) allows you to display biped motion as if it were occurring on a treadmill. Regardless of the distance the biped covers under control of the current motion file, the biped stays within the active viewport when you've turned on In Place Mode.

In Place playback prevents lateral (XY) movement of the biped center of mass during animation playback; vertical motion along the Z-axis is preserved. Biped limbs, footsteps, and center of mass keys can be adjusted in this mode. When the center of mass is moved on the XY axes in this mode, the footsteps move.

Use this feature to view biped playback without requiring a follow camera. In this viewing mode, visible footsteps appear to slide under the biped.

For export to games, this feature is valuable, as many game engines intelligently move the character's center of mass laterally according to game play. In Place mode makes it easy to view, tune, and export animation in a manner that is complementary to game engine playback.

Note: Trajectories do not display using In Place mode.

In Place Mode Options

The In Place Mode button is the default button on a three-button flyout. There are two further options:

-  **In Place X Mode—**Locks center of mass X-axis motion. Use this for game export where the character stays in place but the swinging motion of the hips and upper body along the Y-axis is preserved.
-  **In Place Y Mode—**Locks center of mass Y-axis motion. Use this for game export where the character stays in place but the swinging motion of the hips and upper body along the X-axis is preserved.

Using In Place Mode to Adjust Keyframes

In Place Mode is a good way to adjust keys on a biped that already has animation applied to it. Rather than scrolling the view at different frames to keep a running biped visible, turn on In Place Mode. Now when you scrub the time slider, or use the Next Key and Previous Key buttons, the biped remains visible. A key that needs adjustment can be quickly spotted and corrected.

Tip: Judging lateral center-of-mass motion using In Place mode is difficult. In Place mode limits center-of-mass motion on the XY axes; all sense of body momentum on these axes is suspended during playback. You might want to turn off In Place playback to gain a sense of lateral momentum when setting or adjusting horizontal keys (Body Horizontal Track) for the center of mass.

Procedure

To use In Place mode to adjust keyframes:

1.  Select a biped that has animation, and go to the Motion panel.
2.  On the Biped rollout, click the bar at the bottom to expand the rollout.

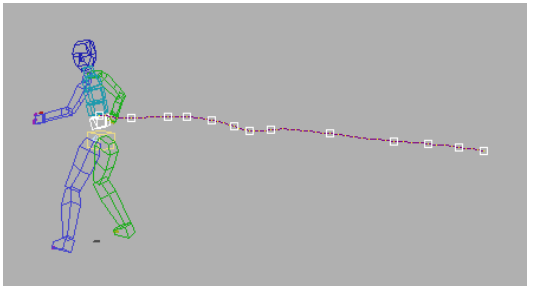
   (If the rollout is already expanded, the bar shows a minus sign at the left, instead of a plus.)
3.  In the Modes group, click In Place Mode to turn it on.
4. 

   At the bottom of the3ds Max window, turn on Auto Key.

5. Find a frame where the biped needs adjustment, and modify or add keys.

Trajectory Display

When a biped is animated, you can view its motion not only using Biped Playback, but you can also see the path, or trajectory, the biped (or selected biped links) follows throughout the motion.



Display of the trajectory of a biped's lower spine

You turn on trajectory display by clicking Trajectories in either of these user-interface locations:

- Motion panel > Biped rollout > Expand Biped rollout. > Display group > Trajectories button
- Motion panel > Key Info rollout > Trajectories button

Tip: Do not use 3ds Max trajectory controls to display or edit biped trajectories. These include the options accessed with Object Properties > Display Properties > Trajectory on the quad menu, and those under Motion panel > Trajectories text button.

Trajectories provide useful visual feedback as you edit keys, showing the effects on the motion path for the parameters you're adjusting. You also can use trajectory display to compare filtered and unfiltered motion capture data.

Changes to Tension, Continuity, Bias, Dynamics Blend, Ballistic Tension, and the overall gravity setting GravAccel are reflected in the trajectories.

Note: Trajectories do not display while you play an animation using In Place mode.

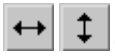
You can customize the trajectory information for the selected biped link by using the *Display Preferences dialog (page 2–944)*. You can choose between Bone Base and Bone Tip, show an entire trajectory, or a moving range of frames.

See also

Biped Color-coded Keys and Trajectories (page 2–1005)

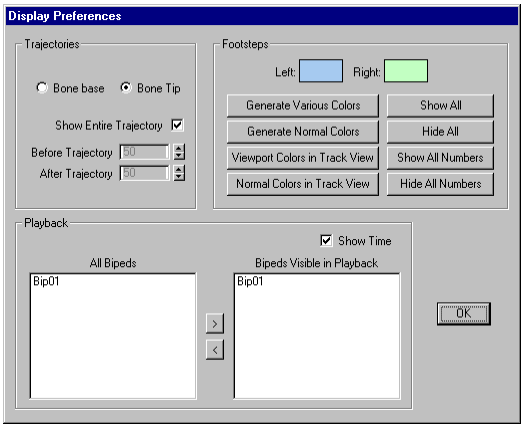
Procedure

To edit biped trajectory keys in a viewport:

1.  Select an animated biped and go to the Motion panel.
2.  On the Track Selection rollout, turn on Body Horizontal or Body Vertical.
3. At the bottom of the Biped rollout, expand the Modes And Display sub-rollout if necessary by clicking its name.
4.  In the Display group, turn on Trajectories.
5. Turn on Sub-Object selection level.
6.  On the main toolbar, turn on Select And Move.
7. Use the Transform gizmo to move the keys on the trajectory.

Display Preferences

The *Display Preferences dialog (page 2–944)* lets you customize how bipeds are displayed while you work with them. This topic introduces some of the commonly used options.



Use controls on this dialog to change footstep, trajectory, and playback display.

In the Trajectories group, you can choose between the Bone Base or Bone tip for trajectory display. You can show the entire trajectory, or define a range of frames for partial trajectory display.

In the Footsteps group, you can choose colors for the left and right footsteps, and generate various colors or standard colors in the viewport or Track View. You can also Show or Hide All Footsteps or Footstep Numbers.

In the Playback group, you can define which bipeds will appear in Biped Playback. (This can be useful for speeding performance if your scene contains a number of bipeds.)

See also

Display group (page 2–939)

Biped User Interface

The Biped user interface is split up into "modes" of operation. You activate these modes by selecting the appropriate button in the Biped rollout on the Motion panel, which is visible when a biped is selected.

There are four modes available:

 Figure mode is used to change the biped skeletal structure and to align the biped to a mesh.

 Footstep mode is used to create and edit footstep animation.

 Motion Flow mode is used to create scripts that combine motion files into longer animations.

 Mixer mode is used to view, save, and load animation created with the Motion Mixer.

As with other parts of 3ds Max, the rollouts change depending on the mode the software is in. When no modes are active, the following rollouts are displayed:

- *Assign Controller Rollout (character studio) (page 2–934)*
- *Biped Apps Rollout (page 2–935)*
- *Biped Rollout (page 2–936)*
- *Track Selection Rollout (page 2–945)*
- *Quaternion/Euler Rollout (page 2–948)*
- *Twist Poses Rollout (page 2–950)*
- *Bend Links Rollout (page 2–952)*
- *Key Info Rollout (page 2–954)*
- *Keyframing Tools Rollout (page 2–962)*
- *Copy/Paste Rollout (page 2–966)*
- *Layers Rollout (page 2–974)*
- *Motion Capture Rollout (page 2–1065)*
- *Dynamics & Adaptation Rollout (page 2–980)*

The Assign Controller, Biped Apps, and Biped rollouts are displayed during all modes. The remaining rollouts depend on the mode.

In Figure mode, *Structure (page 2–984)* is the only additional rollout.

In Footstep Mode, the rollouts displayed are: *Footstep Creation (page 2–988)*, *Footstep Operations (page 2–990)*, and Dynamics & Adaptation.

In Motion Flow mode, *Motion Flow (page 2–1045)* is the only additional rollout displayed.

In Mixer mode, *Mixer (page 2–652)* is the only rollout displayed.

When no modes are active, you can edit tracks and keys, set IK constraints, work with layers, and work with motion capture data. You can also create freeform animation without any of the modes active simply by turning on the Auto Key button and moving or rotating any part of the biped.

Motion Panel (Biped)

Once you have created a biped, use the Biped controls on the Motion panel to animate the biped, load and save Biped files, and fit the biped to a mesh representing your character.

Create a biped if one does not exist, and select any part of the biped to see the controls on the Motion panel.

Depending on the active mode, the Motion panel contains the following rollouts:

Assign Controller Rollout (character studio) (page 2–934)

Biped Apps Rollout (page 2–935)

Biped Rollout (page 2–936)

Track Selection Rollout (page 2–945)

Copy/Paste Rollout (page 2–966)

Footstep Creation Rollout (page 2–988)

Footstep Operations Rollout (page 2–990)

Motion Flow Rollout (page 2–1045)

Mixer Rollout (page 2–652)

Key Info Rollout (page 2–954)

Keyframing Tools Rollout (page 2–962)

Display group (page 2–939)

Layers Rollout (page 2–974)

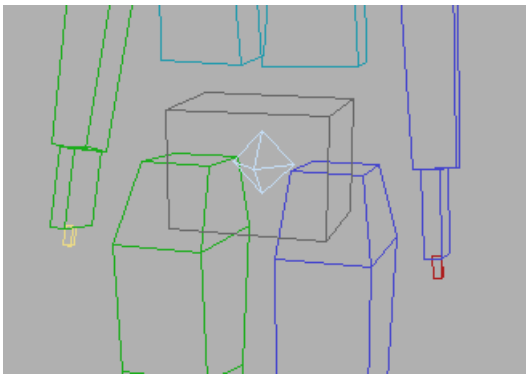
Motion Capture Rollout (page 2–1065)

Dynamics & Adaptation Rollout (page 2–980)

Structure Rollout (page 2–984)

Center of Mass

The parent object of the biped is its center of mass, which appears as a blue octahedron near the center of the biped's pelvis. The name of this object is the root name of the biped (*Bip01* by default).



The Center of Mass (COM) object is the root parent.

The center of mass is sometimes referred to as the COM.

Center of Mass Animation Tracks

Unlike other objects and biped body parts, the center of mass has three separate animation tracks:

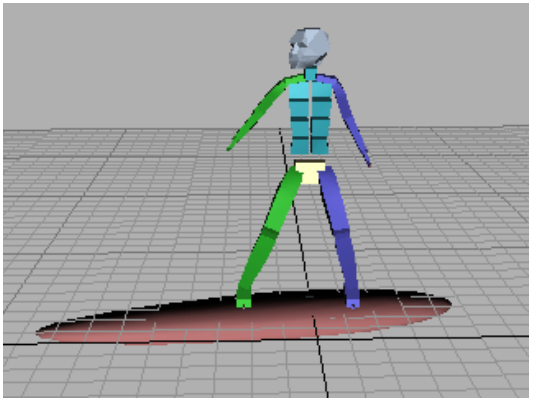
-  Body Horizontal

-  Body Vertical
-  Body Turning (selected with the Body Rotation button)

You can use these buttons, available on the *Track Selection rollout (page 2–945)*, to select each track as well as the COM itself.

Linking the Center of Mass Object

Using the *Select And Link tool (page 2–422)* you can link the center of mass object to another object if you need to reposition an animation sequence. An example of this might be a surfer. You could create an animation of a biped running up and down the surfboard, hanging its toes off the end of the surfboard, and so on, and then link the center of mass object to the board. When you animate the surfboard, the biped animation would travel with it.

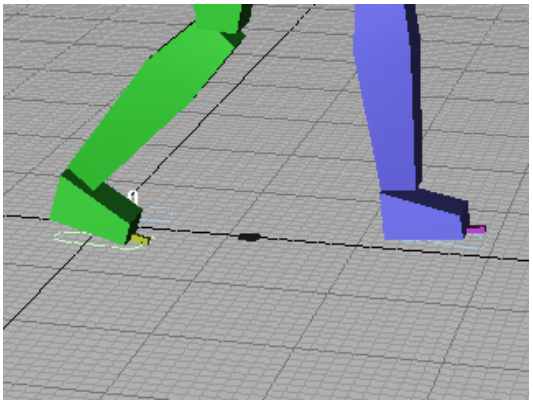


The biped moves with the surfboard because the center of mass is linked to the surfboard.

Because the COM is the root object in the biped hierarchy, it is the only object in the biped that can be linked directly to other objects in the scene with Select And Link. A different technique is used to make other body parts, such as feet and hands, follow other objects in the scene. See *Animating IK Attachments (page 2–908)*.

Center of Mass Shadow

The center of mass shadow object, the circle between the biped's feet on the world plane, provides a sense of where the character's center of mass is positioned relative to the feet. Another use of the center of mass shadow is to link objects to it. For example, a camera and camera target could be linked to the shadow object to follow the character.



The center of mass shadow between the biped's feet.

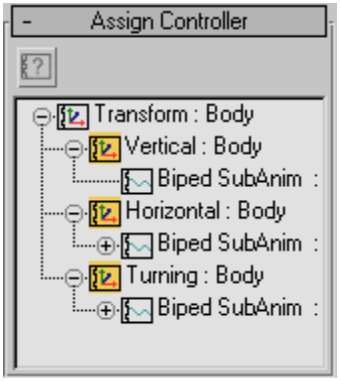
Motion Panel Rollouts (Biped)

Assign Controller Rollout (character studio)

Select a biped's COM. > Motion panel > Assign Controller rollout

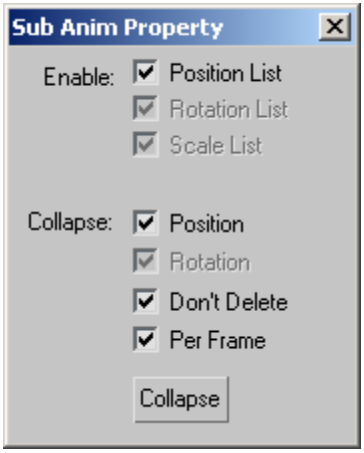
The Assign Controller rollout assigns and appends different transform controllers to individual objects. You can also assign controllers in Track View. Animation controllers are plug-ins that handle all of the animation tasks in the software. For a complete list of available Animation controllers, see *Assign Controller (page 2–546)*.

Interface



Assign Controller— Displays a selectable list of controllers for a selected track.

Once controllers have been added, right-click the Biped SubAnim entry in the list and choose Properties to display the SubAnims dialog.



Enable section

Enable— Selectively activates or disables the three list controllers. Choose any combination of Position, Rotation, and Scale.

Position List— When turned on, selects the Position controller to be collapsed into the Biped SubAnim track.

Rotation List— When turned on, selects the Rotation controller to be collapsed into the Biped SubAnim track.

Scale List— When turned on, selects the Scale controller to be collapsed into the Biped SubAnim track.

Collapse section

Position—Collapses the Position controller when you click Collapse.

Rotation—Collapses the Rotation controller when you click Collapse.

Don't Delete— Prevents the removal of the list controller after collapsing and hides the controller instead.

Per Frame— Creates a key per frame during the collapse, preventing the controller from collapsing onto the key times of the biped and SubAnim controller.

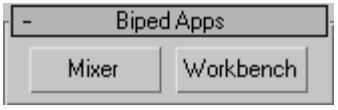
Collapse— Performs the collapse.

Biped Apps Rollout

Select a biped. > Motion panel > Biped Apps rollout

The Biped Apps rollout allows you to choose either of two tools for working with the biped motion.

Interface



Mixer—Opens the Motion Mixer, where you can layer animation files to customize biped motion.

See *Using the Motion Mixer (page 2–604)*.

Workbench—Opens the Workbench, where you can analyze and adjust biped motion curves.

See *Working with the Workbench (page 2–1008)*.

Biped Rollout

Select the biped > Motion panel > Biped rollout

Use controls on the Biped rollout to put the biped into Figure, Footstep, Motion Flow, or Mixer modes, and to load and save *.bip*, *.stp*, *.mfe*, and *.fig* files. You'll find other controls on the Biped rollout, as well.

The Modes group on the Biped rollout lets you turn on the Buffer, Bend Links, Rubber Band, Scale Stride, and In Place mode.

The Display group on the Biped rollout lets you adjust how the biped is displayed, providing controls to show or hide Objects, Bones, Footsteps, and Trajectories.

In addition, the Biped rollout provides controls for converting footsteps into a freeform animation, or a freeform into a footstep animation.

Note: The Modes group, Display group, and Name field are hidden by default. Click on the Modes and Display expansion bar to display them.

Note: In previous versions of **character studio**, the Biped rollout was called the General rollout.

Interface

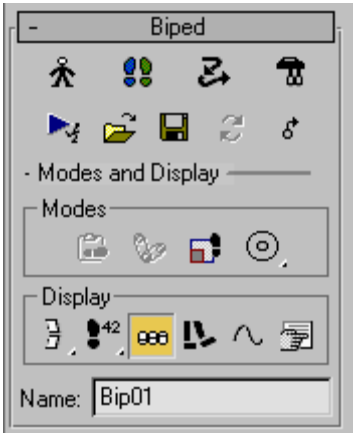


Figure Mode—Use *Figure mode (page 2–982)* to fit a biped to the mesh or mesh objects representing your character. Leave Figure mode on when you attach the mesh to the biped with Physique. Figure mode is also used to scale a biped with a mesh attached, to make biped "fit" adjustments after Physique is applied, and to correct posture in motion files that need a global posture change.

The *Structure Rollout (page 2–984)* appears when Figure mode is active.

Note: When Figure mode is turned on, the biped jumps from its animated position to its Figure mode pose. Animation is preserved when you exit Figure mode.

Footstep Mode—Create and edit footsteps; generate a walk, run, or jump footstep pattern; edit selected footsteps in space; and append footsteps using parameters available in Footstep mode.

Two additional rollouts display on the Motion panel when *Footstep mode (page 2–988)* is active:

- *Footstep Creation rollout (page 2–988)*
- *Footstep Operations rollout (page 2–990)*

 **Motion Flow Mode—**Create scripts and use editable transitions to combine *.bip* files together to create character animation in *Motion Flow mode (page 2–1043)*. After creating a script and editing transitions, use Save Segment on the Biped rollout to store a script as one long *.bip* file. Save a *.mfe* file; this enables you to continue Motion Flow work in progress.

Tip: Use Motion Flow mode to cut motion capture files together.

Note: The *Motion Flow rollout (page 2–1045)* appears when Motion Flow mode is active.

 **Mixer Mode—** Activates any current *Mixer (page 2–604)* animation on the Biped and displays the *Mixer rollout (page 2–652)*.

 **Biped Playback—**Plays the animation for all bipeds unless they are excluded on the *Display Preferences dialog (page 2–944)*. This playback mode usually gives real-time playback, which you may not get if you use Play on the3ds Max toolbar.

Note: In Biped Playback mode, the biped is displayed as bones only, with no other scene objects visible.

 **Load File—**The *Open dialog (page 2–942)* lets you load *.bip*, *.fig* or *.stp* files.

 **Save File—**Opens the *Save As dialog*, where you can save Biped files (*.bip*), figure files (*.fig*), and step files (*.stp*) files.

 **Convert—**Convert a footstep animation to a freeform animation. This works in both directions. Displays the *Convert to Freeform dialog (page 2–999)* or *Convert to Footsteps dialog (page 2–999)* depending on the direction.

- Convert uses biped foot IK Blend values to extract footsteps.
- Use Convert to extract footsteps from an animation saved using Save Segment in Motion Flow mode.
- Convert the animation in either direction depending on how you like to work. Convert to freeform for unrestricted key editing. Convert to footsteps to take advatage of footsteps.

 **Move All Mode—**Allows the biped to be moved and rotated with its relative animation intact. You can transform the biped interactively in the viewport or with the dialog box that opens when the button is active.

When this button is active, the biped's center of mass enlarges to be more easily selected for translation.

The Collapse button on the Move All dialog box allows you to reset the position and rotation values in the Move All dialog to zero, but does not change the position of the biped.

Modes group



Note: The Modes group is hidden by default. To display it, click on the Modes & Display expander in the Biped rollout.

 **Buffer Mode—**Edit segments of an animation in Buffer mode. Copy footsteps and associated biped keys into the buffer using Copy Footsteps on the Footstep Operation rollout first, then turn on Buffer mode to view and edit the copied segment of your animation.

Tip: Paste buffered motion back to the original animation repeatedly to create looping motions.

Edit footstep and biped animation that have been copied into the buffer using Copy Footsteps on the

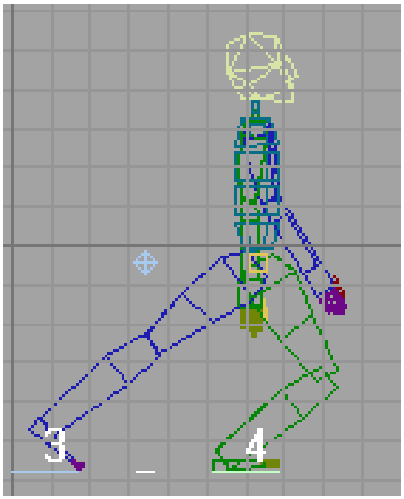
Footsteps Operation rollout. The changes can be pasted back by turning off Buffer Mode, turning on Paste Footsteps on the Footstep Operation rollout, and overlapping the buffered footsteps with the original footsteps. The buffered motion is spliced into the original animation.


Rubber Band Mode—Use this to reposition the biped elbows and knees without moving the biped hands or feet in Figure mode. Reposition the biped center of mass to simulate the physics of wind or weight pushing against the biped. Figure mode must be turned on to enable Rubber Band Mode.

To reposition biped knees and elbows, turn on Figure mode and turn on Rubber Band mode. Select the Move transform tool, then select and drag a biped upperarm or thigh in the viewports. Use this as an aid to fitting a biped to a mesh.

To reposition the biped center of mass relative to the rest of the biped skeleton, turn on Figure mode and turn on Rubber Band mode. Select the Move transform tool, then select and drag the center of mass in the viewports. Use this to account for wind force or pushing against a heavy object. See *Shifting the Biped's Balance (page 2–876)*.

Note: Rubber Band mode behaves differently than Non-Uniform Scale. For example, if you "Rubber-Band" the biped thigh, the thigh and biped calf objects scale proportionally to keep the biped foot stationary. Using Non-Uniform Scale, the calf retains its scale and the foot moves.



Moving the biped center of mass (blue diamond) behind the character, turns this default walk cycle into a struggle against a high wind.


Scale Stride Mode— Footstep stride length and width are scaled to match the stride length and width of the biped figure. Scale Stride mode is on by default.


Displays when Scale Stride mode is *off*.

Scale Stride mode is on by default, so scaling occurs automatically when you load a *.bip*, *.stp*, or *.fig* file. Scaling occurs when you paste footsteps and when you scale the biped's legs or pelvis.

For example, if you load a *.bip* file that was saved from a larger biped, the footsteps come into your current scene scaled to match the selected smaller biped. If Scale Stride mode is off, the footsteps come into the current scene without being scaled down.

If you turn off Scale Stride mode and then go into Figure mode and scale the biped up or down, the footstep stride width and length remains the same when you exit Figure mode.

In Place Mode—Use In Place mode to keep the biped visible in the viewports while the animation plays. Use this for biped key editing or adjusting envelopes with Physique. It prevents XY movement of the biped center of mass during animation playback; however, motion along the Z axis is preserved. This is a three-button fly-out. In Place mode is stored with the 3ds Max file.

In Place X Mode—Lock center of mass X-axis motion. Use this for game export where the character stays in place but the swinging motion of the hips and upper body along the Y-axis is preserved.

In Place Y Mode—Locks center of mass Y-axis motion. Use this for game export where the character stays in place but the swinging motion of the hips and upper body along the X-axis is preserved.

Biped keys for limbs, footsteps, and center of mass can be adjusted using In Place mode. When the center of mass is moved on the XY-axes in this mode, the footsteps move. View biped playback without requiring a follow camera. In this viewing mode, visible footsteps "slide" under the biped.

For export to games, this feature is valuable since many game engines intelligently move the character's center of mass laterally according to game play. In Place mode makes it easy to view, tune, and export animation in a manner that is complimentary to game engine playback.

Tip: Another way of following a moving character is to link a camera and camera target to the center of mass shadow, which is the disc between the biped's feet.

Note: Trajectories do not display when In Place mode is active.

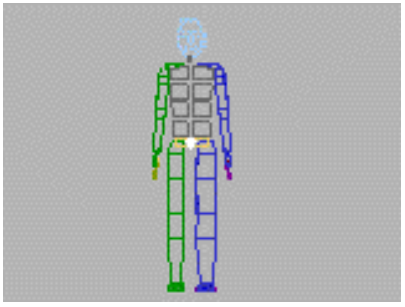
Display group



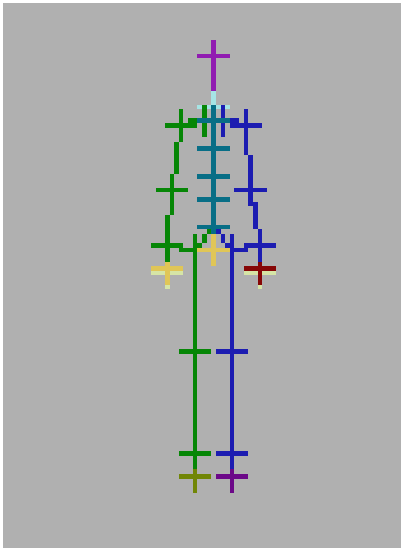
Note: The Display group is hidden by default. To display it, click on the Modes and Display expander in the Biped rollout.

Display Objects flyout—This flyout lets you display bones and objects, together or independently:

- **Objects**—Displays biped body objects; these will render if you do not turn them off before rendering. Hide the biped objects before scene rendering. You can also hide individual body objects by using the standard 3ds Max Hide controls found in the Display panel and Display Floater.



- **Bones**—Displays biped bones. Bones, which do not render, are represented as the color of the corresponding links. Displaying Bones is useful for seeing exactly where the joints fall in relation to the biped objects.



- **Objects and Bones**—Displays bones and objects simultaneously.

Display Footsteps flyout—This flyout lets you display or hide footsteps and their numbers:

- **Show Footsteps and Numbers**—Displays biped footsteps and footstep numbers.

  Footstep numbers specify the order in which the biped will move along the path created by the footsteps. Footstep numbers are displayed in white and do not render, but do appear in preview renderings.



- **Show Footsteps**—Displays biped footsteps in the viewport, but no footstep numbers.

  Footsteps are represented as green and blue foot-shaped outlines by default; these are also visible in preview renderings.

- **Hide Footsteps**—Turns off footsteps and footstep numbers in the viewport.

Twist Links—Toggles the display of twist links used in biped. Default=on.

Leg States—When this button is on, the viewport displays Move, Slide, and Plant at each foot at the appropriate frame.

Trajectories—Displays trajectories for selected biped limbs.

Tip: You can edit keys on the biped's horizontal and vertical track by turning on Trajectories, turning on Sub-Object, selecting the horizontal or vertical center of mass track and transforming keys in the viewports. Use Trajectories when editing keyframe parameters to visualize their influence, and to compare raw and filtered motion capture data.

Display Preferences—Displays the *Display Preferences dialog (page 2–944)* which is used to change footstep colors, trajectory parameters, and to set the number of bipeds to be played back when you use Biped Playback on the Biped rollout. Footstep color preference is a good way to distinguish between the footsteps of two or more bipeds in a scene.

Name Rollout

The Name rollout lets you change the name of the biped. When you change the name in the Biped rollout, the center of mass is renamed and the entire biped hierarchy inherits the new name.

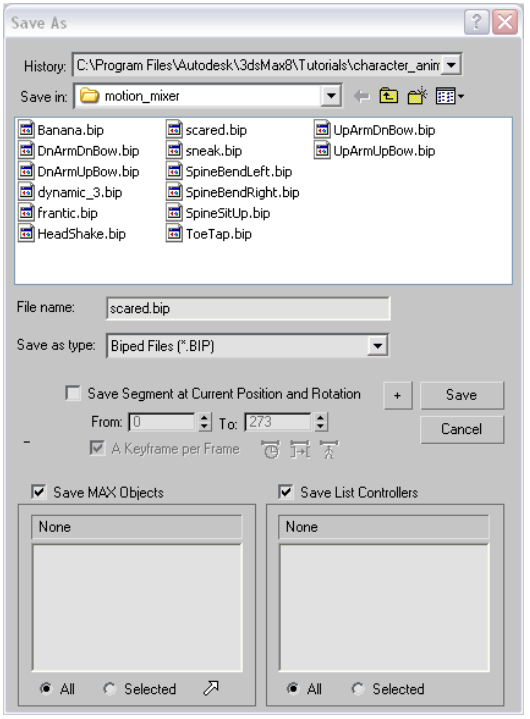
Save As Dialog (Biped)

Select a Biped. > Motion panel > Biped rollout > Save File

The Save As dialog lets you save biped (BIP), figure (FIG), and step (STP) files.

When saving FIG files (in Figure mode), the Save As dialog is a standard Windows Save File dialog. However, when you save BIP or STP files, you can save objects and animation controllers with the file.

Interface



Save As dialog for a FIG, BIP, or STP file

History—This drop-down list lets you choose directories where you have recently saved Biped motion files. See *Motion File History Lists (page 2–920)*.

File Name—The name of your FIG, BIP, or STP file.

Save as type—Sets the file type for your saved file.

-  Saves an incremented version of the segment without overwriting an existing file.
- **Figure file (.FIG)**—Saves the structure and position of a biped in Figure mode. After fitting the biped to a mesh in Figure mode, save a figure file. If the biped is accidentally moved in Figure mode, reload this file.
- **Biped file (.BIP)**—Saves a biped motion file. A BIP file includes footsteps and keyframe data. Biped files store the complete movement and allow you to create libraries of motion. Create your own BIP library by animating the biped and saving a BIP file.
- **Step file (.STP)**—Save footstep timing and location data in an ASCII format. Step files contain no body keys (spine, arms, etc.).

Save Segment—Lets you select segments of your animation for your BIP or STP file.

- **From:** Starts saving at this frame of the named animation.
- **To:** Stops saving at this frame of the named animation.
-  **Active Time Segment:** When on, animation in the active time segment is saved.
-  **Biped Animation Period** When selected, only time containing biped animation is saved.
- **Edited Transition Interval:** Saves frames in a transition only. Select a clip on the Motion Flow Script rollout, open the Transition Editor, and then select Save Segment to activate this control.

  Tip: Game developers can save only the motion in the transition.

A Keyframe per Frame—Saves a key at every frame. Turn this on if you plan to extract footsteps using

Load Motion Capture on the Motion Capture rollout.

Save MAX Objects group

Save MAX Objects—When on, any IK, Head Target, or linked object in the scene is saved with the BIP file.

Object list—Each object associated with the biped in the scene is displayed here along with its object type.

All—When on, all associated objects are saved with the BIP file.

Selected—When on, only the highlighted objects are saved with the BIP file.

Save List Controllers group

Save List Controllers—When on, any animation controllers assigned to the biped are saved with the BIP file.

Controller List—Each animation controller used with the biped in the scene is displayed here along with the object that it controls.

All—When on, all associated animation controllers are saved with the BIP file.

Selected—When on, only the highlighted animation controllers are saved with the BIP file.

Open Dialog (Biped)

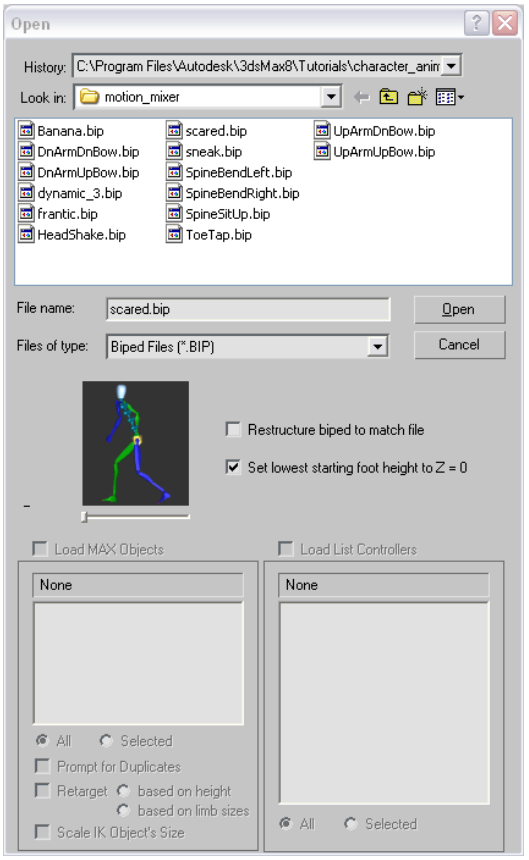
Select a biped. > Motion panel > Biped rollout > Load File

The Open dialog lets you load biped (BIP), figure (FIG), and step (STP) files.

When you load a FIG file (in Figure mode), the Open dialog is a standard Windows Open File dialog. However, when you load a BIP or STP file, you can preview the animation and control how much information is loaded into your biped.

Note: This option is unavailable when you load a BIP file into a clip or onto a biped that is in Edit Clip mode, because all the clips in the motion flow would have to be adapted.

Interface



Open dialog for BIP or STP files

History—This drop-down list lets you choose directories where you have recently opened Biped motion files. See *Motion File History Lists (page 2–920)*.

File Name—The name of your FIG, BIP, or STP file.

Files of Type—Sets the type of file to load.

- **Figure file (.FIG)**—Load a Figure file. Figure mode must be active to load a Figure file.

Figure files allow you to apply the structure of one biped to another. Reload a Figure file if you accidentally lose your biped Figure mode pose; this pose is the biped fitted to a mesh.

Note: When you load a FIG file, the Open dialog is a standard Windows Open File dialog.

- **Biped file (.BIP)**—Load a biped motion file (clip). Motion files include footsteps, keyframe settings, the biped scale, and the active gravity value (GravAccel). IK Blend values for keys and IK objects and their animation are also loaded.

  If you select a BIP file saved with a previous version of **character studio**, the preview will not be visible. After selecting the file for loading, you will be notified that the file is obsolete and should be resaved, unless you have indicated that you no longer want to see this message.

- **Step file (.STP)**—Load footsteps without body keyframes. Using this ASCII file format enables developers to write programs that generate step files for biped motion. Biped will generate body keys for the loaded steps. See *stp.rtf* in the *cstudio\docs* directory for more details on the step file format.

Note: In Figure mode, you can only load FIG files. In other modes, you can load BIP or STP files.

Motion Preview—Lets you manually scrub throught the animation in a BIP or STP file before opening it.

Restructure biped to match file—When on, the structure of the current biped is changed to match the structure in the BIP file.

Set lowest starting foot height to Z=0 (.bip files only)—Sets the lowest starting foot height to Z=0. This is an option in the Load File dialog. Default=On.

In **character studio**, the height of a motion clip can be retained. This is important if you want to retain the height of a motion clip for motions adapted to characters of different sizes. If, for example, the character is jumping off a rock and you want to retain the Z position of the character, you would turn this option off. Leave this option off if Motion Flow motions must be blended that begin and end at different heights, such as three clips that have the character mounting a bicycle, riding the bicycle, and dismounting the bicycle.

Turning off this option can, however, cause a jump in the motion during motion flow transitions. Turn this on for smooth transitions in Motion Flow mode. If adaptation takes place, the height is set so that the lowest foot at frame 0 starts at the Z=0 height. This lines up clips along the Z axis and creates smooth transitions.

Tip: Use Load Motion Capture File on the *Motion Capture rollout (page 2–1065)* to load the raw version of the motion capture *.bip* files that come bundled with **character studio**. These files have no footsteps and keys at every frame. Loading files using Load Motion Capture File allows you to filter the data and extract footsteps.

Load MAX Objects group

Load MAX Objects—When on, **character studio** will import any IK, Head Target, or linked object saved with the BIP data.

Object list—Each object saved with the BIP file is displayed in this list along with its object type. The only way to eliminate an object from the animation is to resave the BIP file without it and reload.

Note: The settings for loading 3ds Max objects are only active if Load MAX Objects is checked and there are objects in the data.

Prompt for Duplicates—Opens a Merge dialog for each object in the BIP file that has a duplicate in the scene.

In this dialog, you can Merge, Skip, or Delete Old.

If Prompt for Duplicates is off, the objects in your scene will be automatically overwritten.

Retarget—When on, the position of the objects in the BIP file adapt to the size of the target biped, based on either height or limb sizes.

Scale IK Object's Size—When on, the imported objects are scaled in size to match the target biped.

Load List Controllers group

Load List Controllers—When on, and sub-animation controllers saved with the BIP file are loaded onto your target biped.

Controller list—Displays the sub-animation controllers in the BIP file.

Note: The only way to eliminate a controller from the animation is to resave the BIP file without it and reload.

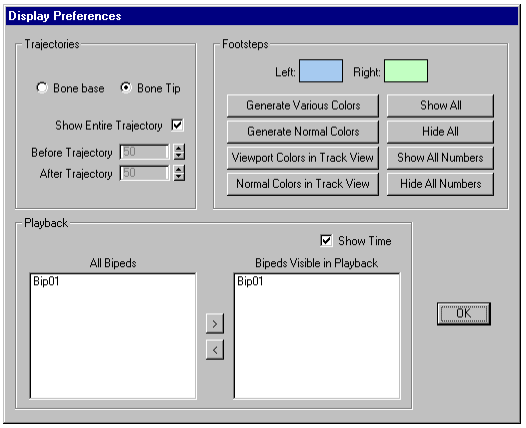
Display Preferences Dialog

Select a biped > Motion panel > Biped rollout > Display group> Display Preferences

Use the controls in the Display Preferences dialog to change footstep colors and trajectory parameters, and to set the number of bipeds to play back when you use Biped Playback on the Biped rollout.

You access the Display Preferences dialog by clicking the Display Preferences button on the Display group in the Biped rollout.

Interface



Trajectories group

Note: Trajectories do not display using In Place mode.

Bone base—Displays bone base trajectories.

Bone Tip—Displays bone tip trajectories (default).

Show Entire Trajectory—Displays trajectories for all animated frames.

Before Trajectory—Sets the number of frames to display trajectories before the current frame.

After Trajectory—Sets the number of frames to display trajectories after the current frame. Using Before and After Trajectory will result in a "traveling" trajectory display that will move with the biped through space.

Footsteps group

Left—Selects the color for left footsteps. Double-click the color swatch next to this selection to display the Color Selector and set the color for left footsteps.

Right—Selects the color for right footsteps. Double-click the color swatch next to this selection to display the Color Selector and set the color for right footsteps.

Generate Various Colors—Asks whether you want different colors for left and right footsteps. Based on your response, generates various colors for footsteps in the viewports. This works with multiple bipeds.

Generate Normal Colors—Changes right footsteps to blue and left footsteps to green, the default. Applies to all bipeds in the viewports.

Viewport Colors in Track View—Displays viewport footstep colors in Track View.

Normal Colors in Track View—Displays normal footstep colors in Track View.

Show All—Shows all footsteps.

Hide All—Hides all footsteps.

Show All Numbers—Shows all footstep numbers.

Hide All Numbers—Hides all footstep numbers. Show and hide are useful when there are multiple bipeds in the viewports.

Playback group

Controls in this group limit the number of bipeds to play back when you use Biped Playback on the Biped rollout on the Motion panel.

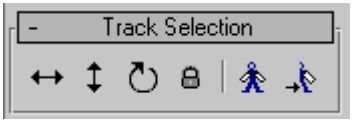
Show Time—Displays frame numbers in the viewport during playback.

All Bipeds—Lists all bipeds in the scene. Select a biped name in the window and click the right arrow to move it into the list of bipeds that will be visible during playback with Biped Playback.

Bipeds Visible in Playback—Lists bipeds visible during playback with Biped Playback. Select a biped name and click the left arrow to *exclude* it from this list.

Track Selection Rollout

Select a biped > Motion panel > Track Selection rollout



The Track Selection rollout give you specialized tools for manipulating the biped *center of mass (COM) (page 2–933)*. The center of mass object is the root of the biped hierarchy, and controls the entire biped structure. Changing its position and orientation affects the biped as a whole, which is an important part of posing and animating the biped. Also available on this rollout are tools for selecting symmetrical and opposite biped body parts.

The COM can be selected in different ways:

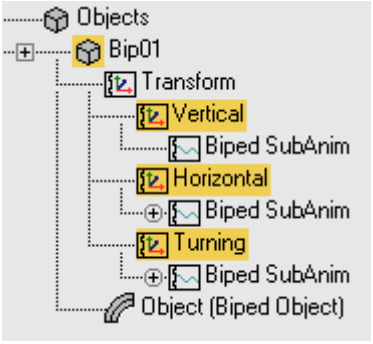
- If the *viewport rendering method (page 3–853)* is set to Wireframe, you can select the COM in your viewport, located near the center of the biped's pelvis.
-  You can turn on Move All Mode on the *Biped rollout (page 2–936)*, which enlarges the COM in your viewport. This is useful if the viewport rendering method is set to Smooth, and the biped's pelvis covers the center of mass.
- You can open the *Select Objects dialog (page 1–78)*, which displays all visible objects in your scene. The first listed biped item (*Bip01*, for example) usually refers to the root; you can select it from there.
- You have access to three tools on the Track Selection rollout (on the Motion panel) to edit the COM's position and orientation: Body Horizontal, Body Vertical, and Body Rotation. Choosing one automatically selects the center of mass.

Once the biped's center of mass is selected, you can move or rotate it using the *Transform gizmo*

(page 3–1025). Controlling the COM using the Track Selection rollout tools is sometimes quicker because selecting the center of mass in your viewports can be difficult if your biped resides inside a bigger rigged model. As long as you select your biped by any limb, activating the COM tools from the Track Selection rollout automatically selects the center of mass with the proper Transform gizmo:

- Body Horizontal and Body Vertical enable the Move gizmo.
- Body Rotation enables the Rotate gizmo.

When you key the COM's position or orientation using *Set Key (page 3–718)* or *Auto Key (page 3–717)*, the animation data is stored within the respective biped transform tracks.



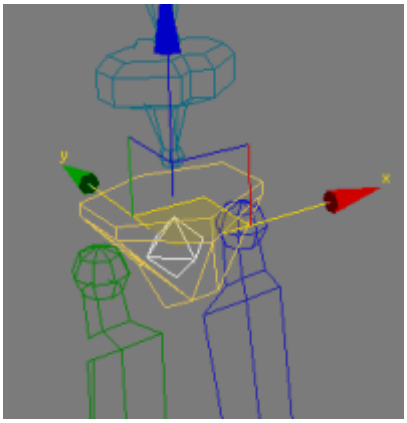
The biped Transform tracks contain the COM keyed data.

Procedures

To Edit the COM's position and orientation:

1. Create a biped.
2. On the Motion panel, expand the Track Selection rollout and turn on Body Horizontal.

   The Move Transform gizmo appears at the COM's location, highlighting the X and Y axes because they are referring to the biped Horizontal Transform track.



The Transform gizmo centered on the biped's COM.

Tip: To avoid deselecting the COM inadvertently, you can use *Lock Selection (page 2–555)* to keep your selection on the center of mass, ignoring all other objects in your scene.

Tip: You can scale the Transform gizmo using the - (hyphen) and = (equal sign) shortcut keys. Scaling the gizmo makes it easier to use, but does not affect the transform values.

3. Select the Move gizmo's Z axis.

   This turns off Body Horizontal and turns on Body Vertical.

4. Turn on Body Rotation on the Track Selection rollout.

   The Move Transform gizmo is replaced by the Rotate Transform gizmo. You can now change the biped's orientation.

To Lock all Biped COM Tools:

This procedure follows the concept from the previous procedure, and centers on selecting multiple COM tracks at the same time, as well as locking them.

1. On the Track Selection rollout, turn on Body Horizontal.

2.  Turn on Lock COM Keying.

   Lock COM Keying allows you to activate all center of mass tracks at the same time.

3.  Turn on both Body Horizontal and Body Rotation.



4. Select any other part of your biped.

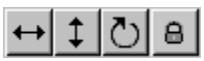
   The COM track controls become inactive. Note that Lock COM Keying is storing the selected controls.



5. Reselect the biped COM either by selecting the blue octahedron near the center of the biped's pelvis, or by clicking one of the three COM track controls.

   Lock COM Keying restores the selected controls.

Interface



Body Track Tools (Center of Mass Tracks)

Body Horizontal and Body Vertical refer to the center of mass's translation axes while Body Rotation focuses on its rotate axes. Once you turn them on, you can move or rotate the COM by either using the Transform gizmo, by entering values in the *Coordinate Display area (page 3–708)*.

Tip: You can also set COM positions using the Tension, Continuity and Bias values on the *Key Info rollout (page 2–954)* (in the TCB group).

When you activate Body Horizontal or Body Vertical, you have access to advanced *Biped Dynamics parameters (page 3–916)*, such as *Balance Factor (page 3–914)* and *Ballistic Tension (page 3–914)*, which you can control from the *Key Info rollout (page 2–954)* and *Dynamics & Adaptation rollout (page 2–980)*.

Note: Tension, Continuity, and Bias affect the COM Body Vertical keys only when the value of *Dynamics Blend (page 3–932)* is less than 1.0. To turn gravity off at a vertical center of mass key, set the value of Dynamics Blend to 0.0.

Tip: To view the way a parameter change affects the animation, Turn on *Trajectories (page 3–1025)*.

Keyed keys for COM tracks are color coded as follows:

- Keys containing Body Horizontal tracks are red.
- Keys containing Body Vertical are yellow.
- Keys containing Body Rotation are green.

  Note: If a frame contains multiple keyed COM tracks, the key's color code is divided accordingly.



A key's color reflects its keyed COM tracks.

 **Body Horizontal—**Selects the center of mass to edit horizontal biped motion.

The Body Horizontal track has a *Balance Factor (page 2–877)* parameter that automatically orients the biped to maintain balance. This saves the animator from having to reposition the pelvis when the biped leans forward, backwards, or sideways.

Note: *Biped Dynamics (page 3–916)* parameters can be animated from no effect to full effect at each keyframe.


Body Vertical—Selects the center of mass to edit vertical biped motion.

The Body Vertical track uses the Dynamics Blend parameter to control gravity in a footstep animation. A Dynamics Blend value of 1.0 uses the value of *GravAccel (page 3–950)* (global gravity value) to calculate an airborne trajectory for the biped. No keyframes are required to position the biped in the air, a trajectory is calculated automatically. A value of 0.0 uses Spline Dynamics for the vertical position of the biped; you must create keyframes to position the biped vertically.

The Body Vertical track also has a *Ballistic Tension (page 3–914)* parameter that controls how much the biped knees bend when the biped lands from an airborne period. This means that keys do not need to be created at the lowest position of the biped after landing; a trajectory is calculated automatically.

Note: You can create keys manually to override the calculated trajectory during the landing period. However, vertical keys must have Dynamics Blend=0.0 in order to fully override the trajectory during the airborne period.


Body Rotation—Selects the center of mass to edit biped rotational motion.


Lock COM Keying—When on, allows you to select multiple COM tracks at the same time. Once locked, the tracks are stored in memory, and are remembered every time the COM is selected.

Note: Locked COM tracks are also honored after a scene reset or a session change.


Symmetrical—Selects the matching object on the other side of the biped. For example, if the right arm is selected, clicking Symmetrical Tracks selects the left arm too. You can then make changes to both sides of the body at once. Symmetrical works for single and multiple biped parts.


Opposite—Selects the matching object on the other side of the biped, and deselects the current object. For example, if the right arm is selected, clicking Opposite Tracks selects the left arm and deselects the right arm. Opposite Tracks can be used for single or multiple objects.

See also

Biped Color-coded Keys and Trajectories (page 2–1005)

Quaternion/Euler Rollout

Select the biped > Motion panel > Quaternion/Euler rollout

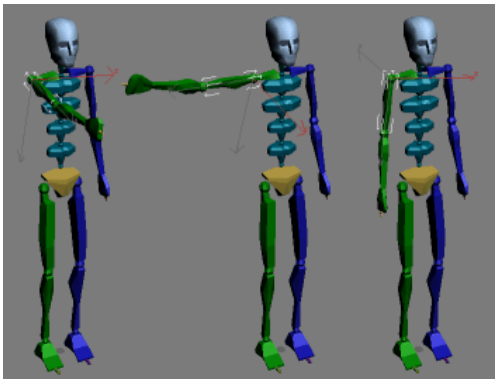
This rollout contains options to switch between Euler or quaternion controllers on biped animations. These choices offer alternative ways to control an animation in the Curve Editor. For more details, refer to *Working with Euler Curves on Biped Animation (page 2–916)*.

Procedures

To switch between quaternion and Euler rotation controllers:

1. Create a biped.
2. 

Turn on Auto Key.
3. 
Select the biped's upper arm and rotate it so it extends forward.
4. 
Go to frame 20 and rotate the upper arm so it extends sideways.
5. 
Go to frame 40 and rotate the upper arm back to its original orientation.

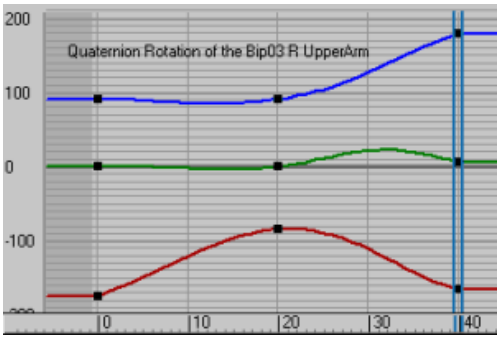


From left to right, the biped at frame 0, 20, and 40.

In the Motion Panel, on the Quaternion/Euler rollout, notice that the Quaternion option is active by default.

Note: If you change options for that biped limb, your choice is preserved in the *3dsmax.ini (page 1–18)* file, from which it is restored after a scene reset or session change.

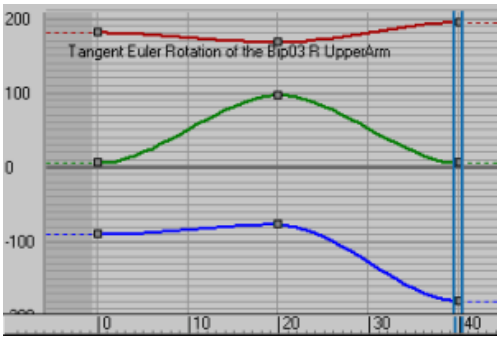
6. On the Biped Apps rollout, click Workbench to open it.



The animation is displayed as three separate TCB rotation curves labeled "Quaternion Rotation".

7. On the Quaternion/Euler rollout, choose the Euler option.

   The curves have been converted to display Euler controlled rotations. You now have access to the curve's tangent handles to change the curve's interpolation.



The animation is displayed as three separate Euler curves labeled "Tangent Euler Rotation".

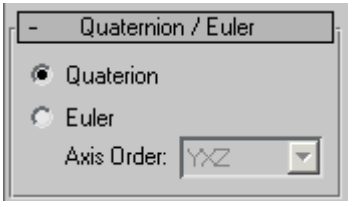
8. Choose a different axis ordering from the Axis Order drop-down list (under the Euler option).

   The curve display changes to reflect the new ordering while preserving the biped posture.

9. Choose the Quaternion option to convert the curves back to a TCB interpolation.

   Note: Converting a curve from Euler to quaternion can affect its interpolation. See *Working with Euler Curves on Biped Animation (page 2–916)* for details.

Interface



The following options set the rotation controller for biped animations.

Quaternion — Converts the selected biped animation to quaternion rotation. If the biped selection has not been animated yet, this starts as the default option.

If you change options for any biped limb, your choice is preserved in the *3dsmax.ini* file, from

which it is restored after a scene reset or session change.

Euler — Converts the selected biped animation to Euler rotation.

Axis Order — Lets you choose the order in which the Euler rotation curves are calculated. Availabe only when Euler is active. Default=YXZ.

Changing the ordering reflects on the curves, which are modified accordingly. However, this change does not affect the animation itself.

Note: This setting is independent of the Workbench Axis Order setting.

Twist Poses Rollout

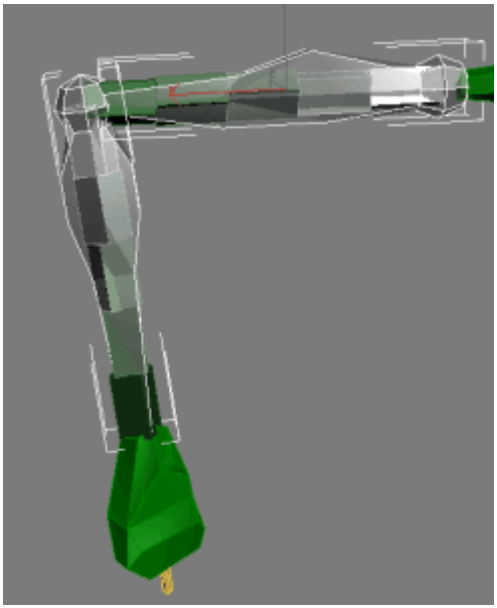
Select a biped. > Motion panel > Twist Poses rollout

The toolset of this rollout lets you create and edit twist poses for a biped's limbs. You can either use the rollout's preset poses (which you can modify and rename), or create your own. When you add a new pose, you establish a reference between the selected's limb relative orientation and the rollout's Twist value of 0 (which means no twist). When more than one reference pose are set for a rotated limb, the limb's twisting is calculated based on the nearest pose.

Tip: It is suggested that you add twist poses for all major limb rotated poses. This ensures that rotating your biped's limb produces well-calibrated twisting.

Note: You don't technically have to enable *twist links (page 2–987)* to use the Twist Poses rollout toolset. However, if your biped does not contain twist links, the Twist and Bias settings are not used.

The Twist Poses rollout tools only affect limbs with three degrees of freedom (DOF), such as upper arms and thighs, because you can control their twist links by rotating them. Two-DOF limbs (forearms and calves, for instance) differ due to the fact that you can only control their twist links if you rotate their child limb.



The upper arm can twist itself while the forearm needs the wrist to twist it.

Note: If you select a two-DOF limb, the only available setting is the Bias value; you can change it to adjust the rotation distribution along the twist links.

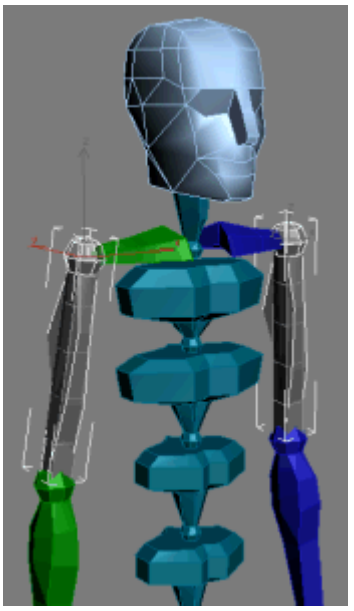
Procedures

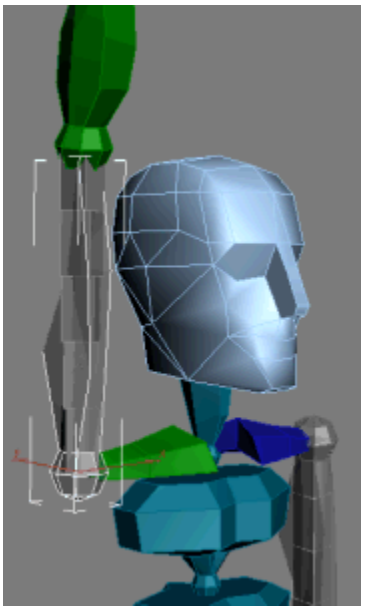
Example: To Add and Edit Twist Poses:

This procedure takes into account the concept of adding *twist links (page 2–987)* to biped limbs and concentrates on using the tools from the Twist Poses rollout to set up basic poses.

1. Prepare a biped with five twist links on each upper arm.

   Tip: You can use *See-Through (page 1–55)* on the biped's limbs to better distinguish the twist links inside.





2.  Exit Figure Mode, and then expand the Twist Poses rollout.
3. Select the right upper arm.

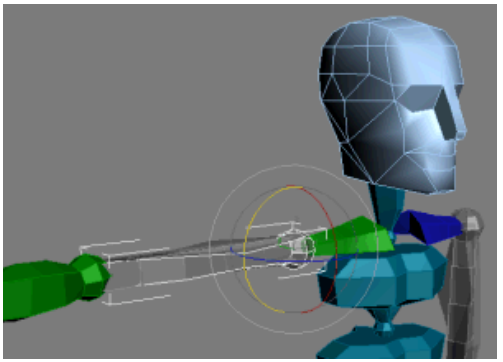
   The rollout controls are enabled because you selected a three-DOF limb.
4.  Use Previous Key and Next Key to cycle through the different pose presets.

   Tip: You can often use these presets as a starting point for posing your limbs.
5. Choose *pose1* from the drop-down list.

   The biped's arm extends upwards and shows some twisting.
6. Click Set.

   This assigns the Twist value of 0 to the upper arm's twist links.

   Note: A twist change is always reflected on both sides (in this case, both upper arms).
7. Rotate the upper arm locally around its Y axis so it extends on the side.

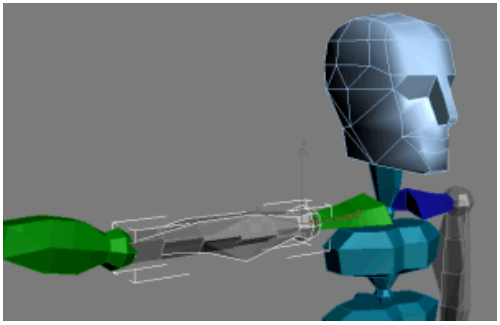


8. Click Add to create a new pose. Rename it **Arm At Side**.

This new pose resets the current Twist value of the twist links.

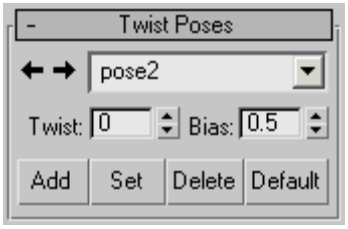
9. Rotate the upper arm locally around its X axis.

   The twisting in the upper arm is calculated based on the limb's proximity to the saved poses.



Rotate the upper arm to verify the twist links.

Interface



Previous/Next Key—Scrolls and selects through the list of twist poses.

Twist Poses List—Lets you choose a preset or saved pose to apply to the biped's selected limb. By default, five twist poses are available for each three-DOF limb: Up, front, lateral, down, and back. You can also rename the current twist pose.

Twist—Sets the amount of twist rotation (in degrees) applied to the twist links linked to the selected limb. The twist links from the opposite side are so affected. Default=0. Range=–180 to 180.

Note: Changing the Twist value automatically resets the current limb's orientation to the active twist pose.

Bias—Sets the distribution of rotation along the twist links. A setting of 1.0 concentrates the twist towards the top link while a setting of 0.0 instead concentrates it towards the bottom link. The default setting is 0.5, distributing the rotation evenly throughout the links. The twist links from the opposite side are so affected.

Note: Changing the Bias value automatically resets the current limb's orientation to the active twist pose.

Note: You can also set a Bias value for Two-DOF limbs.

Add—Creates a new twist pose based on the selected limb's orientation and resets Twist and Bias to their default values.

Note: Adding a new pose for one limb automatically makes it available for the limb on the opposite side.

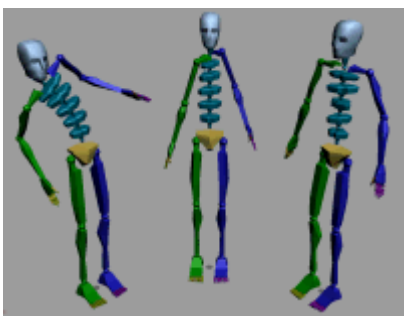
Set—Updates the active twist pose with the current Twist and Bias values.

Delete—Removes the current twist pose.

Default—Replaces all twist poses of all three-DOF limbs with five default preset poses.

Bend Links Rollout

Select the biped > Motion panel > Bend Links rollout



This toolset combines the Bend Links Mode, previously located on the Biped rollout, with other tools to allow an easier control over a chain link, such as a biped spine, neck or tail.

Note: Activating one of the Bend Links rollout modes deactivates any of the others. However, clicking Zero Twist or Zero All maintain any mode currently active.

Note: The Bend Links rollout is displayed in Mixer Mode, Motion Flow Mode or Footstep Modes.

Procedures

To bend a spine naturally using Bend Links Mode:

1. Select any link in the biped's spine.



2. Rotate the spine link. Notice that all chain links rotates the same way.
3.  On the Bend Links rollout, click Bend Links Mode to activate it.
4. Rotate the spine link. The other links in the chain rotates to match the single link's rotation, creating a natural bend along the spine.

To twist a spine naturally using Twist Links Mode:

1. Select any link in the biped's spine.
2.  On the Bend Links rollout, click Bend Links Mode to activate it.
3. Rotate the spine link until you reach a pronounced bend.
4. Rotate the spine link in local X. Notice that all chain links rotates in a weird way due to the fact that they do not maintain their relationship with the two other axis. Undo the rotation.
5.  On the Bend Links rollout, click Twist Links Mode to activate it.
6. Rotate the spine link in local X. The rotation averages up the spine for all links while the spine maintains its existing profile.

To twist a spine link without affecting the chain using Twist Individual Mode:

1. Select the bottom link in the biped's spine, just above the pelvis.
2.  On the Bend Links rollout, click Bend Links Mode to activate it.
3. Rotate the link so the spine's shape follows a noticeable curvature.
4.  On the Bends Links rollout, click Twist Individual Mode to activate it.
5. Select a link in the middle of the spine and rotate it in local X. The rotation only affects the selected link.

To smooth out a spine using Smooth Twist Mode:

1. Select the bottom link in the biped's spine.
2.  On the Bends Links rollout, click Twist Individual Mode to activate it.
3. Rotate the link 45 degrees in local X.
4. Select and rotate the spine's top link –45 degrees in local X
5.  On the Bends Links rollout, click Twist Individual Mode to activate it.
6. Rotate either the bottom or top spine link. The chain links rotates to smooth out the orientation difference between the two link

extremities. Adjust the Smoothing Bias control to distribute the chain's rotation towards the base or top link.

Interface



Bend Links Mode—This mode can be used to rotate multiple links of a chain without having to select all of them beforehand. Bend Links Mode transfers the rotation of one link to the other links, following a natural curvature.

Note: Using Bend Links Mode in Auto Key mode or setting a key after twisting results in keys on all the links of the selected chain. See *Separate Tracks (page 2–962)* for further details.

Twist Links Mode—Similar to Bend Links Mode, this mode takes the rotation in local X applied to the selected link and increment it equally throughout the rest of the chain while maintaining the relationships between the links in the other two axes.

Note: Using Twist Links Mode in Auto Key mode or setting a key after twisting results in keys on all the links of the selected chain. See *Separate Tracks (page 2–962)* for further details.

Twist Individual Mode—Similar to Bend Links Mode, this mode allows a selected chain link to be rotated in local X without affecting its parent or child. The chain therefore maintains its shape while the individual link is adjusted.

Note: Using Twist Individual Mode in Auto Key mode or setting a key after twisting results in keys on all the links of the selected chain. See *Separate Tracks (page 2–962)* for further details.

Smooth Twist Mode—This mode takes into account the orientation in local X of the chain's first and last links in order to distribute the rotation of the other links. This results in a smooth rotation between every chain link. The rotation distribution can be set by adjusting the Smoothing Bias control or by rotating either the first or last link of the chain.

Smoothing Bias—Sets the rotation distribution based on a value between 0.0 and 1.0. 0.0 biases towards the first link and 1.0 towards the last link of the chain. The chain's smoothness can be interactively set by dragging the bias setting between those 2 values.

Zero Twist—Resets every chain link's rotation to 0 in local X based on the current orientation of the chain's parent. This does not change the current shape of the chain.

Zero All— Resets every chain link's rotation to 0 in all axes based on the current orientation of the chain's parent. This adjusts the current shape of the chain so it becomes parallel to the biped.

Key Info Rollout

Select the biped. > Motion panel > Key Info rollout

Tools in the Key Info rollout allow you to do the following:

- Find the next or previous key for the selected biped body part.
- Use the Time spinner to slide a key back and forth in time.
- Change Tension, Continuity, and Bias for a key and display trajectories.
- Adjust biped dynamics.
- Set planted, sliding, or free keys.

- Set IK constraints and pivots for the biped hands and feet.

When the *Body Vertical* COM (center of mass) track is active, you can change the vertical dynamics of the motion, on a key-by-key basis. When the Body Horizontal COM track is active you can change the balance factor for shifts in weight distribution.

Note: On the time slider or in Track View, you can move one Biped key past another. See *Moving Keys (page 2–1004)*.

Activating Parameters

Groups of the Key Info rollout are unavailable depending on what part of the biped is selected and if a key is current. *Body Vertical*, *Body Horizontal*, and *Body Rotation* refer to the three tracks used to animate the biped center of mass. Select one of the three center of mass tracks on the Track Selection rollout, then use Next Key or Previous Key to find a key to edit.

- If Body Vertical is active and a key is current, then parameters for Dynamics Blend, Ballistic Tension, Z Position, Time, and TCB parameters (Tension, Continuity and Bias) are active. Ballistic Tension is available only at keys just before or just after an airborne state, as between footsteps in a run or jump.

  Note: TCB controls are not effective at Body Vertical keys just before and just after an airborne period, between footsteps, if Dynamics Blend=1. Biped Dynamics calculates the airborne trajectory; in this case, lower the value of Dynamics Blend to use the TCB controls. In a walk sequence where footsteps overlap, Dynamics Blend has no effect and TCB controls can be used.

- If Body Horizontal is selected and a key is current, the Balance Factor parameter, XY Position, Time, and TCB parameters are active. Z Position, Dynamics Blend, Ballistic Tension are grayed.

- If Body Rotation is selected and a key is current, only the Time and TCB parameters are active.

- If a biped hand is selected and a key is current, then all parameters are active except parameters in the Body Dynamics group.

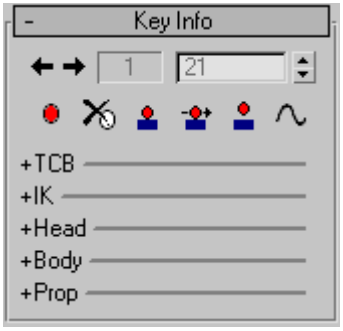
- If a biped foot key is selected and current, then all parameters are active except for parameters in the Body Dynamics group.

- If a biped leg is selected and a key is current, then Time and TCB parameters are active. XYZ Position and Body Dynamics parameters are made unavailable. In a footstep animation, Time is made unavailable at a Touch and Lift key.

- If a biped arm is selected and a key is current, then Time and TCB parameters are active. XYZ Position and Body Dynamics parameters are made unavailable.

Interface

The Key Info rollout is divided up into several groups: TCB, IK, Head, Body, and Prop. You can expand and hide each of these groups by clicking on the line next to its name.



Next Key-Previous Key—Find the next or previous keyframe for the selected biped part.

The field displays the key number.

Time—Enter a value to specify when in time the key occurs.

Use this to fine tune keyframe timing on a character by moving a key backwards and forwards in time.

 **Set Key—**Creates keys at the current frame when you are moving biped objects. This is identical to Set Key on the 3ds Max toolbar.

You can experiment with different biped poses without updating the motion until you find the desired pose. You can also quickly fine tune your motion by setting a key and adjusting the key parameters on the Key Info rollout without having to transform the biped in the viewports.

Note: If a biped key is current, then TCB, XYZ position spinners, and IK Blend parameters can be updated without using Set Key or having the Auto Key button turned on.

Note: If 3ds Max bones using the IK Controller or 3ds Max Particle Emitters are linked to the biped, or if you are displaying 3ds Max trajectories or ghosting, the Auto Key button *must* be on while the biped is positioned. These objects update their parameters in real time as they are positioned.

 **Delete Key—**Deletes the key of the selected object at the current frame.

By default, biped arm, hand, and finger keys are stored in the clavicle track. If you delete keys for any one of these objects, you lose positions for the rest of the arm objects at that frame. If you plan on extensive hand animation, turn on Arms in the Separate Tracks group of the Keyframing Tools rollout. This creates separate tracks for each biped arm object. Deleting an upper arm key will preserve hand and finger keys.

 **Set Planted Key—**Sets a biped key with IK Blend=1, Join To Previous IK Key turned on, and Object selected in the IK group.

In a Footstep or Freefrom animation, all footsteps that do not slide should have Join To Previous IK Key turned on.

Tip: Right-click to open the *Planted Key Defaults dialog*, which lets you set new default values for subsequent TCB keys.

 **Set Sliding Key—**Sets a biped key with IK Blend=1, Join To Previous IK Key turned off, and Object selected in the IK group. This creates a sliding footstep. Sliding footsteps display in the viewports with a line running through the middle of the footstep. Sliding footsteps are understood as footsteps with moving IK constraints.

In a Footstep or Freeform animation, if the foot slides rather than being planted, use Set Sliding Key.

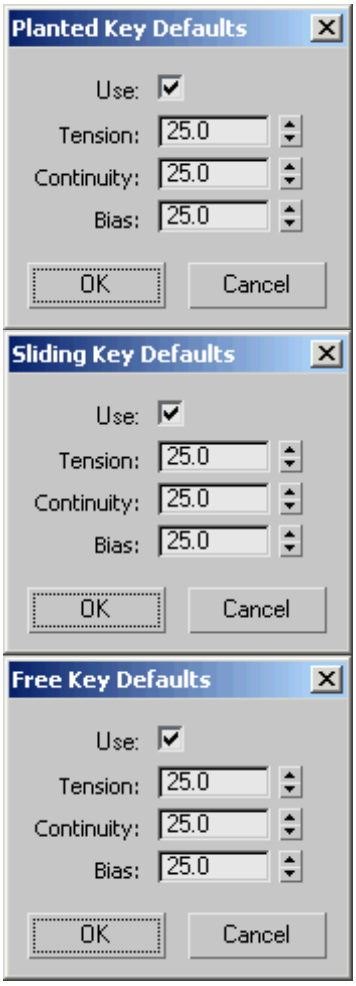
Tip: Right-click to open the *Sliding Key Defaults dialog*, which lets you set new default values for subsequent TCB keys.

 **Set Free Key—**Sets a biped key with IK Blend=0, Join To Previous IK Key turned off, and Body selected in the IK group.

In a Footstep or Freeform animation, a biped leg in a move state should have a "free" key.

Tip: Right-click to open the *Free Key Defaults dialog*, which lets you set new default values for subsequent TCB keys.

To change the default values for TCB keys, right-click either Set Planted Key, Set Sliding Key, or Set Free Key to open the respective dialog.



Use—When on, subsequent keys set with either Set Planted Key, Set Sliding Key, or Set Free Key use the TCB values from the respective dialog. Otherwise, the system uses the default values.

Tension—Controls the amount of curvature in the animation curve. Default=25.

Continuity—Controls the tangential property of the curve at the key. Default=25.

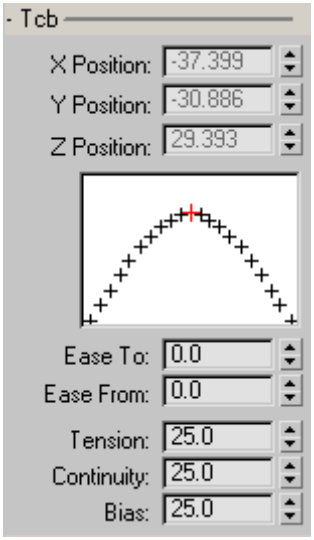
Bias—Controls where the animation curve occurs with respect to the key. Default=25.

Trajectories—Shows and hides trajectories for the selected biped object. You can edit keys on the biped's horizontal and vertical track by turning on Trajectories, turning on Sub-Object, selecting the horizontal or vertical center of mass track, and transforming keys in the viewports.

- You can bend the horizontal center of mass trajectory around selected horizontal keys by using the Bend Horizontal spinner in the Keyframing Tools rollout.
- Display trajectories to view how parameter changes in the Key Info rollout affects the biped motion. Changing Tension, Continuity, and Bias in the Tcb group affects the trajectory around the current key. Changing the value of IK Blend for a hand or foot will affect the trajectory between keys.
- Leave Trajectories on and turn on Show Buffer Trajectories on the Motion Capture rollout to compare a raw motion capture trajectory with the filtered trajectory on the biped. This assumes a motion capture file has been loaded.
- Changing Dynamics Blend for a center of mass vertical key or changing the value of GravAccel will change gravity in a foostep animation and will therefore affect the trajectory.

Simple—Collapses the lower part of the Key Info rollout, for simplified viewing.

TCB group



You can use the TCB controls to adjust easing and trajectories on keys that already exist.

XYZ Position—Reposition the selected biped part using these spinners.

A hand or foot can be repositioned in world coordinate XYZ. The biped center of mass can also be positioned using these spinners.

TCB Graph—Charts the effect that changing the controller properties will have on the animation. The red mark at the top of the curve represents the key. The marks to the left and right of the curve represent an even division of time to either side of the key.

The TCB graph is a stylized representation of the animation around a single key.

Ease To—Slows the velocity of the animation curve as it approaches the key. Default=0.

High Ease To causes the animation to decelerate as it approaches the key.

The default setting causes no extra deceleration.

Ease From—Slows the velocity of the animation curve as it leaves the key. Default=0.

High Ease From causes the animation to start slow and accelerate as it leaves the key.

The default setting causes no change of the animation curve.

Tension—Controls the amount of curvature in the animation curve.

High Tension produces a linear curve. It also has a slight Ease To and Ease From effect.

Low Tension produces a very wide, rounded, curve. It also has a slight negative Ease To and Ease From effect.

The default value of 25 produces an even amount of curvature through the key.

Continuity—Controls the tangential property of the curve at the key. The default setting is the only value that produces a smooth animation curve through the key. All other values produce a discontinuity in the animation curve causing an abrupt change in the animation. Default=25.

High Continuity values create curved overshoot on both sides of the key.

Low Continuity values create a linear animation curve. Low continuity creates a linear curve similar to high tension except without the Ease To and Ease From side effect.

The default setting creates a smooth continuous curve at the key.

Bias—Controls where the animation curve occurs with respect to the key. Default=25.

High Bias pushes the curve beyond the key. This produces a linear curve coming into the key and an exaggerated curve leaving the key.

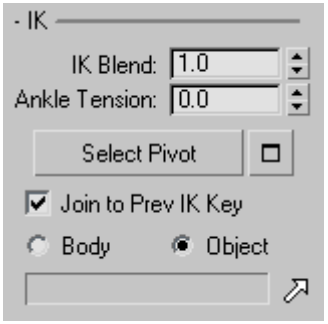
Low Bias pulls the curve before the key. This produces an exaggerated curve coming into the key and a linear curve leaving the key.

The default setting distributes the curve evenly to both sides of the key.

IK group

This group lets you set IK keys and adjust parameters for IK keys.

Note: You can now make and adjust settings for multiple selected IK-capable biped parts. For example, say at the same frame one foot has a planted key and the other has a sliding key. This means they both have (by default) IK Blend values of 1.0 and Ankle Tension values of 0.0. If you select both feet, all of the Key buttons near the top of the Key Info rollout are available, and you can change IK Blend and Ankle Tension values for both.



IK Blend—Determines how the software mixes forward kinematics and inverse kinematics to interpolate an intermediate position. An example of forward kinematics is moving the arm to control the hand. An example of inverse kinematics is moving the hand to control the arm.

Activates when a biped arm or leg (hand and foot) key is current.

- 0 with *Body* chosen is normal biped space (forward kinematics).
- 1 with *Body* chosen is inverse kinematics, which creates more straight-line motion between biped keys.
- 1 with *Object* chosen, but no *IK Object specified*, puts the limb fully into world space.
- 1 with *Object* chosen *and* an *IK Object specified* puts the biped limb into the coordinate space of the selected object; the biped limb follows the specified object.

Ankle Tension—Adjusts the precedence of the ankle joint over the knee joint. When set to 0, the knee takes precedence. When set to 1, the ankle takes precedence.

This effect is only visible between keyframes.

Select Pivot—Activate to designate pivots around which the biped hands and feet should rotate. After clicking a pivot in the viewports, turn off Select Pivot and then rotate the hand or foot.

Select Pivot is available only if the biped hand or foot is in World or Object *coordinate space (page 1–443)*.

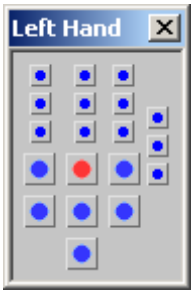
Pivot Selection Dialog—Opens a small dialog that shows the current pivot for the selected limb on the limb's respective hand or foot, and lets you change it. When the chart is green (or blue) and red, indicating that you're on an IK key, the red dot indicates the current location of the pivot. To designate a different pivot, click another dot on the chart. This provides an alternative method to using Select Pivot (see preceding).

The dialog is named according to the displayed hand or foot, depicts the actual number of digits and joints in use, and resizes itself accordingly. The displayed chart uses three different schemes, depending on the context:

- green/blue and red dots for the right/left hand or foot when on an IK key (pivot is editable)
- gray and white dots when in an IK period but not on a key (pivot is visible but not editable)
- no dots if you're in an FK (forward kinematics) period, or no limb or multiple limbs are selected (dialog disabled; no relevant data exists)

These are shown in the following illustration:

Note: When on an IK key, the left hand and foot charts use a blue/red color scheme:



Join to Previous IK Key—When on, places the biped foot in the coordinate space of the previous key. Turn off to put the biped foot into a new reference position.

Turn off and move the biped foot to create a sliding footstep, for example.

Body—The biped limb is in biped coordinate space.

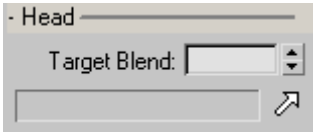
Object—Object Space: the biped limb is either in World coordinate space or the coordinate space of the selected IK object. Coordinate space can be blended between keys.

Select IK Object—Chooses an object for the biped's hand or foot to follow when IK Blend is 1 and Object is selected. The selected object's name is displayed next to the button.

This selection cannot be animated; only one IK object can be active for each hand and foot throughout the animation.

Tip: If you want to change the object the hand or foot follows at different times in the animation, select a dummy object as the IK object, and assign a Link Constraint to the dummy object to make it link to different objects at different times.

Head group



The Head group lets you define a target object for the target to look at.

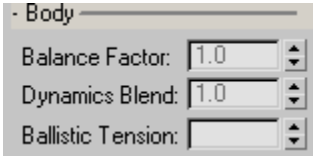
Target Blend—Determines the extent to which the target blends with the head's existing animation.

A setting of 1.0 causes the head to look directly at the target, 0.5 causes the head to blend half of its existing animation with looking at the target, and a setting of 0.0 causes the head to ignore the target, maintaining its existing animation.

Select Look At Target—Click to select an object for the head to look at.

Note: When an object is selected, its name is displayed next to the button.

Body group



These parameters apply to the biped center of mass and are used by **character studio** to calculate the biped's airborne trajectory based on gravity (GravAccel) and time between footsteps, the amount of knee bend on landing (Ballistic Tension) and how the biped objects adapt to maintain balance (Balance Factor).

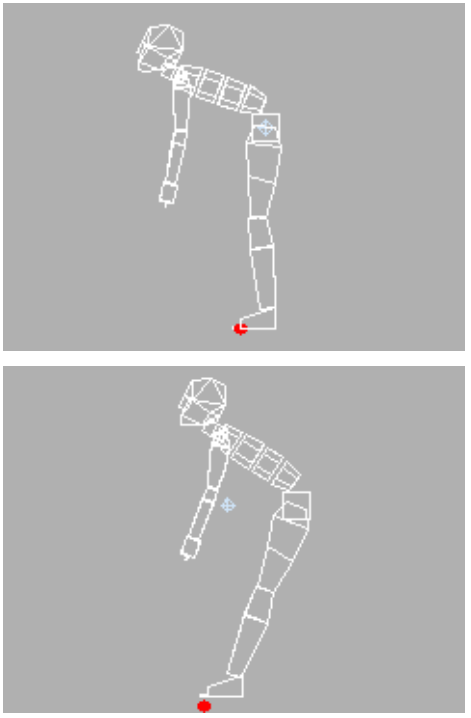
Balance Factor—Position the biped's weight anywhere along a line that extends from the center of mass to the biped's head. This center of mass (Body Horizontal track) parameter can be keyframed. To activate Balance Factor select the Horizontal Track (in the Track Selection rollout),

set a key and enter a value in the Balance Factor field.

For example, to create a *sit then walk* sequence, you could shift the biped's weight (balance) between 0 (the character is supported by the chair) for the sit key and 1 for the stand key (the character's pelvis shifts to maintain balance).

If a character is seated, and reaches across the table, leave Balance Factor at 0; however the character leans, he will pivot from the center of mass. The pelvis will not move back to maintain balance.

In a walking motion, a value of 2 will swing the hips and keep the biped head steady; a value of 0 will keep the hips steady and swing the upper body.



A Balance Factor value of 0 in the first image causes the biped not to compensate for weight.

A Balance Factor value of 2 in the second image causes the biped pelvis to move away from the Center of Mass to compensate for weight.

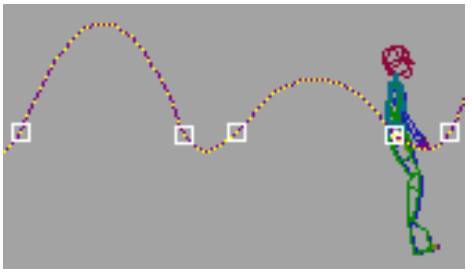
The Balance Factor determines how far the biped's hips will shift forward or backward to compensate for forward or backward bending of the spine. When the biped has a normal weight distribution between the upper and lower body, the default value of 1 causes the hips to swing backward as the biped bends over to compensate for the forward weight.

At times, when the biped leans, you will want the biped's hips to refrain from shifting to compensate for the forward weight. This would be when the biped is sitting or falling down. A Balance Factor of 0 (the minimum value) causes the hips to stay still when the biped leans forward or backward.

A value of 0 places the biped's weight at the center of mass. A value of 1 places the biped's weight above the center of mass. A value of 2 places the biped's weight in the head. Click the spinner up arrow to move the biped's weight distribution toward the head. Range=0.0 to 2.0; Default=1.

Note: You can also shift the center of mass by turning on Figure mode, selecting the center of mass object, and using *Rubber Band mode (page 2–936)* to move the center of mass to a new position. This method cannot be keyframed, but it allows you to move the center of mass outside the biped body. For example, you can simulate pushing a heavy object by moving the center of mass behind the biped.

Dynamics Blend—Select the Body Vertical track (center of mass vertical track) and control the amount of gravity in an airborne period, as in a running or jumping motion. This parameter has no effect on a walking motion where footsteps overlap.



The vertical center of mass track Dynamics Blend parameter is set to 1 for both keys (white squares) in the first jump, and to .5 in the second jump.

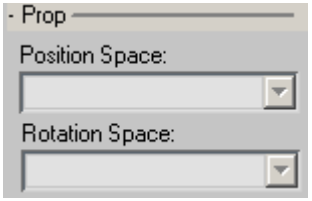
A value of 1 uses the *GravAccel (page 2–980)* value to calculate gravity. A value of 0 removes the effects of gravity calculation and flattens a jumping or airborne motion.

Ballistic Tension—Select the Body Vertical track (COM) and control the amount of spring or tension when the biped lands or takes off from a jump or run step. The change is subtle.

A walk cycle will not activate this value. The biped has to be airborne, *then* the Lift and Touch vertical keys will display a Ballistic Tension value.

If there are more than three vertical keys during a support period, you can also edit Ballistic Tension for the lift-off key; otherwise Biped uses the same value for touchdown and lift-off, since it is assumed that there is only one vertical dip in the motion. Low values are high tension (less dip in the trajectory). Default=0.5; Range=0 to 1;

Prop group



The Coordinate Space pulldowns will set the Prop to refer to the World, Body, Right Hand or Left Hand coordinate space for position and rotation at the current frame.

Position Space—Lets you set the prop position space to World, Body, Right Hand, or Left Hand.

Rotation Space—Lets you set the prop rotation space to World, Body, Right Hand, or Left Hand.

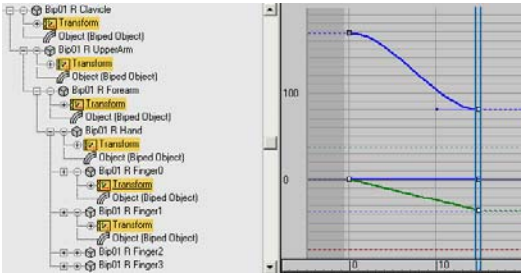
Keyframing Tools Rollout

Select the biped. > Motion panel > Keyframing Tools rollout

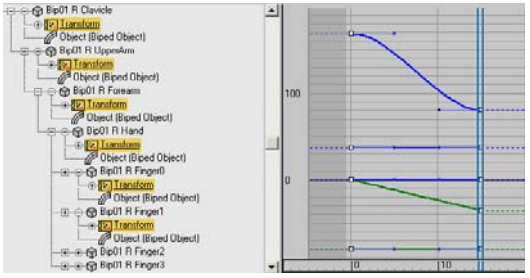
Use controls on the Keyframing Tools rollout to clear animation on a biped or selected parts, mirror biped animation, and cause the Neck to rotate in body space rather than parent space. You can also bend the horizontal center of mass track around a selected horizontal key.

Separate Tracks

By default, 3ds Max uses an optimized method for key storage. For example, keys for the fingers, hand, forearm, and upper arm are stored in the Clavicle transform track. If you prefer instead to have a transform track available for each arm object, use the Separate Tracks group to make these transform tracks available; transform tracks are displayed in Track View.



When the Separate FK Tracks options are on for the Arms and Fingers, only these tracks receive keys.



When the Separate FK Tracks options are off, all limb tracks receive keys.

Tip: Separate Tracks are designed to be used with forward-kinematic (FK) rotations. Using the Move tool to change the position of a biped limb with Separate Tracks requires that you turn on Set Parents Mode. If Set Parents Mode is off, then the limb you move is stored relative to its parent objects, and the new parent positions are *not* stored. The parent objects will snap back to their originally stored location. Turning on Set Parents Mode ensures that the position of the entire biped limb is stored.

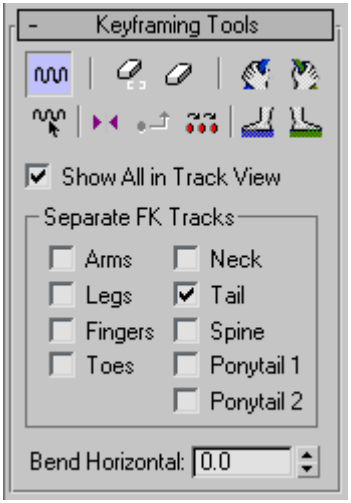
Note: Separate Tracks are intended to be used as a preference. It is better to turn them on once for the tracks you plan to edit by hand. Turning off Separate Tracks that are currently on causes a key to be stored for every biped object in the limb for the frames where any biped keys previously existed.

Procedure

Example: Use separate finger tracks to create a clenching fist:

1. Create a biped with five fingers that have three joints apiece.
2.  Turn on Set Parents Mode.
3. Turn on Separate FK Tracks for Arms.
4. Select the right hand and set a free key at frame 0.
5. Go to frame 30, move the hand up and forward, and then set a planted key.
6. Go to frame 40, set a free key (this will have IK on by default), select the pivot at the base of the hand (only hand pivots are present, not finger pivots), and reset the key.
7. The hand should now move up and then rotate at its base.
8. Select a finger. So far, it has no keys.
9. Animate each finger individually, so they curl up one at a time.

Interface



Enable Subanims—Enables Biped subanims. For more information about Biped subanims, see *Using Controllers (page 2–896)*

Manipulate Subanims—Modifies Biped subanims. For more information about Biped subanims, see *Using Controllers (page 2–896)*

Clear Selected Tracks—Removes all keys and constraints from the selected objects and tracks.

Clear All Animation—Removes all keys and constraints from the biped.

Mirror—Mirrors the entire biped animation.

Set Multiple Keys— Select keys using filters or apply a rotational increment to selected keys. Use this to change periodic motion keys in Track View. Displays the *Set Multiple Keys dialog (page 2–965)*

Set Parents Mode—When a limb key is created, keys are created for the parent objects also, provided Set Parents Mode is turned on. Use Set Parents Mode when you turn on Separate FK Tracks.

Set Parents Mode stores the position of the entire limb when a biped limb is moved using inverse kinematics instead of rotated using forward kinematics. For example, if Set Parents Mode is off and Separate Tracks are turned on for the biped arms, then the arm will snap back to its original position if you transform the biped hand.

If Separate Tracks are turned on for a biped body part, then turn on Set Parents Mode. This lets you use the Move transform to position the biped limbs.

Note: Separate Tracks adds biped object transform tracks.

Anchor Right Arm, Left Arm, Right Leg, Left Leg—Let you temporarily fix the location and orientation of hands and feet. Use anchors when you are setting up animation with inverse kinematics object space, in which the arm or leg follows an object in the scene. Anchors ensure that the arm or leg keeps its alignment until you set the second key that establishes the object-space sequence.

Tip: An alternative to anchors is to use Set Planted Key on the *Key Info rollout (page 2–954)*. When you use Set Planted Key, the limb is positioned to the previous IK key (Join To Previous IK Key).

Show All in Track View—Shows all the curves for the options in the Keyframing rollout in the track view.

Separate FK Tracks group

By default, **character studio** stores a finger, hand, forearm, and upper-arm key in the Clavicle track. The toe, foot, and calf keys are stored in the thigh track. This optimized approach to key storage works well in most cases. If you need extra tracks, turn them on for a specific biped body part.

For example, turn on Arms if you plan to create extensive finger-hand animation; if an arm key is deleted, it will not affect the finger-hand keys.

You must turn on Set Parents Mode for these toggles to take effect.

Arms—Turn on to create separate transform tracks for the finger, hand, forearm, and upper arm.

By default, there is one finger track per hand. All finger keys are stored in the *Finger0* transform track, the first link of the biped thumb.

Neck—Turn on to create separate transform tracks for the neck links.

Legs—Turn on to create separate toe, foot, and calf transform tracks.

Tail—Turn on to create separate transform tracks for each tail link.

Fingers—Turn on to create separate transform tracks for fingers.

Spine—Turn on to create separate spine transform tracks.

Toes—Turn on to create separate transform tracks for toes.

Ponytail 1—Turn on to create separate ponytail 1 transform tracks.

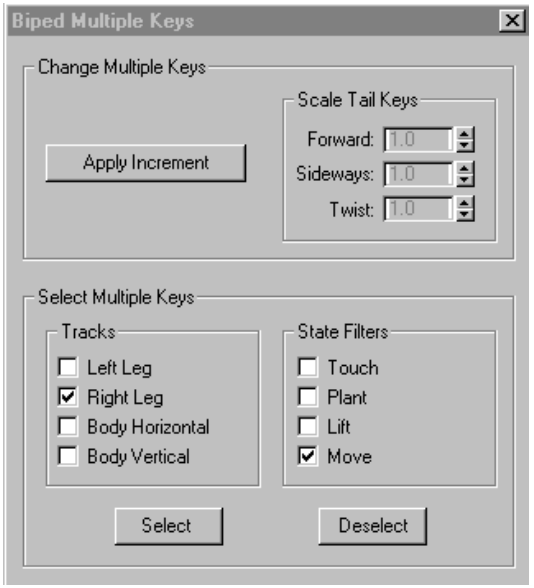
Ponytail 2—Turn on to create separate ponytail 2 transform tracks.

Set Multiple Keys Dialog

Select the Biped > Motion panel > Keyframing rollout > Set Multiple Keys > Set Multiple Keys dialog

Keys can be selected manually in Track View or the Track Bar and an increment applied to the selected keys. State Filters in this dialog select certain biped keys for you, based on foot states (Touch, Plant, Lift, and Move). Select Left Leg and the Move state filter, then click Select to select all the left leg keys in a Move state for example (Move is the leg state between footsteps).

Interface



Change Multiple Keys group

These controls allow you to apply *to a set of keys* the rotation or IK translation of a limb at the current frame. First select the keys in Track View, then rotate or move the limb, then click Apply Increment.

Apply Increment—Adjusts the rotation and/or position of a limb at the selected keys. Use this feature when you need to set the position of a limb over multiple keys.

Scale Tail Keys—Exaggerates or tones down the default motion applied to the biped's tail for three types of rotation. To activate all three spinners, select tail keys in the biped Tail track in Track View.

Forward—Sets the amount of forward and backward swing in the tail.

Sideways—Sets the amount of side-to-side swing of the tail.

Twist—Sets the amount of local X axis rotation of each tail object.

Select Multiple Keys group

These controls allow you to select keys according to the foot state at that frame. This is very helpful when you want to apply an increment to all keys of a particular type in a particular track.

- First, select the tracks you want: Left Leg, Right Leg, Body Horizontal, Body Vertical.
- Then select the foot states: Touch, Plant, Lift, Move
- Then click Select to select all the matching keys. Selected keys are highlighted in white in Track View.

Copy/Paste Rollout

Create a biped or select one. > Motion panel > Copy/Paste rollout (not available in Footstep, Motion Flow, or Mixer modes)

Controls on the Copy/Paste rollout let you copy and then paste posture, pose, or track information from one part of a biped to another, or from one biped to a different one. These are the three categories of information you can copy and paste:

- **Posture—**The positioning of selected biped objects.
- **Pose—**The positioning of the entire biped.
- **Track—**The animation track from the selected biped objects.

Note: Your copied information will be stored in the current copy collection.

Copy Collections

Copy collections are designed to make it easier to manage copied posture, pose, and track information by grouping it together. This improves the way copied animation data is organized when it is transferred between files within a session:

- You can display smaller sets of poses, postures, and tracks in the list.
- You can load more than one *CPY (page 3–925)* file into a single scene.
- You can either append a loaded set to an existing one or replace it entirely.

Note: You must create a copy collection before you can copy a biped's posture, pose or track.

Copy/Paste Buffers

In **character studio**, you can save multiple copy/paste buffers for each of the three modes: Posture, Pose, and Track. These are available from the Copied Postures/Poses/Tracks drop-down list; the active buffer that the Paste buttons use is the one whose name is visible in the field at the top of the list. A thumbnail view gives a preview of what the active buffer has saved.

Buffers are saved with your MAX scene file, and also remain available in your 3ds Max session, even if you reset.

The default name of a buffer depends on which mode you are in. In the Posture and Pose modes, the name of the buffer consists of abbreviated names of the body parts you selected, followed by a sequence number. For example, *RArmRFing1* is the first buffer for the posture or tracks of the biped's right arm and finger. In Pose mode, the name of the buffer is always "Pose," followed by a sequence number. For example, *Pose03*.

To give a buffer a custom name, make it active, highlight its name in the field at the top of the buffer list, and then enter a new name.

For hands-on experience using the Copy/Paste rollout, see the lesson called "Creating a Simple

Freeform Animation" found in the tutorial entitled *Animating with Freeform*.

Tip: When copying poses or postures containing COM data, the *Paste Options (page 2–973)* are activated.

Paste and Paste Opposite

For each mode, there are two paste options: Paste, and Paste Opposite. These are useful in different situations:

- **Paste**—In the Posture and Pose modes, Paste is useful for copying positioning from one biped to a different biped, or for restoring a biped's positioning at a different frame of an animation. In Track mode, Paste is useful mainly for copying movement from one biped to a different biped.
- **Paste Opposite**—In Posture mode, Paste Opposite is useful for making one limb assume the posture of the other, either at the same frame or different frames. In Pose mode, Paste Opposite reverses the biped's pose, or applies that reverse pose to a different biped. In Track mode, Paste Opposite can make a biped's limbs move symmetrically, or apply the opposite of the copied movement to a different biped.

Tip: After pasting a pose or tracks to a different biped, often you have to reposition it. This is easily done with Move All mode.

Note: If a copied pose, posture or track contains sub-animations, they can be pasted to another pose, posture or track as long as the same sub-animations already exist. Otherwise, they are ignored.

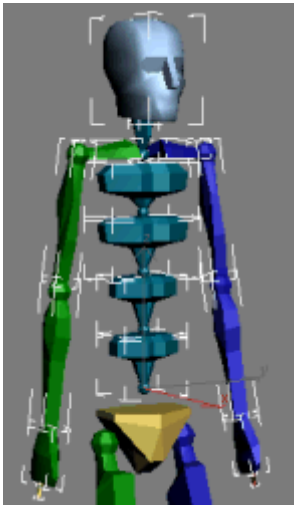
Procedures

Example: To load multiple collections between files within a single session:

1. Create a biped and load some animation onto it.
2. On the Copy/Paste rollout, click Create Collection and rename it **Upper**.

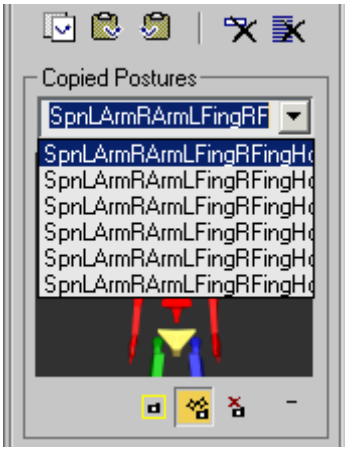


3. Select all of the biped's links from the spine up and, after making sure you're in *Posture mode (page 2–972)*, click Copy Posture. Repeat this step for every 10 frames of your animation.



4. Create a new collection and rename it **Lower**.
5. Go to frame 0. Select all of the biped's links from the pelvis down and copy a posture at every 10 frames.
6. Choose the Upper collection from the Copy Collections drop-down list.and save it. Then, select the Lower collection and save it as well.

7.  Save your scene and reset it (File > Reset).
8.  Create a new biped and load the Upper collection.
9.  For every 10 frames, select the biped's upper body and paste the corresponding posture from the Copied Postures drop-down list. Set a key for each pasted posture.



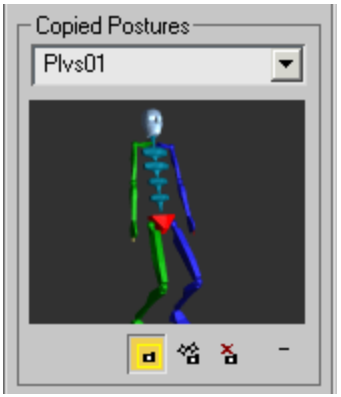
10.  Load the Lower collection and repeat last step with the biped's lower body selected, setting a key for each posture.
11.  Delete the Lower collection.
12.  Click Max Load Preferences and make sure both options are checked. Close the dialog and load the saved 3ds Max file.

    Notice that the Copy Collections drop-down list now contains three collections: two Upper collections (one from the current file and one from the incoming) and one Lower collection from the incoming file.

To capture different snapshots:

1. Under the thumbnail snapshot of the Copied Postures group, click Capture Snapshot from Viewport.
2. Rotate your current viewport to different user views and copy a biped posture.

   Notice that your snapshot matches the viewport angle.



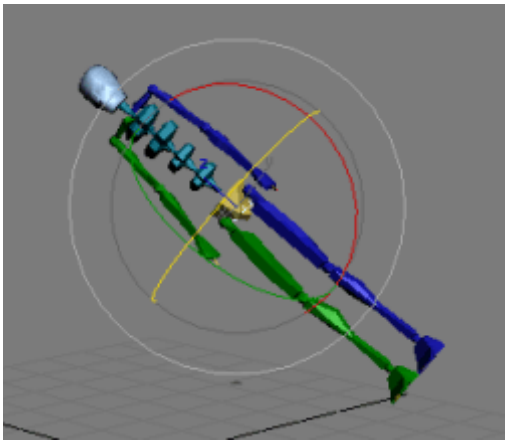
3. Click Capture Snapshot Automatically and copy a posture again.

   Notice that your snapshot is displayed from a frontal view.
4. Click No Snapshot and copy another posture.

   The capture has been replaced by a gray canvas.

Example: To maintain copied COM data when pasting poses:

1.  Create a biped and then create a new collection. Rename it **Poses**.
2.  Make sure you're in Pose mode and click Copy Pose at frame 0. Then, go to frame 30.
3. Using tools from the *Track Selection rollout (page 2–945)*, move your biped away from its current position and rotate it in all three axes.



The biped with offset position and orientation

4. In the *Paste Options group (page 2–973)*, enable all three Paste buttons.
5. Click Paste Pose. The biped resets back to its original position and orientation since its copied COM data is maintained. Undo the step.
6. Disable Paste Rotation but keep the other two buttons enabled.
7. Click Paste Pose. The biped resets back to its original position but remains rotated. This is due to the fact that the copied COM orientation is not maintained on paste. Undo the step.
8. Disabled both Paste Horizontal and Paste Vertical but enable Paste Rotation once again. Click Paste Pose.

   This time, the biped's position is maintained but not its orientation.



Paste Rotation is the only enabled button

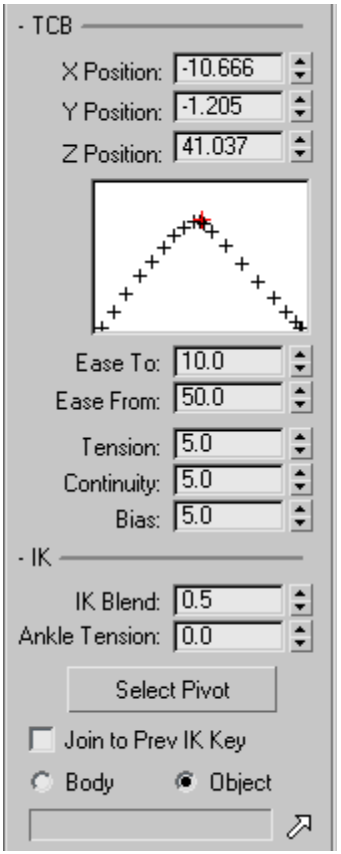
Example: To maintain COM offsets using By Velocity:

1. Create a biped. Then, on the *Track Selection rollout (page 2–945)*, turn on Body Horizontal.
2. Set a key at frame 0. Then, make sure Pose mode is selected and copy your biped's pose.
3. Go to frame 10 and move the Body COM 50 units in the Y axis. Set a new key.
4. In the *Paste Options group (page 2–973)*, enable Paste Horizontal but leave By Velocity unchecked.
5. Go to frame 30 and paste your biped's pose.

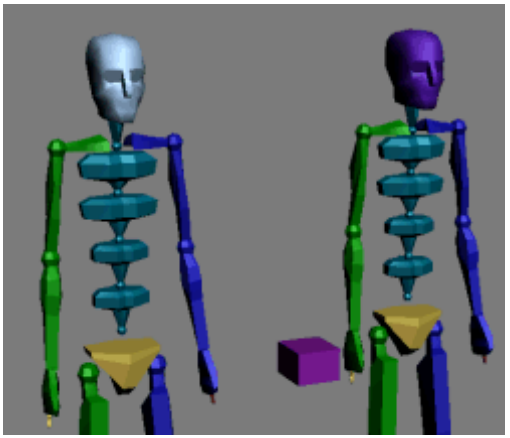
   The biped jumps to the center of your scene. Undo the step.
6. Go back in the Paste Options group and turn on By Velocity. Then, try pasting your pose again. Notice how the biped's COM position is now offsetted by another 50 units in Y.

Example: To set TCB/IK Values when pasting postures:

1. Create a biped and then create a new collection. Rename it Hands.
2. Select the biped's right hand and, on the *Key Info rollout (page 2–954)*, set a key.
3. In the TCB group, set Ease To to 10, Ease From to 50, and all three TCB values to 5.
4. In the IK group, set IK Blend to 0.5 and choose the Object Space option. Then, set another key.

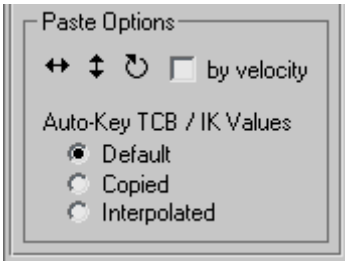


5. Make sure you're in Posture mode and click Copy Posture.
6. Create a new biped next to the original one. Create a box next to its right hand.



The original biped (left) and the new biped with box (right)

7. In the IK group, constrain the biped's hand to the box by clicking Select IK Object and then selecting the box in your viewport. Keep the Body Space option selected.
8. Auto Key On the *Keyframing Tools rollout (page 2–962)*, turn Separate Tracks off for the arms. Then turn Auto Key on.
9. In the *Paste Options group (page 2–973)*, set the Auto-Key TCB/IK Values to Default.



10. Click Paste Posture.

    The hand assumes the copied posture and the key's TCB and IK values are set to default. Keys are also set for the whole arm.
11. Go to frame 30 and turn Separate Tracks on.

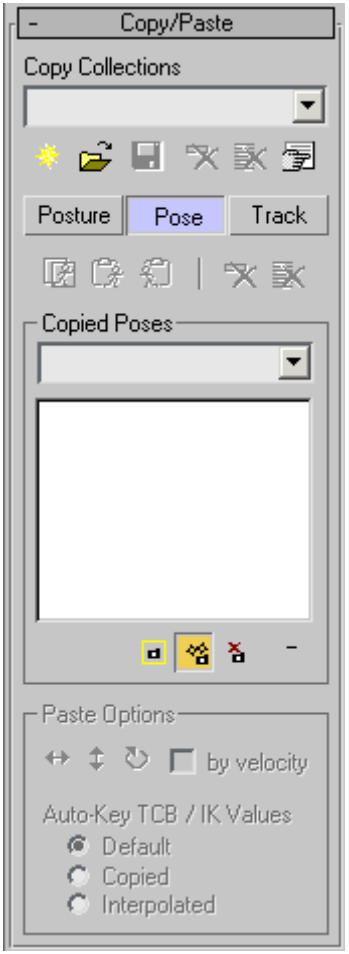
12.  Set Auto-Key TCB/IK Values to Copied and paste the posture.

    The copied value information is transferred to this new key. The key values are now 10 and 50 for the Ease To/Ease From, 5 for the TCB, and 0.5 for the IK Blend. The box is still the IK Object and Object Space remains selected. Also notice that the hand is the only keyed limb.

13.  Go to frame 15, set the Auto-Key TCB/IK Values to Interpolated and paste the posture again.

    This new key averages up the values based on the previous and next keys, setting Ease To/Ease From to 5 and 25, TCB to 15 and IK Blend to 0.25. Also, Body Space is selected from the previous key.

Interface



 **Create Collections—**Clears the current collection name and the poses, postures and tracks associated with it.

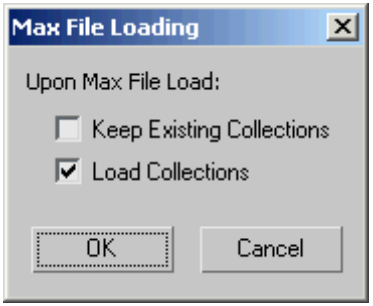
 **Load Collections—**Loads a CPY file and displays its collection name at the top of the Copy Collections drop-down list, making it active.

 **Save Collections—**Saves all the postures, poses and tracks stored within the active collection of the current session in a CPY file.

 **Delete Collection—**Removes the current collection from the scene.

 **Delete All Collection—**Removes all collections from the scene.

 **Max Load Preferences—**Displays a dialog with options for actions to take upon Max file open.



Keep Existing Collections—When on, sets biped to overwrite the existing Copy/Paste Buffer upon load. Default=Off.

Load Collections—When on, sets biped to load the Copy/Paste buffer from the incoming file, appending the current buffer. Default=On.



Posture, Pose, and Track—Choose one of these buttons to choose which kind of information you are copying and pasting. The copy/paste buttons change to indicate the current mode. Default=Pose.

Copy/Paste buttons—These buttons change according to the current mode, as described below in *Posture (page 2–972)*, *Pose (page 2–972)* and *Track (page 2–972)*.

 **Delete Selected—**Deletes the selected posture, pose, or track buffer. The selected buffer is the active one; that is, the buffer whose name is currently displayed in the Copied Postures/Poses/Tracks list.

 **Delete All—**Deletes all the buffers in the Copied Postures/Poses/Tracks list.

Posture mode



Copy Posture—Copies the posture of the selected biped objects and saves it in a new posture buffer.

Paste Posture—Pastes the posture of the active buffer onto the biped.

Paste Posture Opposite—Pastes the posture of the active buffer onto the opposite side of the biped.

Pose mode



Copy Pose—Copies the current pose of the entire biped and saves it in a new pose buffer.

Paste Pose—Pastes the pose of the active buffer onto the biped.

Paste Pose Opposite—Pastes the opposite of the pose of the active buffer onto the biped.

Track mode



Copy Track—Copies the tracks for the selected biped objects and creates a new track buffer.

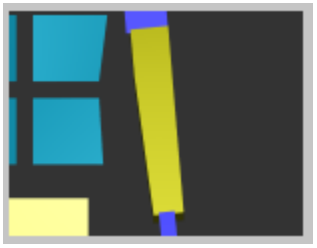
Paste Track—Pastes the track or tracks in the active buffer onto the biped.

Paste Track Opposite—Pastes the track or tracks in the active buffer onto the opposite side of the biped.

Copied Postures group

Copied Postures/Poses/Tracks list—For each of the modes, lists the buffers you have copied. The active buffer is the one that will be pasted by the paste buttons. To make a buffer active, choose it from the drop-down list. To change the name of a buffer, make it active, highlight its name, and enter a new name.

Thumbnail buffer view—For Posture and Track mode, displays a schematic view of the parts of the biped in the active copy buffer. For Pose mode, displays a schematic view of the entire biped. This can help you preview the effect of pasting the active buffer.



Thumbnail showing a copied left arm posture



Thumbnail showing a copied pose



Capture Snapshot from Viewport—When chosen, creates a snapshot of the active 2D or 3D viewport of entire biped.

Capture Snapshot Automatically—When chosen, creates a front view snapshot of the isolated body parts.

No Snapshot—When chosen, replaces the snapshot with a gray canvas.

Show/Hide Snapshot—Toggles maximize/minimize of the snapshot view.

Paste Options group

The Horizontal, Vertical and Rotation copy options found in previous versions of **character studio** have been replaced by the Paste Options. By default, copying poses or postures with the COM selected copies all three COM tracks.

 **Paste Horizontal/Vertical/Rotation buttons—**When turned on, the COM's Body Horizontal, Vertical or Rotation data are set to be pasted next time you perform a paste operation. Choosing to paste COM data maintains the world space positions and orientation of the copied COM onto the current data tracks. Otherwise, the COM's current positions and orientation are maintained when the copied data tracks are pasted. See *Example: To maintain copied COM data when pasting poses: (page 2–968)*

Tip: This choice can be seen as pasting absolute or relative COM animation. The former will overwrite the current data with the copied one while the latter will maintain the current data, ignoring all copied information.

By Velocity—When on, determines the values of the active COM tracks based on the previous trajectory of the COM through the scene. This option is only enabled when the Paste Horizontal/Vertical/Rotation buttons are active. See *Example: To maintain COM offsets using By Velocity: (page 2–969)*

Note: If there is no previous trajectory from which to gather the velocity data, the COM result is considered as absolute.

Auto-Key TCB/IK Values—Active only in *Auto Key mode (page 3–717)*.

Note: If part of the copied biped is linked to an IK object, that object it is not transferable using the copy/paste process. However, the Join to Prev IK Key option is maintained, as well as they Body or Object options. If a copied pose or posture contains an IK Object, the data becomes baked into the biped's link while keeping the key in Object space.

Note: If part of the pasted biped is linked to an IK object, that object is maintained even while the new pose or posture is honored. In other words, if an incoming pose requires a change in an existing link between a biped limb and an IK object, the link is maintained and adjusted accordingly.

- **Default**—Sets (on the *Key Info rollout (page 2–954)*) the TCB Ease To and Ease From to 0, and Tension, Continuity and Bias to 25. These settings are unrelated to what was copied or where it is pasted. Any IK key values already set are maintained. Otherwise, values will be obtained from the position in time between the previous and next keys.
- **Copied**—Sets the TCB/IK values to match those from the copied data. If a copied pose or posture is not on a key, the TCB/IK values are based on interpolation from the previous and next keys.
- **Interpolated**—Sets the TCB values to an interpolation from the animation on which you are pasting. If an existing key is being pasted, its TCB values are maintained. Along the same lines, existing IK values are also maintained when pasted. If no key is present, the IK values are set from the position in time between the previous and next keys.

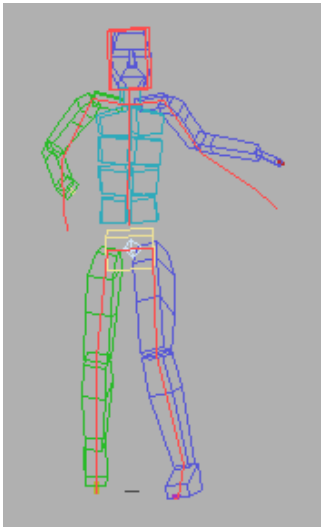
Layers Rollout

Select the Biped > Motion panel > Layers rollout

Controls on the Layers rollout allow you to add layers of animation above the original biped animation. This is a powerful way of making global changes to your character animation. For example, add a layer to a run cycle and rotate the spine forward at any frame, and it becomes a crouched run. The original biped motion is kept intact and can be viewed by switching back to the original layer. You can view layers individually or as a composite of all the animation in all the layers. Layers behave like freeform animation; the biped can adopt any position.

Tip: You can globally translate both footstep and freeform animation by doing a layered edit on the *center of mass (page 2–933)*. For example, by adding a layer and moving the center of mass you can move a freeform or footstep animation.

With layers, you can easily adjust raw motion-capture data, which contains keys at every frame. Simply add a layer and keyframe the biped. The original layer is displayed as red bones.

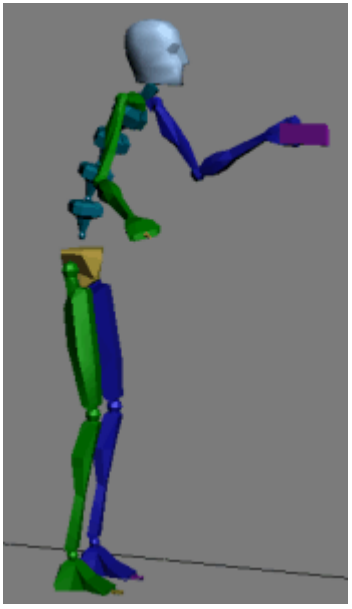


Honoring IK Constraints Across Layers

You can maintain a biped's IK constraints across layers by retargeting its hands and feet to the original layer. This assures that the biped's constrained body parts are honored and locked in place while you make animation changes on the upper layers.

You can also choose to retarget a biped using another one as reference. This is useful when you need to precisely match hands and feet positioning between two animated bipeds with different body proportions.

You can save an animation containing honored constraints across layers in a *BIP (page 3–916)* file. However, if the layered animation links to another biped as reference, that link is not saved with the animation.



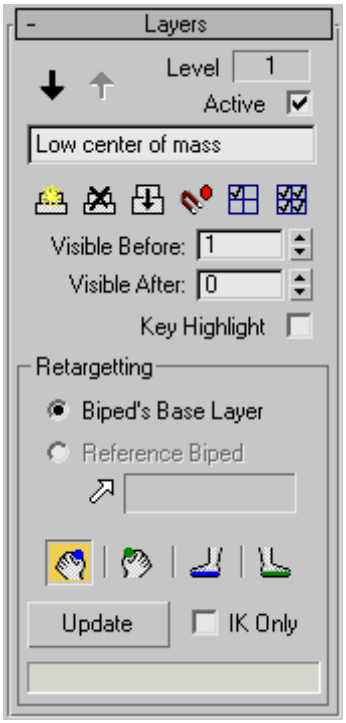
Procedures

Example: To Maintain IK Constraints Across Layers:

This procedure takes into account the fundamentals of animating across layers and expands on the concept of honoring IK constraints.

1. Prepare a biped with IK keys on its feet, and its left hand constrained to an IK object.
2. Select your biped and expand the Layers rollout on the Motion panel.
3.  Create a new layer and rename it **Low center of mass**.

   The Retarget Left Arm button in the Retargeting group becomes active, which indicates that the current layer honors the IK constraint of the base layer for this body part.



4. Auto Key Turn on *Auto Key mode (page 3–717)*.

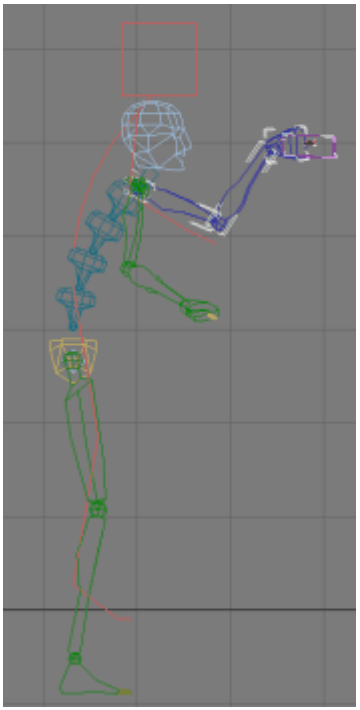
5. On the Track Selection rollout, turn on Body Vertical.

   This selects the biped's center of mass.

6. Lower the center of mass on the Z axis until the biped's head is under the base layer's head display (represented as a red box).

   Tip: To avoid deselecting the center of mass, you can lock it by clicking the *Selection Lock Toggle (page 3–707)* (or press the spacebar).

   The entire biped is lowered except its left hand, which remains locked to the object because the current layer retargets it to match the respective IK constraint of the base layer.

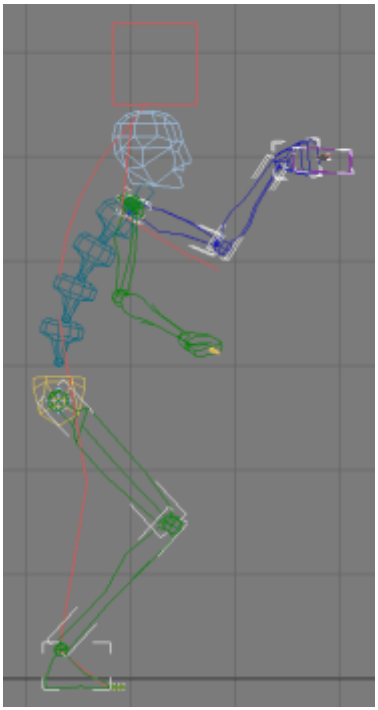


Only the biped's left hand constraint is honored.

7. In the Retargeting group, turn on both Retarget Left Leg and Retarget Right Leg. Then, click Update.

   The biped's IK feet are adjusted to match those of the base layer. The animation keys are updated to reflect the current layer's retargeted feet.

   Tip: If some body parts assume odd positions, simply drag the *time slider (page 3–701)* a few frames past your current frame, and then drag it back home.



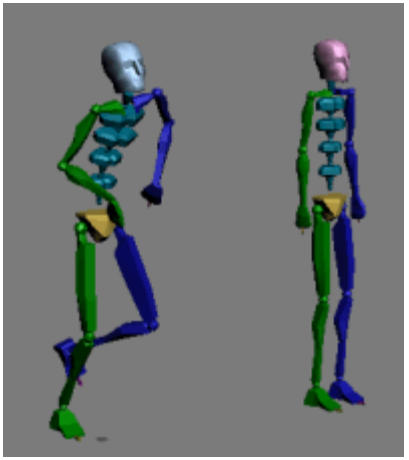
Both the biped's feet and left hand constraints are honored.

8. You can continue animating the biped to your liking. Collapse your layers when you are satisfied.

Example: To Maintain IK Constraints From a Reference Biped:

This procedure centers on using an animated biped as a retarget reference for another biped with disproportionate body parts. This method is often used when a motion is imported from raw data and adapted to bipeds with different proportions.

1. Prepare an animated biped or load a *BIP (page 3–916)* file onto one.
2. Create a second biped alongside and rename it **Disproportionate Biped**.

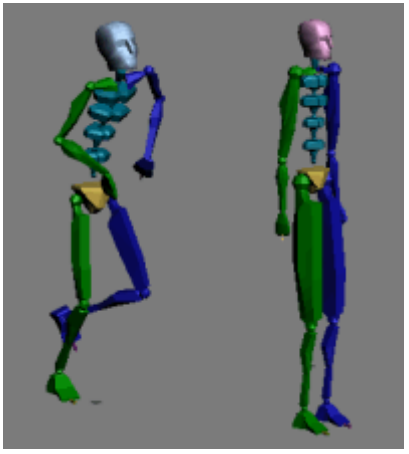


The animated biped on the left and Disproportionate Biped on the right

3. Select *Disproportionate Biped* and enter *Figure Mode (page 2–982)*.

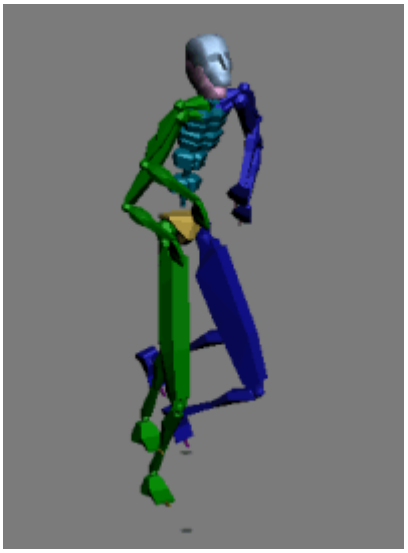
   You can now change the biped's structure.
4. Scale its upper arms and thighs to make them disproportionate from the rest of its body.

   Tip: You can quickly select body parts on both sides of a biped if you first select the body part, then click Symmetrical on the *Track Selection rollout (page 2–945)*.



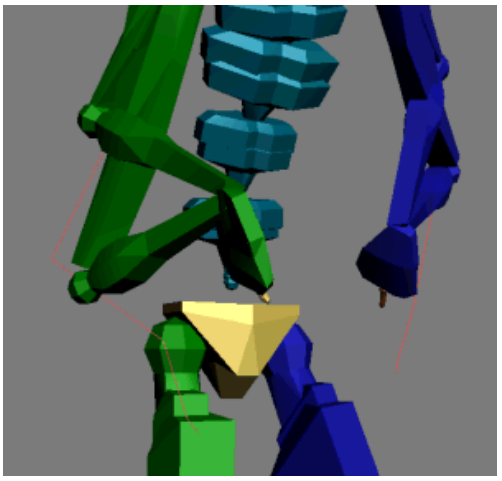
5.  Exit Figure Mode.
6. Keep *Disproportionate Biped* selected and expand the Layers rollout.
7.  Click Select Reference Biped and select your original biped in the viewport.

   Disproportionate Biped adopts the animation from the reference biped, whose name is now displayed next to the Select Reference Biped button.



8.  Add a new layer and rename it *Retargeted Biped*.
9.  Turn on both Retarget Left Arm and Retarget Right Arm, and then click Update.

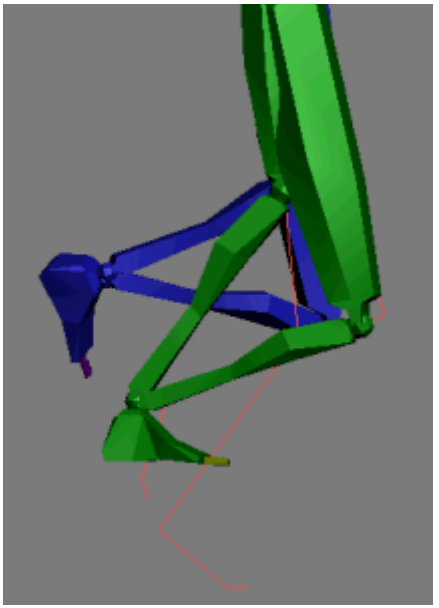
   Both hands precisely match those from the reference biped. The animation keys are updated to reflect the current layer's retargeted hands.



Both hands are retargeted to honor the base layer IK constraints.

10.  Turn on both Retarget Left Leg and Retarget Right Leg. Then, click Update.

    The feet are correctly retargeted to the reference biped, updating the respective animation keys.

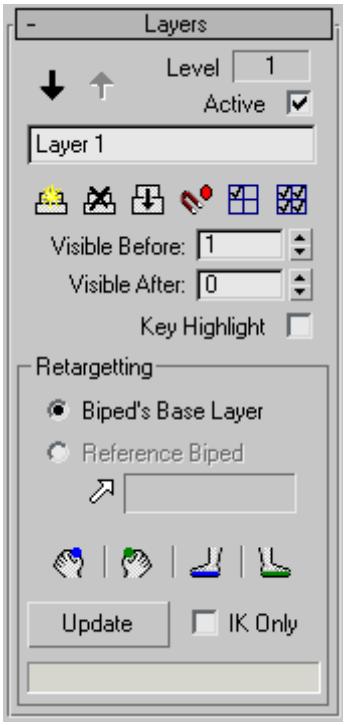


The feet are retargeted to honor the base layer IK constraints.

11.  Use Previous layer and Next Layer to switch between the original and *Retargeted Biped* layers. The base layer displays the non-targeted motion while *Retargeted Biped* shows an adjusted motion that matches both hands and feet from *Disproportionate Biped* with the original biped.

12.  You can continue animating the biped to your liking. Collapse your layers when you are satisfied.

Interface



 **Previous-Next Layer**—Navigate through the layers using the up and down arrows.

Level—This field displays the current layer (Level).

Active—Toggles the displayed layer on and off.

Name Field—Type a name to easily identify a layer.

 **Create Layer**—Creates a layer, and the Level field increments.

Position the biped to create keys in a layer.

 **Delete Layer**—Deletes the current layer.

All layer numbers 'above' the one deleted are decremented by one.

 **Collapse Layers**—Collapses all the layers into layer 0.

Legs that stray from the original footsteps in higher layers are "pulled in" to the original footsteps.

 **Snap Set Key**—Snaps the selected biped part to its original position in layer 0 and creates a key.

Use this in higher layers to return the selected biped part to the original motion. If a layer has a posture key that bends the character forward at frame 2 and you want to return the biped to its original posture motion at frame 50, use Snap Set Key at frame 50 with a spine object selected. The character will interpolate from its forward posture position to its original posture position between frame 2 and 50.

 **Activate Only Me**—View the animation in the selected layer.

Choose Play after turning on Activate Only Me to view a layer's keys in motion.

 **Activate All**—Activates all the layers.

Playing the animation shows a composite of all the layers.

Visible Before—Sets the number of preceding layers to display as stick figures.

Visible After—Sets the number of succeeding layers to display as stick figures.

Key Highlight—Displays keys by highlighting the stick figures.

Retargeting Group

The tools in this group let you animate a biped across layers while maintaining the IK constraints of the base layer. You can also choose to use a different biped in your scene to use as reference for your biped's retargeted hands and feet.

Biped's Base Layer—Choose this method to use the IK constraints from the selected biped's original layer as retarget reference. Default=active.

Reference Biped—Choose this method to use the biped whose name is displayed next to the Select Reference Biped button as retarget reference.

 **Select Reference Biped—**Lets you choose a biped to use as retarget reference for the selected biped. The chosen biped's name is displayed next to the button.

 **Retarget Left Arm—**When on, the IK constraints of the base layer are honored for the biped's left arm.

 **Retarget Right Arm—**When on, the IK constraints of the base layer are honored for the biped's right arm.

 **Retarget Left Leg—**When on, the IK constraints of the base layer are honored for the biped's left leg.

 **Retarget Right Leg—**When on, the IK constraints of the base layer are honored for the biped's right leg.

Update—Calculates the selected biped's hands and feet positions for every set key based on the retarget method (Base Layer or Reference Biped), the active retarget body parts buttons, and the IK Only option.

IK Only—When on, the biped's constrained hands and feet are retargeted only during the frames on which they are IK controlled. When off, the hands and feet are retargeted during both IK and FK keys. Default=off.

Dynamics & Adaptation Rollout

Select the biped. > Motion panel > Dynamics & Adaptation rollout

The controls on the Dynamics & Adaptation rollout let you specify the way you want to create biped animation. The parameters let you modify gravity strength, dynamic properties for keys generated by newly created footsteps, the number of transform tracks available on the biped, and prevent key adaptation.

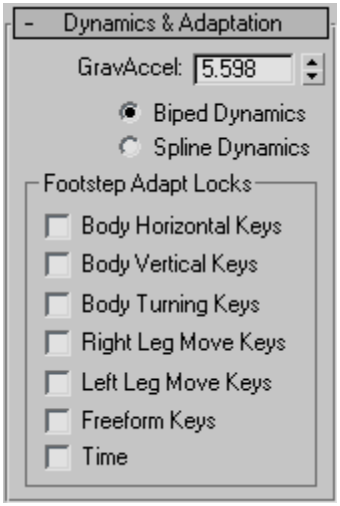
Biped Dynamics and Spline Dynamics

These parameters specify how new biped center of mass keys are created and therefore how you want to work with the biped. Turning on Spline Dynamics will create center of mass keys, without gravity and balance calculation, for newly created footsteps. This may feel more familiar to new users who are already familiar with spline interpolation.

Biped Dynamics calculates biped airborne trajectory, knee bend on landing and positions the biped to maintain balance when the spine is rotated. When parameters change, the biped adapts. Turn on Biped Dynamics and use this adaptation to your advantage.

You can always change from one method to the other on a per-key basis or for the entire animation at any time.

Interface



GravAccel—Sets the strength of the gravitational acceleration used to calculate the biped's motion.

By default, this parameter is set to accurately simulate Newtonian gravity as found on the Earth's surface.

At a value of 0, the biped still runs but the feet hardly get off the ground.

Biped Dynamics—Creates new center-of-mass keys using Biped Dynamics.

Keys for the center of mass (COM) Balance Factor and Dynamics Blend parameters are set to a value of 1. Biped calculates airborne trajectories and biped balance.

Note: Use Dynamics Blend on the Key Info rollout > Body group to set a vertical COM key that is a blend between Dynamics and Spline Interpolation; a value of 1 is full Dynamics, a value of 0 is full Spline interpolation.

Spline Dynamics—Creates new center-of-mass keys using full spline interpolation. All new vertical COM keys are created with Dynamics Blend=0, and all new horizontal COM keys are created with Balance Factor=0.

Footstep Adapt Locks group

Lock specified tracks to prevent automatic adjustments being made to those tracks when footsteps are moved in space or edited in time. All the locks except for Time work for footstep editing in space. Time locks upper body keys when footsteps are edited in time (Track View). Adapt Locks only applies to a Footstep animation, not a freeform animation.

When you move a footstep in space or adjust footstep timing, Biped automatically adapts existing keyframes to match the new footsteps. Adapt locks allows you to preserve the exact position of already created keys for a selected track.

Adapt Locks does not need to be on all the time. For example, if you want to raise all the footsteps along the world Z-axis, without changing the upper body position, turn on Footstep Adapt Locks group > Body Vertical Keys, turn on Footstep mode, select all the footsteps and move them up along the world Z-axis. The footsteps are repositioned, the legs are adapted, but the upper body retains the same motion rather than being raised with the footsteps. Now turn off Body Vertical Keys, and the upper body still retains its original motion.

Use Footstep Adapt Locks > Time to retain upper body motion while editing footstep duration in Track View. When the duration of a footstep is changed, the biped leg will adapt by re-timing the touch, plant, and lift keys. The biped upper body keys will retain their exact motion.

Body Horizontal Keys—Turn on to prevent adaptation of body horizontal keys when footsteps are edited in space.

Body Vertical Keys—Turn on to prevent adaptation of body vertical keys when footsteps are edited in space.

Body Turning Keys—Turn on to prevent adaptation of body turning keys when footsteps are edited in space.

Right Leg Move Keys—Turn on to prevent adaptation of right leg move keys (a leg move key, is a leg key between footsteps) when footsteps are edited in space.

Left Leg Move Keys—Turn on to prevent adaptation of left leg move keys (a leg move key, is a leg key between footsteps) when footsteps are edited in space.

Freeform Keys—Turn on to prevent adaptation of a freeform period in a footstep animation. The biped's position during a freeform period will not move if footsteps after the freeform period are moved further away.

Time—Turn on to prevent adaptation of upper body keys when footstep duration is changed in Track View.

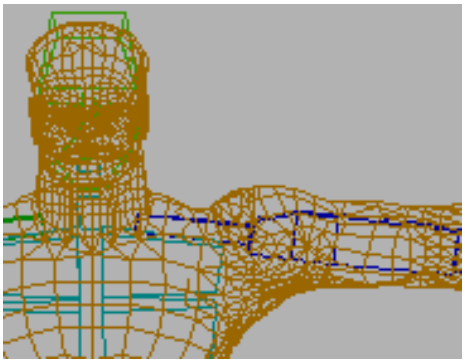
Note: Leg and foot keys for frames in which the foot is in contact with the ground are always automatically adapted.

See also

Adjusting Vertical Motion (page 2–878)

Figure Mode

Select the Biped > Motion Panel > Biped rollout > Figure mode



While Figure mode is active, you can change biped structure and fit that structure to a character mesh. It can be used for a variety of other procedures as well.

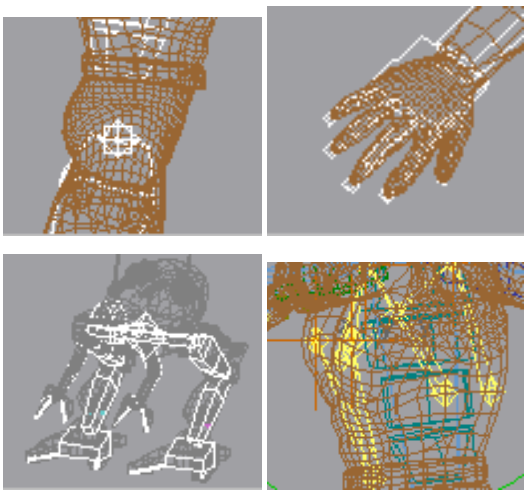
- **Figure mode is a reference position to fit a biped to a mesh.** Use Figure mode to fit a biped to the mesh representing your character. This "reference" or Figure mode position, in which the biped is aligned to the mesh, is necessary when a mesh is linked or attached to the biped with Physique. After the biped is positioned to fit within a mesh, leave Figure mode on during the process of attaching a mesh to the biped with Physique, or when using Select And Link on the 3ds Max toolbar to link the mesh objects of a character to the biped.

  The relationship or "fit" position between the biped and the mesh can always be restored by turning on Figure mode, regardless of which motion file happens to be loaded. After fitting a biped to a mesh in Figure mode, use Save on the Biped rollout to save a figure file (*.fig*). If you accidentally reposition the biped in Figure mode, load the figure file.

- **Figure mode is used for biped adjustment after a mesh is attached to correct biped joint location.** After using Physique to attach a mesh character to the biped, you may want to reposition a biped limb relative to the mesh. For example, if the biped shoulder joint is too far out relative to the mesh shoulder, then the Physique modifier must be inactivated, the biped limbs adjusted, and then a Reinitialize in Physique must be performed before reactivating the Physique modifier.
- **Figure mode is used for biped adjustment after a mesh is attached to correct posture in a motion file.** Figure mode is also used to make adjustments after a character is attached or linked to the biped. After loading a *.bip* motion file, for example, you may find that the character is hunched too far forward during the entire animation. Rotating the biped's spine objects in Figure mode will correct the character's posture for the entire animation. This is a basic procedure where you simply rotate the biped limbs in Figure mode and then exit Figure mode; the posture will be corrected for the entire animation.
- **Figure mode is used to define biped structure.** The *Structure rollout (page 2–984)* is displayed when you work in Figure mode, allowing you to tailor the biped to your mesh character. After creating a biped, make all of your biped structure changes on the Structure rollout. For example, you may want to use one toe with one toe link if your character is wearing shoes or if your character's toes do not need to be keyframed individually. Set the biped structure before "fitting" the biped to the your mesh character.
- **Turn Figure mode on to scale a character.** Use the height spinner on the Structure rollout to scale a complete character (a complete character has a biped and mesh attached with Physique).
- **Reverse-Knee Characters.** If your character mesh has reverse knees, rotate the biped calves or thighs along the local X axis 180 degrees in Figure mode; the biped local X axis is along the *length* of the limb. **character studio** assumes you want a reverse knees character if the calves or thighs are rotated past 90 degrees in the local X axis. When Figure mode is turned off, the biped walks, runs, and jumps with reverse knees.

Notes on Fitting the Biped to a Mesh in Figure Mode

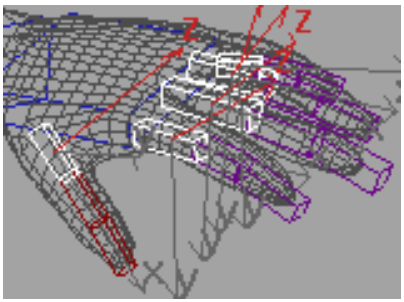
These are quick notes designed to give you a general sense of the issues involved when a biped is fitted to a mesh.



- Use the Structure rollout to set the number of toes and fingers; specify the number of links per finger and toe. One toe with one toe link is often sufficient if your character wears shoes, or if animating individual toes will not be necessary.
- Put the lowest biped spine object at the character's belt-line.
- Scale the biped fingers to slightly protrude from the character's hand.

- Rubber Band mode and scale are used to size the biped limbs to fit the biped to a mesh.
- Use the options from the *Twist Links group (page 2–987)* to transfer twisting animation in the biped's associated mesh.
- Use Props to represent weapons or tools attached to your character.
- Use Select And Link on the 3ds Max toolbar to link non-deformable (mechanical) objects to the biped. Do this after Physique is applied to prevent Physique from generating extra links (Envelopes). Superflous envelopes (links) can be turned off in Physique however, so this is not critical.

  Note: Objects like eyeballs and weapons should be linked to the biped after Physique is applied; otherwise links (Envelopes) will extend to these objects when Physique is applied.
- Reposition and use Ponytails on the Structure rollout to animate a character's jaw, ears, hat, hair and ponytails.
- A saved *.fig* file can be reloaded if the biped is repositioned in Figure mode by mistake.



Move the first link on each finger to position the fingers relative to the mesh; use local and world coordinate systems for this. Scale the finger links to position the joints. After positioning the thumb, rotate the first thumb link around the local X-axis until the local Z-axis creates a natural rotation for the thumb (refer to the image). A User view and toggling back and forth between a shaded and wireframe display is helpful when fingers are positioned.

Structure Rollout

Select the biped. > Motion panel > Biped rollout > Turn on Figure mode > Structure rollout

Turn Figure mode on to enable parameters on the Structure rollout. The Structure rollout provides parameters for changing the biped's skeletal structure to match your character mesh (dinosaur, robot, human and so on). You can also add up to three props to your biped, to represent tools or weapons.

After setting parameters to suit your character, use the Height parameter to scale the biped to the size of the mesh representing your character. This is the first step in fitting a biped to a mesh in Figure mode.

Procedures

To Scale a Biped and a Physique Mesh:

-  Select the biped, turn on Figure mode, and then change the biped height on the Structure rollout.

  The biped and mesh scale together.

To Skin a Biped with Twist Links:

This procedure takes into account the process of skinning a character and expands on the concept of using twist links to better deform a skinned mesh.

1. Position a biped within the mesh.
2.  Select the biped and turn on Figure Mode.
3. On the Structure rollout, turn on the Twists option.

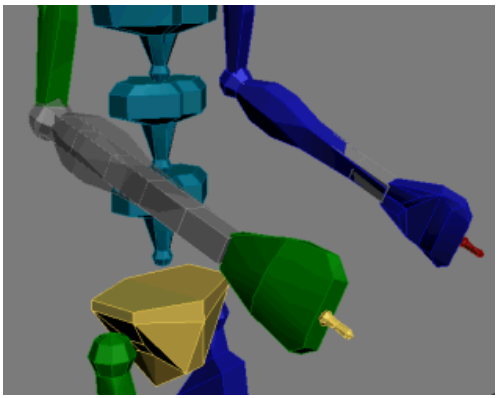
   This enables the input fields for all biped limbs.

Note: Horse Link is only available if your biped has four leg links.

4. Set Forearm to 5.

   Both forearms have five twist links.

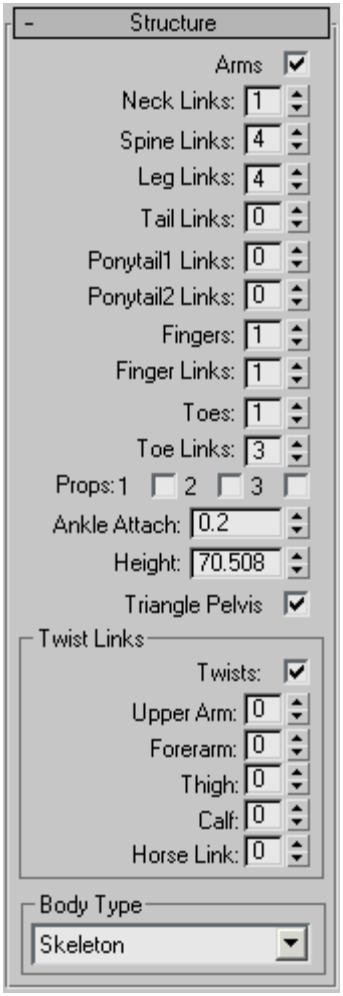
   Tip: To get a better view of your twist links, you can select the forearms and turn on See-Through in the *Display Properties rollout (page 1–55)* on the Display panel (or press Alt+x).



See-Through is turned on for the right forearm only.

5. Add a *Skin modifier (page 1–791)* to the mesh.
6. Unfreeze all the biped twist bones.
7. On the Skin Parameters rollout of the Motion panel, add all the biped bones to the skin except the forearms.
8. Select and freeze the twist bones again.
9. Turn on Edit Envelopes on the Parameters rollout.
10. Select and adjust the Envelopes of the twist bones until the desired behavior is achieved, moving and rotating the hand to test.

Interface



Arms—Sets whether or not arms will be generated for the current biped.

Neck Links—Sets the number of links in the biped neck. Default=1. Range=1 to 25.

Spine Links—Sets the number of links in the biped spine. Default=4. Range=1 to 10.

Leg Links—Sets the number of links in the biped legs. Default=3. Range=3 to 4.

Tail Links—Sets the number of links in the biped tail. A value of 0 specifies no tail. Default=0. Range=0 to 25.

Ponytail1/2 Links—Sets the number of Ponytail Links. Default=0. Range=0 to 25.

You can animate hair with ponytail links. Ponytails are linked to a character's head and can be used to animate other appendages. Reposition ponytails in Figure mode and use them to animate a character's jaw, ears, nose, or anything that should move with the head.

Unlike the process of selecting the biped hand and dragging to reposition the entire arm, ponytails must be keyed using rotational transformations.

Fingers—Sets the number of biped fingers. Default=1. Range=0 to 5.

Finger Links—Sets the number of links per finger. Default=1. Range=1 to 3.

Toes—Sets the number of biped toes. Default=1. Range=1 to 5.

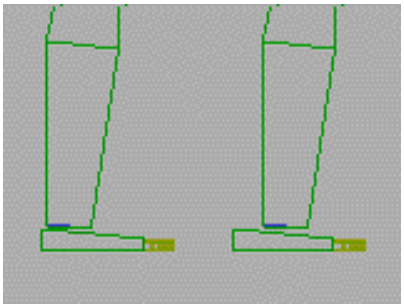
Toe Links—Sets the number of links per toe. Default=3. Range=1 to 3.

Tip: Characters wearing shoes might need only one toe with one link.

Props 1/2/3—Turns on up to three props, which can be used to represent tools or weapons attached to the biped. By default, prop 1 appears next to the right hand, prop 2 appears next to the left hand, and prop 3 appears centered between in front of the torso.

You can animate props throughout the scene, like any object. You can use 3ds Max controllers to animate props and can then collapse animation onto the prop's transform controller for use in Mixer, Motion Flow, and Layer editing.

Ankle Attach—Specifies the ankles' point of attachment along the corresponding foot block. The ankles can be placed anywhere along the centerline of the foot block, from the heel to the toe.



Ankle Attach=0.25 and Ankle Attach=0.5

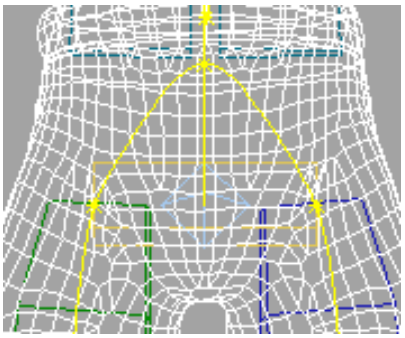
A value of 0 places the ankle attachment point at the heel. A value of 1 places the ankle attachment point at the toes. Click the spinner up arrow to move the ankle attach point toward the toes. Range=0 to 1.

Height—Sets the height of the current biped.

Use to size the biped to your mesh character before Physique is attached. This parameter is also used in a procedure to scale your character after Physique is attached.

Triangle Pelvis—Turn on to create links from the upper legs to the lowest biped spine object when Physique is attached. Normally the legs are linked to the biped pelvis object.

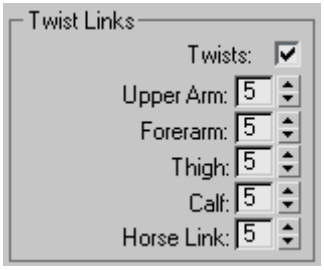
The pelvis area can be a problem when the mesh is deformed with Physique. Triangle Pelvis creates a more natural spline for mesh deformation.

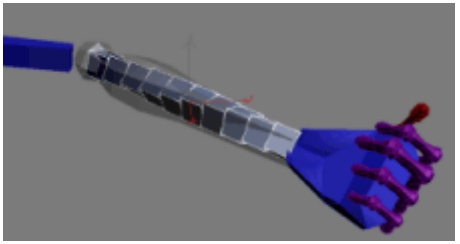


Triangle Pelvis creates two links that extend from the legs to the lowest biped spine object. A link from the biped pelvis to the lowest spine object is also created. This provides natural deformation to this area after Physique is applied and the character is moving. If you are working on a new character, turn this on before applying Physique. If Bones is turned on in the Display rollout, links from the legs to the lower spine object are visible.

Twist Links group

The bone twist option, previously limited to the biped's forearm, has been expanded to include all limbs. These settings allow better mesh deformation on skinned models (using *Physique (page 2–834)* or *Skin (page 1–791)*) when twisting occurs on animated limbs. Use the *Twist Poses rollout (page 2–950)* for better control over your twist links.





Note: If a limb has twists, the twist links will control the skin deformation while the base link (the biped's forearm, for example) drives the animation. Twist links can not be animated.

Twists—Enables twist links for biped limbs. When enabled, twist links become visible but remain frozen. You can unfreeze them using Unfreeze By Name or Unfreeze By Hit on the *Freeze rollout (page 1–54)*.

Upper Arm—Sets the number of twist links in the upper arms. Default=0. Range=0 to 10.

Forearm—Sets the number of twist links in the forearms. Default=0. Range=0 to 10.

Thigh—Sets the number of twist links in the thighs. Default=0. Range=0 to 10.

Calf—Sets the number of twist links in the calves. Default=0. Range=0 to 10.

Horse Link—Sets the number of twist links in the horse link. Default=0. Range=0 to 10.

Note: You must set Leg Links to 4 to enable Horse Link.

Body Type group

The Body Type group lets you select the biped's body type:

- **Skeleton—** The skeleton body type provides a realistic skeleton which fits naturally into mesh skins.
- **Male—** The male body type provides a silhouette mold based on basic male proportions.
- **Female—** The female body type provides a silhouette mold based on basic female proportions.
- **Classic—** The classic body type is the same as the biped object from older versions of **character studio**.

Footstep Mode Rollouts

Footstep Mode

Select a Biped. > Motion panel > Biped rollout > Footstep Mode

When Footstep mode is active, you can create or edit footsteps to generate a walk, run, and jump footstep pattern. You also edit selected footsteps in space and append footsteps using parameters available in Footstep mode.

If footsteps are extracted during motion capture import, turn on Footstep mode to edit footsteps in the viewports.

A powerful feature in Biped is the ability to adapt keyframes when footsteps are edited in space or time. The following tracks are influenced in the vicinity when a footstep is edited in space:

- Body Horizontal keys change to step or hop within range of new footstep locations.
- Body Vertical keys change to match possible changes in stride length, since the body must be lower in order to step longer distances.
- Body Rotation keys change to bank into turns created by changes in path curvature or body speed.
- Right or left leg keys in a move state must be adapted to step between new locations.

  Note: If for some reason you do not want the adaptation to occur, use the Footstep Adapt Locks settings on the *Dynamics & Adaptation rollout (page 2–980)* to keep the biped from correcting the body position.

Two additional rollouts display when Footstep mode is active: *Footstep Creation rollout (page 2–988)* and *Footstep Operations rollout (page 2–990)*.

Footstep Creation Rollout

Select a Biped with footsteps > Motion panel > Biped rollout > Footstep Mode > Footstep Creation rollout

The Footstep Creation rollout, available on the Motion panel when Footstep mode is on, provides controls for creating and editing footsteps. Create a walk, run, or jump footstep pattern using these controls.

Tip: All footsteps created here are inactive; you activate them using Create Keys for Multiple Footsteps on the *Footstep Operations rollout (page 2–990)*.

The timing parameters at the bottom of the rollout work with Create Footsteps (append) and Create Footsteps (at current frame) to change timing for newly created footsteps. These change depending on whether you select Walk, Run, or Jump mode.

Procedures

To create multiple footsteps:

1.  On the Biped rollout, click Footstep Mode.

   You are now in Footstep mode, and can create, activate, or edit footsteps.
2. On the Footstep Creation rollout, choose the gait you want to create — Walk, Run, or Jump.
3. Click Create Multiple Footsteps.

   The Create Multiple Footsteps dialog displays for the selected gait.
4. Set multiple footstep parameters, and then click OK.
5.  On the Footstep Operations rollout, click Create Keys for Inactive Footsteps to activate the footsteps.

To create footsteps manually, beginning at the current frame:

1. On the Footstep Creation rollout, click Walk, Run, or Jump and set timing parameters for that gait.
2.  Click Create Footsteps (at current frame).
3. Click in a viewport to create a footstep. Continue clicking to create more footsteps.
4.  On the Footstep Operations rollout, click Create Keys for Inactive Footsteps.
5. Play the animation.

To append footsteps onto the existing footsteps:

1.  On the Footstep Creation rollout, click Create Footsteps (append).
2. Click in a viewport to create a footstep. Continue clicking to create more footsteps

   By default, footsteps alternate from foot to foot. The first click creates a right footstep, the next click creates a left footstep, and so on. Look at the prompt line and the cursor to see which type of footstep will be created next.

Interface



 **Create Footsteps (append)—**Turn on Create Footstep mode. Manually place footsteps by clicking in any viewport. Hold the mouse button down after clicking to move a footstep. Release the mouse button to place the footstep.

Each new footstep is appended to the end of the biped's footstep sequence. Create Footsteps alternates right and left footsteps as you create new ones. Press Q to toggle between a left and right footstep.

Newly created footsteps are bright green for right footsteps and bright blue for left footsteps. Once the footsteps have been activated, the footsteps change color to pastel green and pastel blue.

 **Create Footsteps (insert at current frame)—**Create footsteps at the current frame. Footstep creation alternates between left and right footsteps.

 **Create Multiple Footsteps—**Create a walk, run, or jump footstep pattern automatically. Select a gait type before using Create Multiple Footsteps. Displays the Create Multiple Footsteps dialog. The dialog differs slightly depending on the gait (walk, run, jump) you have selected.

Create Multiple Footsteps Dialog: Walk (page 2–992)

Create Multiple Footsteps Dialog: Run (page 2–995)

Create Multiple Footsteps Dialog: Jump (page 2–997)

 **Walk:** Sets the biped gait to Walk. Any footsteps you add will have walk characteristics until you change to another mode (run or jump). Each new footstep will start before the end of the previous footstep on the opposite side.

 **Run—**Sets the biped gait to Run. Any footsteps you add will have run characteristics until you change to another mode (walk or jump). Each new footstep will start after the end of the previous footstep on the opposite side.

 **Jump—**Sets the biped gait to Jump. Any footsteps you add will have jump characteristics until you change to another mode (walk or

run). Each new footstep will start at the same time as the most recent footstep on the opposite side. Alternately, it may start after the end of the previous footstep.

Timing parameters



Use these parameters with Create Footsteps (append) or Create Footsteps (at current frame) to apply timing to newly created footsteps. These parameters are different for each gait and change as you select a different gait.

Walk Footstep —(Walk only): Specifies the number of frames a new footstep will be on the ground during a walk.

Double Support—(Walk only): Specifies the number of frames when both feet will be on the ground at the same time during a walk.

Run Footstep—(Run only): Specifies the number of frames a new footstep will be on the ground during a run.

Airborne—(Run and jump only): Specifies the number of frames when the body will be in the air during a run or a jump.

2 Feet Down—(Jump only): Specifies the number of frames when two sequential footsteps, on opposite sides, will be on the ground during a jump.

Footstep Operations Rollout

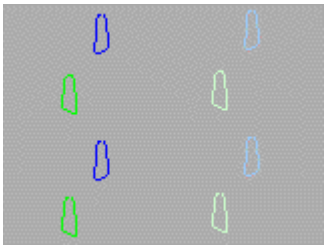
Select the Biped > Motion Panel > Biped rollout > Footstep Mode > Footstep Operations rollout

Once footsteps are created on the Footstep Creation rollout, use parameters on the Footstep Operations rollout to activate and deactivate footsteps, and to adjust the footstep path.

Interface




Create Keys for Inactive Footsteps—Activates all inactive footsteps. Activation creates default keys for any footsteps that do not have them. If a footstep does not have keys, it is displayed as bright green (right foot) or bright blue (left foot). After keys are created for the footsteps, the footsteps change color to pastel green and pastel blue.



Inactive footsteps on the left and activated footsteps on the right

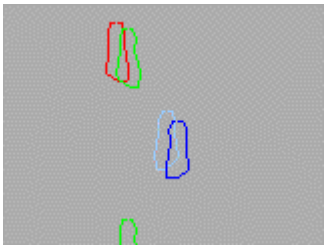

Deactivate Footsteps—Removes the body keys assigned to the selected footsteps, making those footsteps inactive. The footsteps themselves remain in the scene.


Delete Footsteps—Deletes the selected footsteps.


Copy Footsteps—Copies the selected footsteps and biped keys to the footstep buffer. Biped will only copy a continuous sequence of footsteps (2,3,4,5...). You can't copy discontinuous footsteps (3,4,7,8...).

If any footsteps exist that have not been activated, the Copy button is grayed. Activate the footsteps first, then try again.

Tip: Turn on Buffer mode on the Biped rollout to view and edit only the buffered footsteps and biped motion.



A pasted footstep is inactive until it overlaps a like foot and turns red. Mouse up to activate the new footsteps.

 **Paste Footsteps—**Pastes footsteps from the footstep buffer into a scene. The footsteps appear inactive in the scene. Move them so they overlap the active footsteps. When a footstep turns red, mouse up and the pasted footsteps will activate.

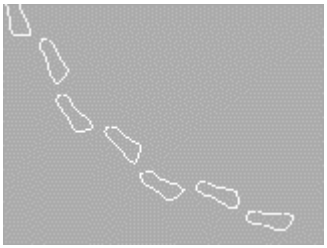
Warning: If any footstep in the buffer overlaps in time with a footstep previous to the one onto which you are pasting, a message appears and the paste is not performed. This can occur if the first footstep you are pasting and the original footstep you are pasting onto both have *double support* periods during the same duration of the footstep. The second pasted footstep and the footstep prior to the one you are pasting onto may overlap in time.

To correct the problem, you may want to move the right edge of the preceding footstep from the original scene to the left, or move the left edge of the second pasted footstep to the right while in buffer mode. Which option you choose depends on what footstep timing and support relations you want to end up with.

Tip: To create a cycle of a motion with alternating footsteps, you must copy and paste at least three footsteps. All body keys between the start of the first step and end of the last step are also copied and pasted.

Bend—Bends the path for the selected footsteps. The path is bent to the left or right as you move the spinner. Other footsteps after the selected footsteps will be moved to maintain their positions relative to the repositioned footsteps.

The Bend spinner is grayed if inactive footsteps are present.



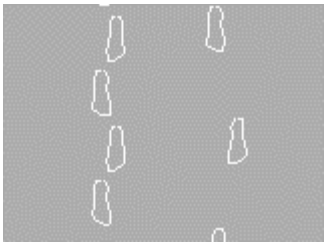
Bend=0.5

Scale—Changes the width or length for the selected footsteps. The selected footsteps are scaled around the first footstep in the selection.

First select the width or length box (or both), then use the spinner to set the amount of scaling.

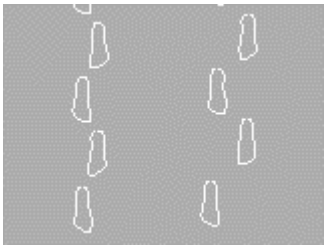
If any footsteps exist that have not been activated, the Scale spinner is grayed out. Activate the footsteps first, then try again.

Length—When Length is selected, the Scale spinner changes the stride length of the selected footsteps. Length and Width may both be active at the same time.



Length=1 and Length=2

Width—When Width is selected, Scale changes the stride width of the selected footsteps. Length and Width may both be active at the same time.



Width=1 and Width=2

Create Multiple Footsteps Dialog: Walk

Select the Biped. > Motion Panel > Biped rollout > Footstep Mode > Footstep Creation rollout > Walk > Create Multiple Footsteps > Create Multiple Footsteps Walk dialog

The Create Multiple Footsteps dialog for the walk gait creates a sequence of walking footsteps using a series of parameters.

 This dialog displays when you select the Walk gait on the Motion panel's Footstep Creation rollout, then click Create Multiple Footsteps.

Note: The Create Multiple Footsteps dialog appears differently depending on the gait chosen. For other gaits, see *Create Multiple Footsteps Dialog: Run (page 2–995)* and *Create Multiple Footsteps Dialog: Jump (page 2–997)*.

Procedures

To make the biped walk up or down stairs:

- In the First Step group, set Actual Stride Height to a value other than 0. A value higher than 0 will cause the biped to step up, while a value below 0 will make it step down.

To make the biped walk in place:

- In the First Step group, Set Parametric Stride Length to zero.

To make the biped walk backward:

- In the First Step group, set Parametric Stride Length to a value less than zero.

  The absolute value of the Parametric Stride Length is still the length of the stride.

To make the biped speed up as it walks:

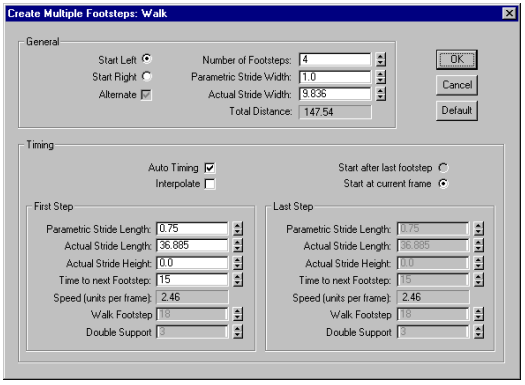
1. In the Timing group, click Interpolate.

   The controls in the Last Step group are enabled.

2. In the Last Step group, set Time To Next Footstep to be a value less than Time To Next Footstep in the First Step group.

   You can adjust the values in either group, or both. The important thing is to make the Last Step a shorter time than the First Step.

Interface



Start Left—Starts the footstep sequence with a left step.

Start Right—Starts the footstep sequence with a right step.

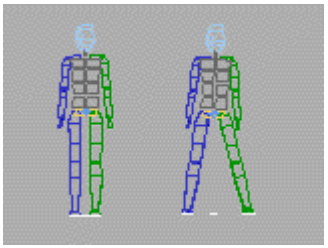
Alternate—Footsteps will alternate between right and left. Alternate is always selected when the

Walk gait is selected. You can only clear Alternate when Run or Jump gaits are selected.

Number of Footsteps—Determines the number of new footsteps to be created.

Parametric Stride Width—Sets the stride width as a percentage of the pelvis width. A value of 1.0 produces a stride width equal to the pelvis width. A value of 3.0 produces a wide, waddling stride. Changes to this setting automatically change the Actual Stride Width.

Parametric describes the parameter in terms of biped anatomy, and Actual describes the value in 3ds Max units.



Stride Width=1 and Stride Width=3

Actual Stride Width—Sets the stride width in modeling units. Changes to this setting automatically change the Parametric Stride Width.

Total Distance—Displays the total distance the footsteps will travel with the current settings. This value cannot be changed directly.

OK—Creates footsteps with the current settings.

Cancel—Cancels footstep creation.

Default—Resets the values on the dialog to default values.

Timing group

Auto Timing—Sets timing parameters automatically. Auto Timing affects the following timing parameters for the Walk gait:

- Walk footstep, Double Support

  When Auto Timing is selected, these parameters are automatically adjusted to reasonable values. Control the footstep sequence by adjusting the Stride Length and Time To Next Footstep parameters.

  When Auto Timing is off, you can control the footstep sequence by adjusting the gait timing parameters, but you can't change the Time To Next Footstep parameter.

Interpolate—Control acceleration or deceleration of the series of footsteps. When this box is selected, a second set of step parameters under Last Step is enabled.

Biped creates footsteps starting with the values of the parameters under First Step and ending with the values of the parameters under Last Step.

By interpolating between the two, Biped produces a footstep series that changes over time.

When Interpolate is cleared, the Last Step parameters are grayed out. Biped creates all the footsteps using only the parameters under First Step.

Start After Last Footstep—Appends the newly created footsteps to the end of the existing footstep sequence.

Start at Current Frame—Inserts the newly created footsteps at the current frame after the existing footstep sequence, allowing you to make a gap in time before the footsteps start again.

First Step and Last Step groups

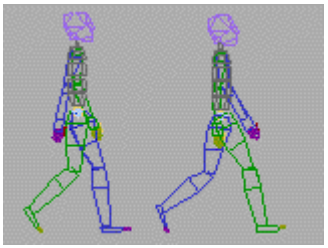
Parametric Stride Length—Sets the stride length for the new footsteps as a percentage of the length of the biped's leg. The default value of 0.75 gives an average stride of normal proportions.

A value of 1.0 will produce a stride length equal to the leg length, which makes the biped stretch slightly to reach the next step. A value of 0.0 will

make the biped walk in place. A negative stride length will make the biped walk backwards.

When a biped walks backwards, it does not simply reverse the forward movement but maintains the correct foot-state sequence with the toe touching the ground first, followed by the heel.

Adjusting Parametric Stride Length automatically changes the value for Actual Stride Length.



Stride Length=0.75 and Stride Length=1

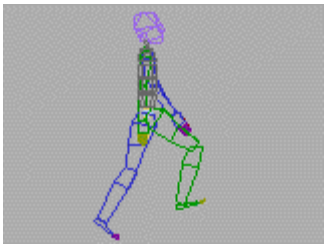
Actual Stride Length—Sets the stride length for new footsteps in 3ds Max units.

The same rules apply as for Parametric Stride Length (described above).

Adjusting Actual Stride Length automatically changes the value for Parametric Stride Length.

Actual Stride Height—Sets the rise or fall between footsteps. You can use this parameter to create a set of footsteps going up or down a slope or a stairway.

The value for Actual Stride Height is the difference in height in units between each of the new footsteps. Positive values step up and negative values step down.



Stride Height=5 units

Time to Next Footstep—Specifies the number of frames in each foot movement cycle. A cycle starts with the frame that a foot touches the ground, continues as the foot lifts and moves, and ends with the frame before the foot touches the ground again. This parameter is only enabled if Auto Timing is on.

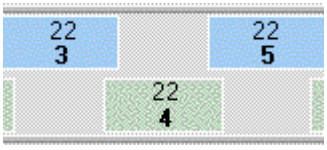
Speed—Displays the number of units the biped will move per frame. It changes in response to changes in the other parameters but cannot be adjusted directly.

The following two parameters are only enabled when Auto Timing is off.

You can use these parameters instead of Auto Timing to control the speed of the forward motion over the series of footsteps. However, because these parameters both affect the footsteps' time in contact with the ground, using them to slow down a walk gives the walk a hesitant, 'stop-go' quality.

Walk Footstep—Specifies the number of frames each footstep will be on the ground during a walk.

The higher the number, the longer each biped foot remains in contact with the ground and, consequently, the slower the speed of the walking motion.



Footsteps 3 through 5 are on the ground for 22 frames each

Double Support—Specifies the number of frames both feet will be on the ground at the same time during a walk.

The higher the number, the longer the period during which both feet remain in contact with the ground during each walk cycle and, consequently, the slower the speed of the walking motion.



The dotted line surrounds the double-support period (6 frames)

Create Multiple Footsteps Dialog: Run

Select the Biped > Motion Panel > Biped rollout > Footstep Mode > Footstep Creation rollout > Run > Create Multiple Footsteps > Create Multiple Footsteps Run dialog

The Create Multiple Footsteps dialog for the run gait allows you to create a sequence of running footsteps by setting a series of parameters. This dialog will display when Run is selected on the Footstep Creation rollout first, and then Create Multiple Footsteps is clicked.

Interface

Start Left—Starts the footstep sequence with a left step.

Start Right—Starts the footstep sequence with a right step.

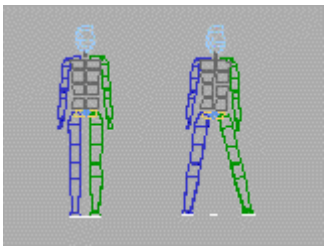
Alternate—Turn on to alternate between right and left footsteps. When this box is cleared, all of the footsteps will be either right or left, producing the effect of a hopping run on one foot.

Number of Footsteps—Determines the number of new footsteps to be created.

Parametric Stride Width—Sets the stride width as a percentage of the pelvis width.

A value of 1.0 produces a stride width equal to the pelvis width. A value of 3.0 produces a wide, waddling stride. Changes to this setting automatically change the Actual Stride Width.

Parametric describes the parameter in terms of biped anatomy, and Actual describes the value in 3ds Max units.



Stride Width=1 and Stride Width=3

Actual Stride Width—Sets the stride width in modeling units. Changes to this setting automatically change the Parametric Stride Width.

Timing

Auto Timing—Sets timing parameters automatically.

Auto Timing affects the following timing parameters for the Run gait:

- Run footstep, Airborne

  When Auto Timing is selected, these parameters are automatically adjusted to reasonable values. You control the footstep sequence by adjusting the Stride Length and Time to Next Footstep parameters.

  When Auto Timing is off, you can control the footstep sequence by adjusting the gait timing parameters, but you can't change the Time to Next Footstep parameter.

Interpolate—Control acceleration or deceleration of the series of footsteps. When this box is selected, a second set of step parameters under Last Step is enabled.

Biped creates the footsteps starting with the values of the parameters under First Step and ending with the values of the parameters under Last Step.

By interpolating between the two, Biped produces a footstep series that changes over time.

When Interpolate is cleared, the Last Step parameters are grayed out. Biped creates all the footsteps using only the parameters under First Step.

Start After Last Footstep—Appends the newly created footsteps to the end of the existing footstep sequence.

Start at Current Frame—Inserts the newly created footsteps at the current frame after the existing footstep sequence allowing you to make a gap in time before the footsteps start again.

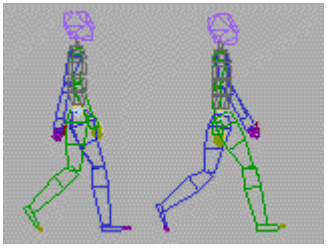
First Step and Last Step

Parametric Stride Length—Sets the stride length for the new footsteps as a percentage of the length of the biped's leg. The default value of 0.75 gives an average stride of normal proportions.

A value of 1.0 will produce a stride length equal to the leg length, which makes the biped stretch slightly to reach the next step. A value of 0.0 will make the biped run in place. A negative stride length will make the biped run backwards.

When a biped runs backwards, it does not simply reverse the forward movement but maintains the correct foot-state sequence with the toe touching the ground first, followed by the heel.

Adjusting Parametric Stride Length automatically changes the value for Actual Stride Length.

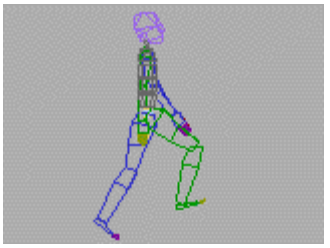


Stride Length=0.75 and Stride Length=1

Actual Stride Length—Sets the stride length for the new footsteps in 3ds Max units.

The same rules apply as for Parametric Stride Length (described above).

Adjusting Actual Stride Length automatically changes the value for Parametric Stride Length.



Stride Height=5 units

Actual Stride Height—Sets the rise or fall between footsteps. You can use this parameter to create a set of footsteps going up or down a slope or a stairway.

The value for Actual Stride Height is the difference in height in units between each of the new footsteps. Positive values step up and negative values step down.

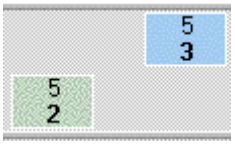
Time to Next Footstep—Specifies the number of frames in each foot movement cycle. A cycle starts with the frame in which a biped foot touches the ground, continues as the foot lifts and moves, and ends with the frame before the foot touches the ground again. This parameter is only enabled if Auto Timing is on.

Speed—Displays the number of units the biped will move per frame. It changes in response to changes in the other parameters but cannot be adjusted directly.

The following two parameters are only enabled when Auto Timing is off. You can use these parameters instead of Auto Timing to control the speed of the forward motion over the series of footsteps.

Run Footstep—Specifies the number of frames each footstep will be on the ground during the run.

The higher the number, the longer the biped foot remains in contact with the ground and, consequently, the slower the speed of the running motion.



Footsteps 2 and 3 are on the ground for 5 frames each

Airborne—Specifies the number of frames the body will be in the air between footsteps.

The higher the number, the longer the biped hangs in the air for each step and, consequently, the slower the speed of the running motion.

Create Multiple Footsteps Dialog: Jump

Select the biped. > Motion Panel > Biped rollout > Footstep Mode > Footstep Creation rollout > Jump > Create Multiple Footsteps > Create Multiple Footsteps Jump dialog

The Create Multiple Footsteps dialog for the jump gait allows you to create a sequence of jumps by setting a series of parameters. It displays when you've selected Jump on the Footstep Creation rollout on the Motion panel, and then click Create Multiple Footsteps.

Interface

Start Left—Starts the footstep sequence with a left step.

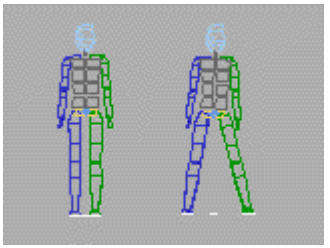
Start Right—Starts the footstep sequence with a right step.

Alternate—When this box is selected, the footsteps will alternate between right and left. When this box is cleared, the footsteps will be either right or left steps, causing the biped to hop on one foot.

Number of Footsteps—Determines the number of new footsteps to be created.

Parametric Stride Width—Sets the stride width as a percentage of the pelvis width. A value of 1.0 produces a stride width equal to the pelvis width. A value of 3.0 produces a wide, waddling stride. Changes to this setting automatically change the Actual Stride Width.

Parametric describes the parameter in terms of biped anatomy, and Actual describes the value in 3ds Max units.



Stride Width=1 and Stride Width=3

Actual Stride Width—Sets the stride width in modeling units. Changes to this setting automatically change the Parametric Stride Width.

Timing

Auto Timing—Sets timing parameters automatically.

Auto Timing affects the following timing parameters for the Jump gait:

- 2 Feet Down, Airborne

  When Auto Timing is selected, these parameters are automatically adjusted to reasonable values. You control the footstep sequence by adjusting the Stride Length and Time To Next Footstep parameters.

  When Auto Timing is off, you can control the footstep sequence by adjusting the gait timing

parameters, but you can't change the Time To Next Footstep parameter.

Interpolate—Allows you to control acceleration or deceleration of the series of footsteps. When this box is selected, a second set of step parameters under Last Step is enabled.

Biped creates the footsteps starting with the values of the parameters under First Step and ending with the values of the parameters under Last Step.

By interpolating between the two, Biped produces a footstep series that changes over time.

When Interpolate is cleared, the Last Step parameters are grayed out. Biped creates all the footsteps using only the parameters under First Step.

Start After Last Footstep—Appends the newly created footsteps to the end of the existing footstep sequence.

Start at Current Frame—Inserts the newly created footsteps at the current frame after the existing footstep sequence, allowing you to make a gap in time before the footsteps start again.

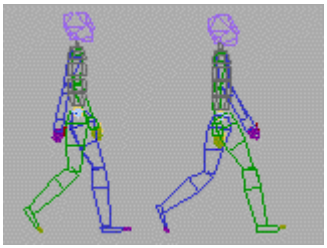
First Step and Last Step

Parametric Stride Length—Sets the stride length for the new footsteps as a percentage of the length of the biped's leg. The default value of 0.75 gives an average stride of normal proportions.

A value of 1.0 will produce a jump length equal to the leg length, which makes the biped stretch slightly to reach the next step. A value of 0.0 will make the biped jump in place. A negative stride length will make the biped jump backwards.

When a biped jumps backwards, it does not simply reverse the forward movement but maintains the correct foot-state sequence with the toe touching the ground first, followed by the heel.

Adjusting Parametric Stride Length automatically changes the value for Actual Stride Length.



Stride Length=0.75 and Stride Length=1.0

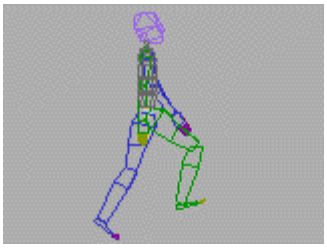
Actual Stride Length—Sets the stride length for the new footsteps in 3ds Max units.

The same rules apply as for Parametric Stride Length (described above).

Adjusting Actual Stride Length automatically changes the value for Parametric Stride Length.

Actual Stride Height—Sets the rise or fall between footsteps. You can use this parameter to create a set of footsteps going up or down a slope or a stairway.

The value for Actual Stride Height is the difference in height in units between each of the new footsteps. Positive values step up and negative values step down.



Stride Height=5 units

Time to Next Footstep—Specifies the number of frames in each foot movement cycle. A cycle starts with the frame a particular foot touches the ground, continues as the foot lifts, moves, and ends with the frame before the foot touches the ground again. This parameter is only enabled if Auto Timing is on.

Speed—Displays the number of units the biped will move per frame. This changes in response to changes in the other parameters but cannot be adjusted directly.

The following two parameters are only enabled when Auto Timing is off. You can use these parameters instead of Auto Timing to control the speed of the forward motion over the series of footsteps.

2 Feet Down—Specifies the number of frames that the left and right footsteps will be on the ground during the jump. The higher the number, the longer the pause between each jump and, consequently, the slower the speed of the jumping motion.



The dotted line surrounds the 2 feet down period (15 frames)

Airborne—Specifies the number of frames the body will be in the air during the jump.

The higher the number, the longer the biped hangs in the air for each jump and, consequently, the slower the speed of the jumping motion.



The dotted line surrounds the airborne period (13 frames)

Convert to Freeform or Footsteps Dialogs

Select the Biped. > Motion panel > Biped rollout > Convert

When you click Convert on the Biped rollout of the Motion panel, a Convert To dialog displays:

Convert to Freeform or *Convert to Footsteps*, depending on the animation method of the currently loaded motion. Use Convert to Freeform for unrestricted key editing or Convert to Footsteps to take advatage of footsteps.

Note: By default, Jumping or airborne periods between footsteps are calculated by **character studio** based on gravity strength (GravAccel) and time between footsteps. Biped elevation during these jumping and airborne periods will be lost when converting to freeform from footsteps (unless you create a vertical COM key in the airborne period in the footstep animation), a freeform animation uses spline interpolation for the center of mass position and does not account for gravity. Elevation in airborne periods is restored if you convert back to footsteps.

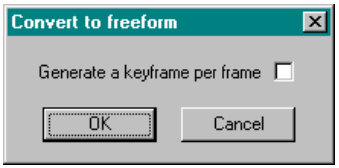
When to Use Convert

- Use Convert after using Save Segment in Motion Flow mode to save a script as a *.bip* file. Exit Motion Flow mode, load the *.bip* file and click Convert to extract footsteps. Save Segment applies IK Blend values of 1 to the biped foot keys for the keys at footsteps. If a freeform animation is part of the script in Motion Flow mode, and you want to convert the entire script motion to footsteps, then use the Save Keys at every Frame option in Save Segment and use Load Motion Capture File on the Motion Capture rollout. This will allow you to extract footsteps using the proximity method.
- Use Convert if you are working on a footstep animation and want to switch to a freeform animation.
- Use Convert if you have started a freeform animation and want to covert to a footstep animation. In order to convert a freeform animation to a footstep animation, the file must be properly set up by locking the feet to world space using IK Blend before converting to

footsteps. When creating freeform animations, you should set your leg keys to IK Blend=**1.0** during the periods you want the feet to be planted, and set the move keys to IK Blend=**0.0**. This will insure that the feet are locked and rid of unnecessary foot sliding as the body is moved. When converting, if the leg keys have been set up this way, Biped will extract footsteps during any duration where IK Blend=1.0.

Interface

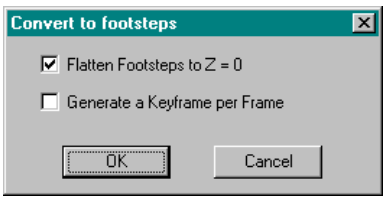
Convert to Freeform dialog



Generate a keyframe per frame—Creates keys at every frame.

Note: Converting footsteps to freeform creates foot IK Blend values of 1 for the biped feet for the original footstep keys. This simplifies keyframing by putting the feet into world coordinate space, which prevents them from sliding when the biped is moved. These foot IK Blend values are also used when you click Convert (on the Biped rollout) to convert back to footsteps.

Convert to Footsteps dialog



Generate a key per frame—Creates a key at every frame, and extracts footsteps based on foot IK Blend values equal to 1. Save Segment in Motion Flow mode stores the active script as a *.bip* file without footsteps. The biped foot keys are assigned IK Blend values of 1 for the original footstep keys. After loading a *.bip* file saved using Save Segment in Motion Flow mode, use Convert (on the Biped rollout) to extract footsteps.

Convert is also useful if you have already converted from footsteps to freeform in which case the feet are assigned an IK Blend value of 1 for the keys that represented footsteps.

Flatten Footsteps to Z=0—The entire biped is repositioned to place the footsteps at Z=0.

Note: An IK Blend value of 1 for the feet puts them into world coordinate space and prevents them from slipping while setting biped keys in a freeform animation.

Footstep Mode Dialog

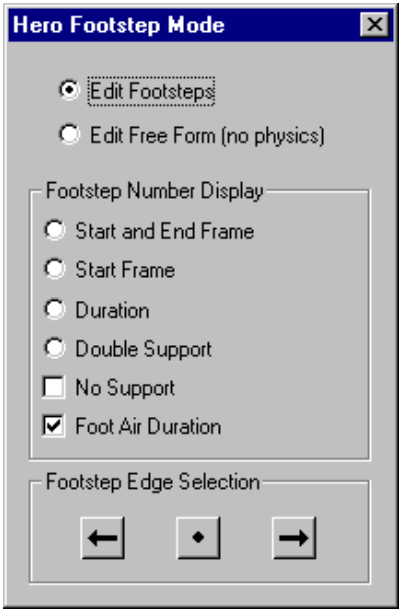
Graph Editors menu > Track View – Dope Sheet > Right-click on a footstep track.

Right-click anywhere in the Footstep track in Track View (Dope Sheet) to display the Footstep Mode dialog. This is a modeless dialog and can remain visible while you continue to work in Track View.

This dialog allows you to:

- Control the duration information displayed with the footsteps.
- Create a freeform period between footsteps.
- Select the right or left edges of footstep blocks or the whole block.

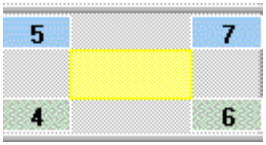
Interface



Edit Footsteps—In this mode (which is the default), you can edit the biped's footsteps to change their duration; start and end frames; airborne duration; and so on.

Edit Free Form (no physics)—In this mode, you can edit the biped's body keys for the frames at which the biped is airborne.

Free Form suspends the physically based dynamics that normally control biped motion. This is essential when you want to make the biped fly, or sit down, or fall over. Note that the freeform period must be between footsteps.



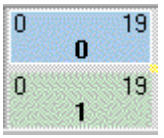
The freeform area appears as a yellow block between the footsteps.

Footstep Number Display group

The footstep blocks can have any one of four time settings displayed for them. You can only show one of these at a time.

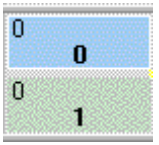
The footstep step number is always displayed on the footstep blocks (in boldface).

Start and End Frame—Displays the start and end frames of the footstep (from Touch to Lift). Biped displays just the start frame if the footstep block is too small to show both. Zoom in to see both start and end numbers.



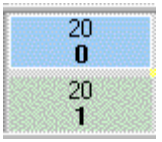
Both footsteps 0 and 1 begin at frame 0 and end at frame 19.

Start Frame—Displays just the start frame number (Touch).



Both footsteps 0 and 1 begin at frame 0.

Duration—Displays the number of frames that the foot is in contact with the ground (from Touch to Lift).

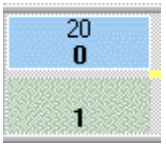


Both footsteps 0 and 1 are in contact with the ground for 20 frames.

Double Support—Displays the number of overlapping frames in which both feet are in contact with the ground.

You can also turn on the following two numbers for the intervals between the footsteps. You can

display both numbers at the same time by selecting both boxes.



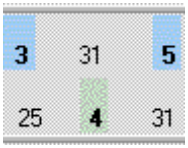
Footsteps 0 and 1 share a double-support period of 20 frames.

No Support—Displays the number of frames that the whole biped is airborne; that is, the frames in which neither of the feet have any contact with the ground.



The no-support period between frames 3 and 4 and 4 and 5 is 12 frames.

Foot Air Duration—Displays the number of frames that each foot has no contact with the ground.



The foot air duration for the left footstep track between frames 3 and 5 is 31 frames.

Footstep Edge Selection group



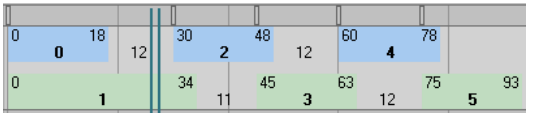
Select Start of Footstep—Selects the leftmost key for the current footstep selection.

Select Entire Footstep—Selects the entire footstep.

Select End of Footstep—Selects rightmost key for the current footstep selection.

Track View (Biped)

Graph Editors menu > Track View - Dope Sheet or Track View - Curve Editor



In this Track View - Dope Sheet view of a walking motion, footsteps are represented by colored squares, footstep numbers appear in boldface inside the squares, in between the start and end frames of the footstep. The numbers between the colored boxes represent foot-air duration.

Biped and Track View

The Dope Sheet can be used to create and adjust biped keyframes, edit biped footsteps, and specify freeform periods. Biped keys display as dots on the tracks. Footsteps display as squares that can be moved in time. Biped dynamics and footsteps work together; if footsteps are moved in time, Biped adapts leg keys and vertical positions to account for the editing. There are a couple of restrictions: you can't overlap footsteps on the same track, or move footsteps into negative time.

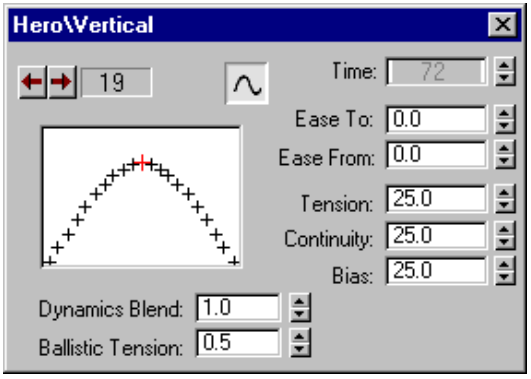
The Curve Editor can be used to make adjustments to biped keyframe animation. The Curve Editor is a Track View mode that allows you to work with motion expressed as function curves on a graph. It lets you visualize the interpolation of the motion, the object transformations that the software creates between the keyframes. You can see and control the motion and animation of the objects in the scene using the keys found on the curves. You can even draw curves directly on the graph.

Footsteps in Track View - Dope Sheet

By default, left leg footsteps are blue and right leg footsteps are green. Inactive footsteps are more saturated values of blue and green, active footsteps are pale blue and green. The left edge of each footstep block indicates when a foot touches a footstep (Touch State). The right edge of each footstep block indicates the last frame the foot is on a footstep (Lift State). The space between two footsteps indicates an airborne state for the foot (Move State). The period between the Touch and Lift States is the Plant State.

Editing Biped Keys in Track View

Track View is often used to change multiple keys simultaneously. For example, to change Dynamics Blend for all the vertical centers of mass keys, select all the vertical keys, display the TCB (Tension, Continuity, and Bias) dialog by right-clicking over the selected keys and change the Dynamics Blend parameter.



Right-click a key in Track View to display the TCB dialog.

Separate Tracks

By default, Biped stores all of the toe, foot, and calf keys in the thigh track. The finger, hand, forearms, and upper-arm keys are stored in the clavicle track. All the spine keys are stored in the spine 01 track. Although you can see all of these objects in the Track View hierarchy, they have no transform track, unless you enable them in the Separate Tracks group of the *Keyframing Tools rollout (page 2–962)*.

For example, if you rotate a biped foot, a key is created in the biped thigh track. This optimized approach works well in many cases. If your animation requires extensive hand and finger keyframing, turn on Arms on the Keyframing Tools rollout; all of the arm transform tracks are now enabled, down to the first finger link on each finger. Now if you delete an upper arm key, your finger-hand animation is preserved.

How Dynamics and Footsteps Relate

When footsteps are created, a footstep track displays in Track View – Dope Sheet. Footsteps appear as blue and green blocks, laid out to show their exact placement in time. If you compare the center of mass tracks to the footsteps, you will notice that frames where leading and trailing edges of the footsteps occur also have keys for the center of mass Body Vertical and Horizontal tracks. These keys contain dynamics information used by **character studio** to calculate airborne body position relative to gravity, and leg bend on landing and balance. Biped Dynamics is the reason you do not need a vertical center of mass key at the top of a jumping motion or at the bottom of dip when the biped lands from an airborne period.

Click and drag the middle of a footstep to move it in time. Click and drag one edge of a footstep to stretch the footstep in time.

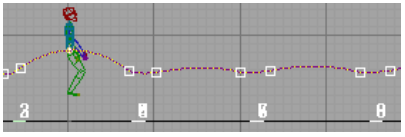
Note: Changing the duration of footsteps or moving them relative to one another may change the "support relationships" of the footsteps. Whenever the support relationships change, Biped generates new keys and deletes any existing leg keys in the airborne period between the edited footsteps.



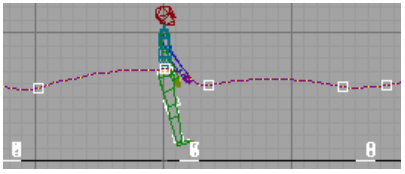
*Because of Biped Dynamics, no keys are necessary for the highest part of this jumping motion or for the dip when the biped lands; **character studio** calculates the trajectory of the Body. This image shows the center of mass trajectory. Keyframes are white squares on the trajectory.*



Simply change gravity (GravAccel on the Dynamics & Adaptation rollout), and the jumping motion is flattened. The biped looks like it's hopping on the moon.



By selecting half the footsteps in Track View - Dope Sheet and moving them to the right in time, the biped has to jump higher to account for the added time to the next footstep. Notice that the yellow dots, representing frames, are tightly bunched together; there are more frames in this airborne period.



*Here, the center of mass is moved in the Z plane. Now the biped heel never hits the ground; the biped appears to do a jump using just its toes. Here, **character studio** understands that you want to override the calculated trajectory and position the keyframe yourself.*

Moving Keys

You can use the time slider to move one Biped key past another. The result is as if the original key were deleted from its original frame and copied to the new frame. Simply drag the key to the new frame, then release the mouse to complete the move. The key's various Biped properties (ease values, TCB values, IK blend, and so on) are not changed.

You can move a key past either the following key or the previous one. You can also move it past multiple keys. If you move a key onto a frame where a key is already present, it replaces the original key.

In Track View, you can move a key past another using the transform track. You can also scale keys forward in time such that some keys are now past the keys that weren't scaled. If keys move into negative time, a warning reminds you that keys in negative frames can't be used for footstep animation: a workaround is to set new keys in the negative frames, instead of moving existing ones.

If you change the key to Euler control (on the Quaternion/Euler rollout) and then drag it past another, the key's tangent type is maintained.

If you add a layer (on the Layers rollout), set two keys, and then move the first one past the second one, this also works.

Freeform Animation

It is left to you to create all the keys in a freeform animation; Biped Dynamics is not active and does not recalculate body position. Balance Factor is active in a freeform animation. A completely freeform animation contains no footsteps.

To start an entirely freeform animation, simply create a biped and begin keyframing.

Note: Once you initiate a freeform animation, you cannot add footsteps to it in Footstep mode. You can, however, convert your freeform animation to a footstep animation using Convert on the Biped rollout.

You will often want a freeform period in a footstep sequence; for a walk then fall type of motion for

example. In cases like this, a freeform period is specified between footsteps in Track View using the *Footstep Mode dialog (page 2–1000)*. A combination of footsteps and freeform is often required when motion capture data is imported. A freeform period is created using controls in the Footstep Mode dialog and display as a yellow boxes between footsteps in Track View.

Right-click the footstep track to display the Footstep Mode dialog; here you specify a freeform period, select multiple footstep edges, and set footstep numbering display options.

Biped Color-coded Keys and Trajectories

Biped uses color coding for IK keys in three places: Track View, Track Bar, and when trajectories are displayed in the viewports. Also, depending on the types of keys, trajectories themselves can be color coded. The color coding helps you visualize when IK and FK are in use, and transitions between them.

This topic describes color coding of non-*COM (page 3–920)* tracks. For information about COM key color coding, check *here*.

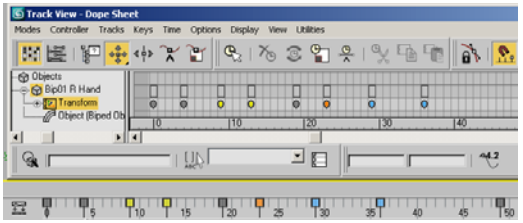
Note: Most of the information presented in this topic describes new functionality in 3ds Max 9 Extension 1. The only legacy functionality discussed here concerns color coding of IK keys in Track View.

Color-coded Keys in Track View and on the Track Bar

Biped IK keys as displayed in Track View and on the track bar use the following color scheme:

- orange for planted keys
- yellow for sliding keys
- blue for body-space IK keys with a non-zero value for IK Blend

All other keys use the default color: gray.



Color-coded biped keys in Track View and on the track bar

Visualizing Pivot Trajectories and Keys



To view trajectories for Biped parts in the viewports, go to the Motion panel and on the Biped rollout > Modes And Display sub-rollout, or on the Key Info rollout, turn on Trajectories.

To make Biped pivots easier to interpret and manipulate, Pivot trajectories and their keys are represented in the viewport during IK periods and display in the color of the associated key type:

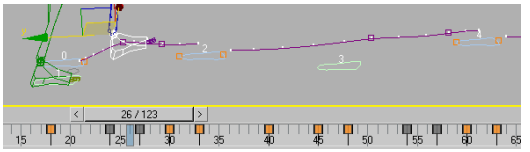
- orange for planted keys
- yellow for sliding keys
- blue for body-space IK keys

Tip: One way to create a body-space IK key is to set a regular IK planted or sliding key, and then, at that frame, on the Key Info follout > IK section, choose Body. Or you can convert an FK key that's already set to Body by increasing its IK Blend value above 0.0.

During FK periods the node trajectory and its keys display in purple (the same keys are gray on the track bar).

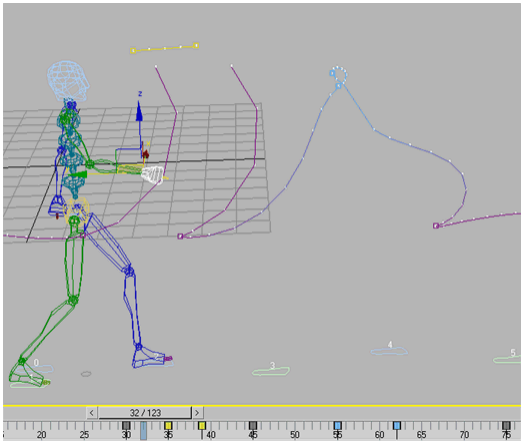
During FK/IK blend periods the trajectory of the resulting blend is displayed and its color is determined by the amount of the blend; for example, a blend of 0.5 to a sliding key results in a trajectory that is a 50:50 blend of yellow and purple.

Because pivots are located in all extremes of the hands and feet, their trajectories discontinuous when both FK and IK periods are employed. The following figure shows a simple example of this. During the planted IK periods, the pivot trajectories show just the keys on the heel and toe (because the keys are planted, each pivot stays in place), and during the FK periods the trajectories are derived from the foot's node pivot location. This gives you more-detailed information about the animation that is driving the foot.



Footstep trajectory

In following illustration, while the biped walks, the right hand reaches up to touch something in IK, as shown by the yellow trajectory, made by two sliding IK keys. Later, the biped touches its own face, using body space IK, as shown by the bright blue keys. The color transition in the trajectory from purple to blue indicates a changing blend from FK to body-space IK, showing the interpolated value of the IK Blend parameter.



Hand trajectory

Following is a chart showing the result of going from one key type to another. the last three rows of the chart show transitions between object and body space that produce pure FK. In the past, it was hard to know what type of trajectory these combinations produced. The new trajectory colors clarify the result.

In the chart, the letters have the following meanings:

- O: object space
- B: body space
- IK: a key where IK Blend > 0. It is always accompanied by the space: O or B.
- FK: a key where IK Blend = 0, no matter what space it's in

From -> To	Result	Trajectory Drawn	Trajectory Color
OIK -> OIK	OIK	Pivot	Yellow
BIK -> BIK	BIK	Node	Blue
FK -> FK	FK	Node	Purple
BIK -> BFK	BIK/BFK blend	Node	Blue -> Purple
OIK -> OFK	OIK/OFK blend	Node	Yellow -> Purple
OIK -> BIK	FK	Node	Purple
OIK -> BFK	FK	Node	Purple
BIK -> OFK	FK	Node	Purple

Note: When separate tracks exist for either a limb or its digits, the FK trajectory of the bone base is always drawn, regardless of whether the trajectory display is set to Bone Base or Bone Tip on the *Display Preferences dialog (page 2–944)*.

Biped Shortcuts

The table in this topic shows the default keyboard shortcuts for **character studio**. To customize your shortcuts, use the path specified above.

 Use the *Keyboard Shortcut Override toggle (page 3–872)* on the main toolbar to enable the **character studio** keyboard shortcuts.

All **character studio** keyboard shortcuts activate when the Motion panel is active and the Keyboard Shortcut Override button is active.

See Also

Keyboard Shortcuts (page 3–871)

Keyboard Panel (page 3–793)

Customize User Interface Dialog (page 3–792)

Action	Shortcut	Description
Activate Layer (Toggle)		
Bend Links Mode		
Body Horizontal		Selects the center of mass to edit horizontal biped motion
Body Rotation		Selects the center of mass to edit biped rotational motion
Body Vertical		Selects the center of mass to edit vertical biped motion
Buffer Mode		
Change Leg State	Alt+Ctrl+S	Toggles foot states of the selected leg at the current frame. View the state change in the leg states displayed on the Biped rollout.
Clear All Animation		
Clear Selected Tracks		
Collapse Layer		
Collapse Move All Mode Changes	Alt+M	
Copy Posture	Alt+C	Copies the posture of the selected biped objects to the clipboard.
Create Layer		
Delete Layer		
Display Biped Trajectory		
Display Bones		
Display Footstep Number		
Display Footsteps		
Enable Subanims		
Figure Mode		
Fix Graphs	Alt+Ctrl+F	Searches for any problem in the motion and prompts you whether or not to fix problems it encounters. Problems it looks for include overlapping keys, keys past the end of the footstep range, keys at negative frames, or illegal footstep timing. Click OK when prompted to fix these problems automatically.
Footstep Mode		
Free Mode		
In Place Mode		
In Place X Mode		
In Place Y Mode		
Load .bip File		
Lock Selected Keys (toggle)	Alt+Ctrl+L	Toggles locked keys (red) on and off for the selected leg or vertical track key at the current frame. In Track View, you can watch the key turn from red to gray, and back again, as you lock and unlock it.

Action	Shortcut	Description
Manipulate Subanims		
Mirror		
Mixer Mode		
Motion Flow Mode		
Move All Mode		
Next Layer		
Paste Posture	Alt+V	Pastes the posture from the clipboard onto the currently selected biped.
Paste Posture Opposite	Alt+B	Pastes Posture Opposite for the currently selected biped.
Play Biped	V	Toggles Biped playback.
Previous Layer		
Reset all limb keys	Alt+K	
Rubber Band Mode		
Save .bip File		
Scale In Transform (toggle)	Alt+Ctrl+E	This toggle strips the scale from the biped. Developers should use this when exporting biped objects as regular 3ds Max links through the 3ds Max SDK. Animators should not use this shortcut.
Scale Stride Mode		
Set Free Key		
Set Key	0	Sets a biped key.
Set Layer Snap Key		
Set Multiple Keys		
Set Planted Key		
Set Range	Alt+R	Resets the total animation length to the length of the current biped footstep range.
Set Sliding Key		
Smooth Twist Mode		
Toggle Biped Keys in Track Bar	Alt+T	
Track View Select end of footstep	Alt+D	Selects all right edges of the selected footsteps in Track View.
Track View Select entire footstep	Alt+S	Selects both edges of the selected footsteps in Track View.
Track View Select start of footstep	Alt+A	Selects all left edges of the selected footsteps in Track View.
Twist Links Mode		

* Not available for customization in the Customize User Interface dialog.

Using the Workbench

The Workbench is a customized version of Track View that you use for correcting and improving biped animation. It extends the functionality of existing curve editors by giving you different options for visualizing and manipulating curves, and provides filters to perform general rotation, position and other biped-specific operations. Curves can be analyzed for error conditions and filters applied to the identified keyframes, or selected biped tracks can be filtered to generally improve the motion of a track without errors.

See also

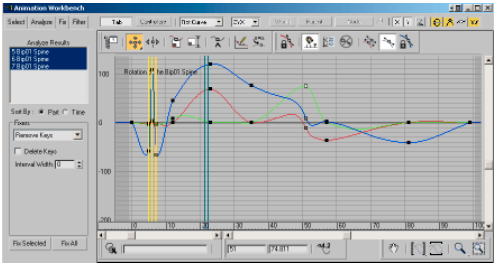
Navigating the Workbench (page 2–1010)

Workbench Workflow

Here are the overall steps you would follow to use the Workbench to analyze and fix a biped for problems

1. Create your animation and play it. Visually note anything that jumps out at you as being wrong or incorrect.
2. With the biped selected in the viewport, open the Workbench by clicking the Workbench button in the Biped Apps rollout on the Motion panel.
3. Select the biped body part that seems to have the problem. You can select it in the viewport, or from the list in the Select panel.

   Be aware that some biped objects share tracks. For example, all the objects in a leg or an arm share a single track.
4. On the Analyze panel, choose an analyzer from the drop-down list. Set the analysis conditions, then click Analyze. Any errors the analyzer encounters will be displayed in a list at the bottom of the Analyze panel. The Curve view also indicates the problem area by displaying a vertical bar.
5. On the Fix panel, select the errors you want to fix. Select the filter for the fix, adjust its parameters, and then click Fix Selected. You can select the errors one at a time, or in any combination. You can sort the errors by time or body part, depending on what kind of fix is needed.



Errors displayed as yellow lines in Curve View

   If the results aren't satisfactory, try changing the parameters and fixing again. Repeat the process until you find the combination of parameters that correct the movement. Or try a different fixer or filter (these are described later in this topic).

Animation Workbench Tools

The Animation Workbench provides three types of tools to use with biped animation: analyzers, fixers, and filters.

Analyzers

These are used to find out-of-the ordinary properties in your animation. Certain analyzers parse the waveform of your motion tracks and detect variation from the curves. There are two detectors available that do this; Noise and Spike. There are also analyzers for finding and fixing knee conditions, specifically knee wobble and extension.

For more information, see *Analyze Panel (page 2–1017)*.

Fixers

Fixers address problems found by the analyzer. The fixers work by running a specified operation over just the analyzer result. They don't change the complete animation, just areas around the problem.

For more information, see *Fix Panel (page 2–1020)*.

Filters

Filters are operations that are performed on the specified animation tracks. Unlike fixers, they affect the complete specified time interval, not just the analyzer result.

The Rot Smoothing, Rot Blurring, Rot Boosting, Adv Rot Smoothing, Pos Smoothing, Pos Blurring and Pos Boosting filters are all variants of one

type of filter. They all share the same parameters: Width and Damping.

For more information on these filter types, see *Filters Panel (page 2–1023)*.

Navigating the Workbench

The Animation Workbench, a customized version of Track View, uses some of the standard Track View controls, and adds new ones of its own. The areas of the Workbench interface include:

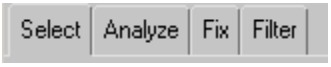
- The Tab panels: Select, Analyze, Fix, and Filter
- The Curve view, with Track View toolbars for key selection, manipulation, and view navigation.
- The Workbench toolbar, with tools for choosing Coordinate space, display of sub-animations, and tools to hide the Tab area and controller list.

Each section of the Workbench User has specific navigation techniques.

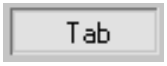
See also

Selecting Workbench Tracks (page 2–1011)

Navigating the Tab Panel



Simply click the tabs to move from one panel to another. The tools displayed change with each panel. The overall workflow is left to right for the tabs, and top to bottom within each panel. Selections are displayed in lists: clicking a name in a list selects that entry. Operations are usually performed by clicking a button at the bottom of the panel.

Tip: 

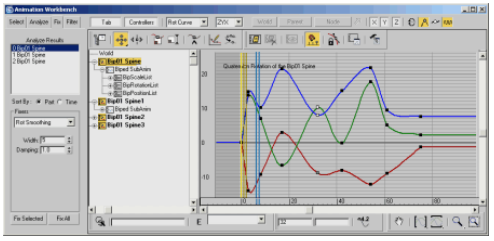
The Tab panels can be hidden once the operations are performed and you don't need to see the panels any more. Click the Tab button in the Workbench toolbar to hide the panel display.

Displaying the Controller List

Click the Controllers button in the Workbench toolbar to display the Controller List. When the controller list is displayed you can select multiple tracks as you like, right-click them and assign controllers using the pop-up menu. You can also get to properties for controllers such as Noise that don't have keyframes.



Click the Controllers button in the Workbench toolbar to display the embedded curve editor's Controller window. The biped body tracks shown are in the list are the ones that are selected in the *Select panel (page 2–1016)*, in addition to any scene objects that are attached to those body parts or are used as IK objects for those body parts. The controller list is similar to the one in the curve editor.



Controller list displayed

Navigating the Curve View

Use the standard 3ds Max Track View navigation tools to adjust the view of curves in the Curve view window. Zoom, Zoom Extents Horizontal and Vertical, and Pan are often used to get a better view of an entire curve, or a portion of a curve. As a default, the curve view automatically displays the curve of whatever biped object is selected in the viewport.

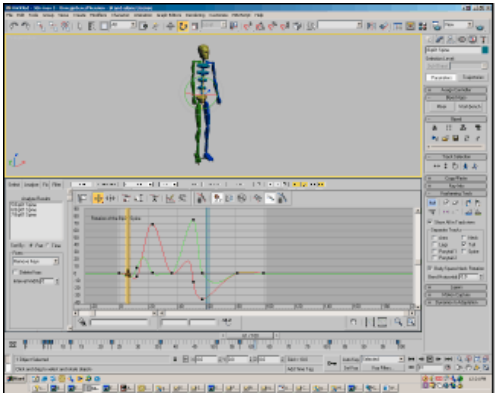
The Workbench window can be resized to make it easier to work with curves. You can float or hide

toolbars to give you even more room. You can dock toolbars left and right as well.

Docking the Workbench

You can dock the Workbench in a viewport. With the Workbench open, right-click the viewport's label and choose Views > Extended, then choose Biped Animation Workbench. The open workbench docks in the viewport.

Tip: This works best if you change your layout so there is at least one horizontal viewport.



Note: Unlike the 3ds Max Track View – Dope Sheet and Curve Editor, the Workbench cannot be docked below the viewport trackbar.

Show Layered Edit


You can affect a layered range of keys by turning on Show Layered Edit. This is a special Workbench mode that extends the edit to affect keys surrounding the one you are editing. It works like soft selection in 3ds Max but gives you greater control over the displacement. It can help you blend your edits into existing motion.

Selecting Workbench Tracks

As a default behavior, whatever biped part is selected in the viewport will have its curves displayed in the Curve view. You can also use the *Select panel (page 2–1016)* to select biped parts by name. The functionality of the Select panel is derived from the 3ds Max Select Objects dialog; you can invert selection sets or enter names into the Selection field to find a track.



You can also use the controller window as a selection mechanism. With the Controllers list visible you can click on track entries there to work on the curves.

Note: Sometimes a curve will not appear in the display at first. Click Zoom Extents Horizontal and Zoom Extents Vertical to make sure the curve can be seen.

The Workbench toolbar lets you choose the type of curve to display and the coordinate space in which it is displayed. For example, if you know you are dealing with rotational errors, choose an *Rot* or *Quat* type curve. If it is an error of positional data, choose *Pos*.

If you open the Track View – Curve Editor or expand the track bar while you have the Workbench open, you will see a simultaneous display. The views are all synchronized, so selecting a track in one graph editor selects it in all.

When using footstep animation, you can see the footstep keys by opening a Track View – Dope Sheet window.

Analyzing Curves

Once you have selected curves, you can analyze them to discover error conditions. You choose between four kinds of detectors; Noise, Spike, Knee Wobble or Knee Extension. You can also

select only a the portion of a curve to analyze. Perform the analysis by clicking the Analyze button. The errors are found based on the detector's parameter settings. If no errors are found, try adjusting the parameters and clicking Analyze again. Lowering the Standard Deviation value will yield more errors.

When an error has been detected, the biped part name appears in the Analyze Results list. If multiple errors are detected, the entry displays a number that tells how many errors were found. Errors are also displayed in the Curve View as vertical brown lines.

Once errors have been located, you can use the tools on the Fix panel to correct those keys and the keys around them. Or you can use the error location merely to identify the keys, then manually make changes using standard interpolation techniques. Right-clicking a key in the Curve view will display the properties dialog that allows for such adjustments.

 Results of the Analyze operation can be loaded or saved as a file using the Load or Save buttons at the bottom of the interface.

 Results of the last analysis can be cleared using the Clear button, also found at the bottom of the panel.

Fixing Curves

After you have selected curves and analyzed them for error, you can use the tools found on the Fix panel to process the curves and reduce the errors.

The fixes are determined by which type of fixer you choose, and the parameters you set for that fixer. Clicking Fix Selected or Fix All at the bottom of the panel performs the Fix operation.

In general, fixes are made either by changing the position or value of a key, or by removing the key.

The basic fixing techniques involve smoothing, blurring, boosting, or key deletion. This is similar to audio-editing software, where you view music as a waveform, and then edit the waves in various ways to alter the sound. Similarly, the motion of the biped body parts, as defined by position and rotation tracks, can be evaluated for error conditions regarding speed, angle, acceleration, or change of direction.

Of course, you can also fix curves manually. You can select the key on the curve and move it, using the standard Track View key buttons duplicated in the Workbench, such as Move, Slide, Scale or Delete Keys.

 You can achieve a layered range of keys by turning on Show Layered Edit. This is a special Workbench mode that extends the edit to affect keys surrounding the one you are editing. It works like soft selection in 3ds Max but gives you greater control over the displacement. It can help you blend your edits into existing motion.

Workbench User Interface

Select a biped body part. > Motion panel > Biped Apps > Workbench

The Animation Workbench is a tool for working with biped function curves. It is a customized version of Track View, with many of the same toolbars and buttons, plus some new ones.

In addition to the standard Track View features, the workbench has tools to help you select curves, analyze them for errors, and automatically fix those errors based on various settings. This is especially useful when working with motion-capture data, or other animation that has many keys.

Of course, you can also display and manipulate function curves for bipeds in the standard 3ds Max Track View and expanded track bar, as well. However, you have access to the automatic error-analysis tools only in the Workbench

Note: On the time slider or in Track View, you can move one Biped key past another. See *Moving Keys (page 2–1004)*.

The Animation Workbench displays rotation curves in local biped space for both quaternion and euler rotations. This improves the speed of curve manipulation since the whole biped doesn't need to be calculated in order to draw a single curve.

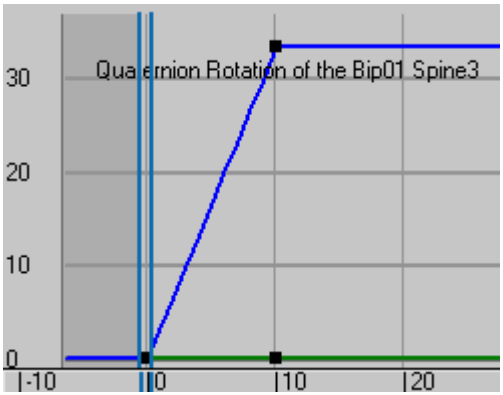
Note: This does not affect position curves, which are still evaluated in world space.

Note: Curve display will only change for certain objects in COM space: base of spine, neck, clavicles, upper arms, upper thighs, and feet.

Procedures

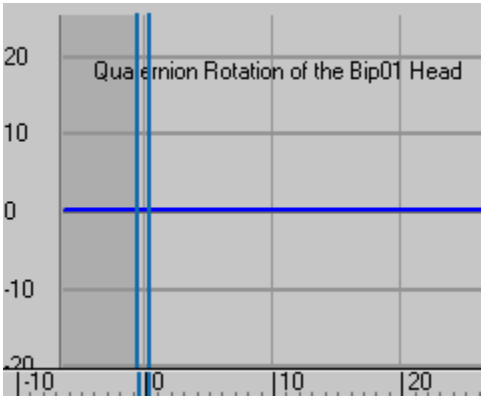
To display the local rotation curve of an animated spine link:

1. Create a biped and select its upper spine link.
2.  Set a key at frame 0.
3.  Go to frame 10 and rotate the spine link 30 degrees in the local Z axis.
4. Open the animation workbench.



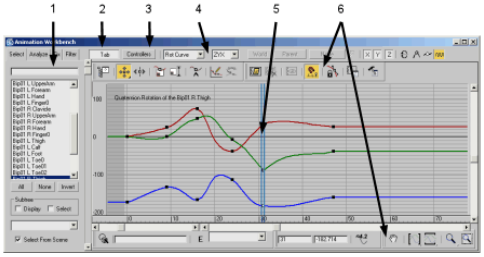
Only one curve is displayed because the upper spine rotates in local space by default.

5. Select the biped's head.



Notice how there are no curves in the workbench because the head doesn't inherit the spine link's rotation and doesn't have an internal rotation of its own.

Interface



1. Tab panels (Select, Analyze, Fix, and Filter).

2. Display Tab panels

3. Display Controllers list

4. Workbench toolbar

5. Curve View

6. Curve View toolbars (same as Track View toolbars).

Animation Workbench Tab Panel

The Animation Workbench Tab panel consists of four panels:

- **Select panel—**Provides tools to select curves for biped body parts. For more information on the Select panel user interface, see *Select Panel (page 2–1016)*.
- **Analyze panel—**Evaluates the curves for error conditions. For more information on the Analyze panel user interface, see *Analyze Panel (page 2–1017)*.
- **Fix panel—**Provides a variety of methods to be applied to the errors located by the analysis. For more information on the Fix tab user interface, see *Fix Panel (page 2–1020)*.
- **Filters panel—**These are a special type of fixer that perform an operation over each key that is selected. When no keys are selected it performs the operation over the whole curve rather than using the results of the last error analysis. Use this when you have errors visible in the motion, but the fixer refuses to find them. For more information on the Filter panel user interface, see *Filters Panel (page 2–1023)*.

Use the Tab button on the Workbench toolbar to display and hide the Tab panel.

Curve View

To the right of the Tab area is the Curve View, where keys and curves for the selected biped parts are displayed. This is the same as the Track View Key window.

Curve View Toolbars

Directly above and below the Curve View are some of the standard 3ds Max Track View toolbars. They give access to tools used to navigate the view and manipulate the keys. The toolbars below the Curve View are used for selection and navigation; toolbars above the Curve View are used to move, slide, and scale keys; to scale values; and to add new keys.

Animation Workbench Toolbar

The Workbench toolbar extends the toolset found in the 3ds Max Track View. The Workbench toolbar includes:

- **Tab—**Hides or displays the Workbench Tab panel, which contains the tools for automated error analysis and correction fixers and filters.
- **Controllers—**Hides or displays a controller-window hierarchy list. This is the same list you see in standard Track View with the same right-click menus used to assign controllers.
- **Curve Type—**Determines which curves are displayed for the selected object.

  Rot Curve, Rot Speed, Rot Accel, and Rot Jerk curves represent rotation.

  Pos Curve, Pos Speed, Pos Accel and Pos Jerk curves represent position.

  Note: Choosing *Rot Curve* or *Pos Curve* turns on Show Biped Rotation Curves or Show Biped,

Position Curves, respectively, in the Biped toolbar of the *Curve Editor (page 2–535)*.

Noise curves represent random position and rotation.

Each biped part can show any of these curves, except for the Horizontal and Vertical tracks, which show only positional curves, and the Turning track, which shows only angular curves.



The Axis order drop-down lets you choose the order in which the rotation curves are calculated. Curves in previous version of **Character Studio** were always drawn in XYZ. This ordering follows the same rules as the Euler Axis Order on the *Quaternion/Euler rollout (page 2–948)*.



When manipulating position curves, the Workbench allows the curves to be displayed relative to the World, the Biped Root, or any object in the scene as picked via the arrow button.

World—Displays curves relative to the World.

Biped Root—Displays curves relative to the biped root object (COM) .

Scene Root—Displays the curves relative to a selected object. Turn on the arrow button, then click the object in a viewport. The name of the object will be displayed on the button.

The **X, Y, and Z buttons —**Choose which curves of the selected object are displayed in the Workbench.

Note: Toggling X, Y, or Z also toggles the corresponding Biped toolbar button of the *Curve Editor (page 2–535)*.



The **180 limit—**Sets the Quat Curve to display as "clamped" between 180 and –180 degrees, instead of accumulated. Default=off.

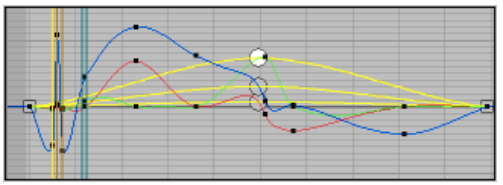


Draw While Moving—Sets the Workbench to update the curves as you move keys. When off, curves display as optimized lines until you release the mouse after moving or changing a key. Default=on.



Show Layered Edit —When on, displays a graphic tool for adjusting sets of keys along a curve within a range.

To set the range, click to highlight one of the square handles, and then drag it. As you drag, the handle follows the curve.



Handle follows curve.

To set a layered offset, click to highlight the circular handle, and then drag it to the desired position. The circular handle follows the curve unless you hold down the ESC key. If you drag while holding down ESC, you can move the handle off the curve, which scales the keys within the range.



Enable Sub Anims—When on, sets the Workbench to display the curves of the selected part's subanim in combination with the position or rotation curves. Subanims are list controllers which can be added to provide additional animation controls to biped body parts. Default=off.

Biped Key Manipulation

Interactive key manipulation behaves in the Workbench much as it does in Track View. The difference is that you can't move keys beyond other keys. This has always been a restriction of biped animation.

See also

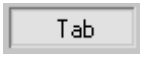
Working with the Workbench (page 2–1008)

Select Panel

Select a biped body part. > Motion panel > Biped Apps > Workbench > Select panel

The Select panel of the Animation Workbench provides tools for selecting bipeds or biped components. The selected biped part's animation tracks can be manipulated using curves displayed in the Workbench, or by using the Analyze, Fix, and Filter panels to automatically identify errors and fix them.

The objects whose curves might be displayed are set using the selection drop-down list. To change the curves that are displayed go into the Controller window and select the curves you want to see



Once the selection has been made, you can hide the entire Tab area by clicking the Tab button on the Workbench toolbar. This will give you more room to work on the curves in Curve View.

The functionality of the Select tab is similar to the standard 3ds Max *Select Objects dialog (page 1–78)*.

See also

Analyze Panel (page 2–1017)

Procedures

To display a biped body part curve in the Workbench, do one of the following:

- On the Select panel, make sure Select From Scene is turned on, then in a viewport, click any biped object.

  The object's curve is displayed in the Workbench Curve window.

- On the Select panel of the Animation workbench, click to highlight the name of the body part in the scrollable Selection list.
- On the Workbench toolbar, turn on Controllers. In the hierarchy list, highlight the body part track whose curve you want to see.

To hide or unhide the Tab panel:

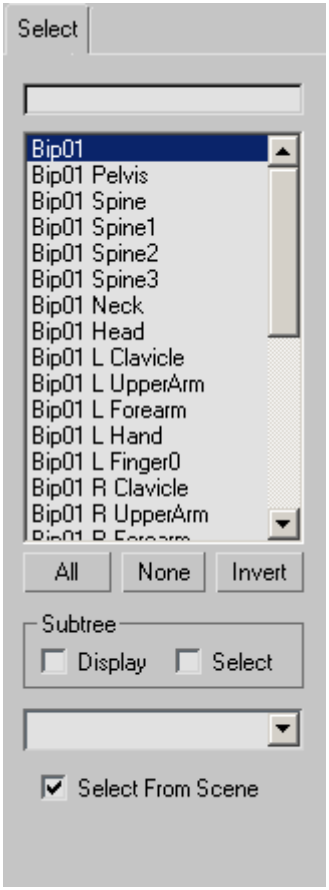
1. Click the Tab button on the Workbench toolbar.

   The Tab panel disappears from view.

2. Click the Tab button again.

   The Tab panel returns to view.

Interface



Selection field—Enter the name of the biped part you want to locate. That name (if found) will be highlighted in the Selection list. This is useful when you can't find an entry by scanning the list.

You can use the question mark and asterisk characters as wildcards, in order to select multiple objects at once.

Selection list—The Selection list displays all the body parts of all the bipeds in the scene.

You can use the Shift, Ctrl, and Alt keys to build selection sets, as in the rest of 3ds Max.

All—Selects every entry in the list.

None—Clears the selection so nothing is highlighted.

Invert—Selects everything not highlighted in the list.

Subtree group

Display—Indents the list to show the biped hierarchy. Default=off.

Select—When on, selecting one biped part selects both it and all of its children. Default=off.

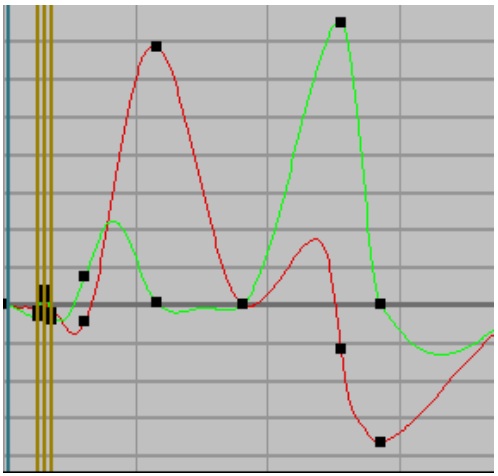
Named selection list— When you create a named selection in the scene, the named selection becomes available.

Select From Viewport—When Select From Viewport is on, you can click the biped body parts directly in the viewport, rather than from the list. The curves are immediately displayed in the Workbench Curve View. This is handy if you don't know the name of the part you need to select.

Analyze Panel

Select a biped body part. > Motion panel > Biped Apps > Workbench > Analyze panel

The Analyze panel provides tools to evaluate the curves for the selected biped parts, and review them for certain error conditions. It can spot spikes and noise in the curves, and locate specific keyframes that are responsible for discontinuous motion. The errors are displayed as brown lines over the curves, and are also listed at the bottom of the Analyze panel. You can then use the Fix panel to automatically fix errors, or you can fix errors by manually adjusting keys and curves in Curve view.



Brown lines on the curve show the errors.

Procedure

To analyze a curve:

1. On the Select panel, choose the biped object whose curve you want to correct.

   You can select the object from the Select list, in a viewport, or from the Controllers hierarchy.

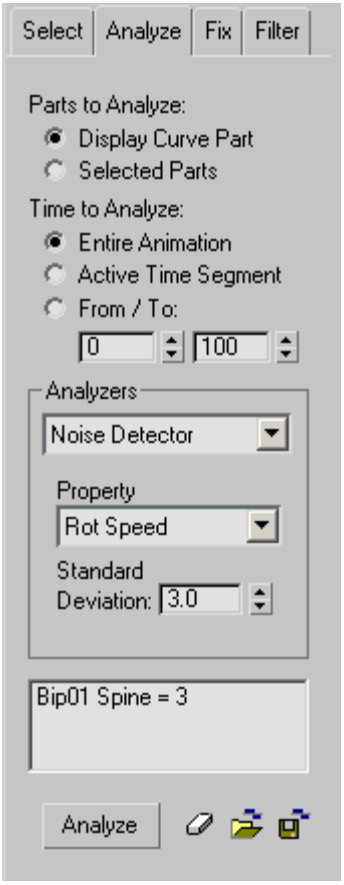
2. Click the Analyze tab, then on the Analyze panel, choose the analyzer you want to use from the drop-down list.

By default, there are two choices: Noise Detector and Spike Detector.

3. If you chose Noise Detector, choose the appropriate Property from the Property drop-down list.
4. Click Analyze.

   Any errors found are displayed in the error result list at the bottom of the panel. The errors also appear in Curve View as brown lines.
5. If no errors are found, try lowering the standard deviation value or make sure you have used the appropriate property when using the noise detector. For example, if you have a problem with the rotation of a leg, but you choose a Pos property, that error will not be detected.

Interface



Parts to Analyze— These options choose which parts to analyze:

- **Display Curve Part—**When this is turned on, the curve displayed is analyzed.
- **Selected Parts—**When this is turned on, the parts selected in the list or viewport are analyzed. Use this when you want to analyze an entire biped without displaying all of the curves in Curve View.

Time to Analyze—Sets the range to be analyzed. You can choose either:

- **Entire Animation—**Analyzes the entire animation.

Note: This is independent of the active 3ds Max time segment.

- **Active Time Segment—**Choose this to use the active time segment, as set by the 3ds Max Time Configuration dialog.
- **From / To—**These values let you specify a range with a particular start and end frame.

Analyzers group

Analyzers drop-down list—Lets you choose which analyzer will be used to evaluate the curves. Each analyzer can present its own individual settings. The default choices are Noise Detector and Spike Detector.

- **Noise Detector—**Finds any large change in the animation, either rotational or positional as determined by the standard deviation value. Can operate on specific curves as determined by the selection in the Property drop-down list, regardless of what's currently visible in the Workbench.

  Note: It's important to understand that noise isn't always bad. For example, in an animation of a person waving his hand, then suddenly punching: the movement from the wave to the punch, plus the pullback from the punch might be analyzed as noise, but it is noise that you want. The Workbench will identify the magnitude of these changes by showing you the acceleration and the jerk curves. By being selective with the fixers, you can smooth out only what you want to improve, and maintain the rest.
- **Spike Detector—**Finds any large change in the animation that also contains a change in direction (in quaternion space or position space). Can operate on all curves in the animation, regardless of what's currently visible in the Workbench. The analysis can be adjusted for Standard Deviation, which is the degree to which the animation departs from its overall pattern.
- **Knee Wobble —**Finds knees that wobble or shake when a foot is planted. Uses Frame and Fluctuation parameters to determine what is a wobble error.
- **Knee Extension—**Finds knees that overextend when a foot is planted. Uses a Knee Angle parameter to determine extension errors.

Property drop-down list—Lets you choose the criterion that the Noise detector uses to evaluate the curves errors. Use Rotational (Rot) properties to analyze rotational errors; use Position (Pos) properties to analyze move transform errors.

Note: This control is not available when using the Spike detector.

Options in this drop-down list include:

- **Rot Speed—**Looks for noise in the speed of the rotational angle.
- **Rot Accel—** Looks for noise in the acceleration of the rotational angle.
- **Rot Jerk—**Looks for noise in the jerk of the rotational angle.
- **Pos Speed—**Looks for noise in the speed of the position.
- **Pos Accel—**Looks for noise in the acceleration of the position.
- **Pos Jerk—**Looks for noise in the jerk of the position.

Standard Deviation—Lets you set the degree to which the animation can depart from its overall pattern. For example, in data that is somewhat noisy by design, the Noise detector can look for instances that are noisier. Lower deviation numbers mean higher sensitivity, and thus more errors found.

Error results list—Displays all errors that the analyzer finds. Each entry lists the name of the body part, and the number of errors found.

Errors are also graphically highlighted with a vertical brown line in Curve View. Clicking an error entry in the list highlights that error in Curve View.

Analyze—Click to perform the analysis.

 **Clear Results—**Deletes the results of the latest analysis.

 **Load Analysis File—**Loads the results of a previous analysis.

 **Save Analysis—**Saves the results of the latest analysis.

See also

Fix Panel (page 2–1020)

Fix Panel

Select a biped body part. > Motion panel > Biped Apps > Workbench > Fix panel

The Fix panel provides access to the tools that can be used to automatically fix the errors found by the *Analyze panel (page 2–1017)*. There are a variety of methods available to automatically correct errors found in the curves. Smoothing, blurring, and removing keys are all options. Or, if you prefer, you can correct the error manually by adjusting the interpolation per key, using standard 3ds Max methods. You can move the key on the curve, or right-click the key to access interpolation controls.

A fix is performed as a specific operation on a particular biped body part or set of parts, at a particular time. In general, fixes either change the value of a key or remove keys to produce corrected motion.

The results of last analysis are displayed in the Analyze Results list. Clicking an error in the list once displays the error as a yellow line over the curve; clicking the error twice updates the viewport scene, making the error the current frame. Errors can be sorted by Part or by Time.

Fixes can be applied to individual errors, selected errors, or to all errors at once.

See also

Filters Panel (page 2–1023)

Procedures

To fix an individual error:

1. Click to highlight the error in the Analyze Results list.
2. Choose the Fixer type you want to use to automatically correct the problem.
3. Click Fix Selected.
4. Observe the change in the Curve View.
5. Play the animation in the viewport to see if the error has been corrected.
6. If the error is still visible, try changing the parameters for the fixer, or choosing a different fixer. The Rot and Pos fixers attempt to solve the problem by changing the values of the keys; if they don't give good results, it might be more useful to use the Remove Keys fixer.

To fix multiple errors:

1. Highlight all the errors in the Analyze Results list.

   You can use the Ctrl and Shift keys to add to your selection.
2. Choose the Fixer type you want to use to automatically correct the problem.

3. Click Fix All.
4. Check your results as you would if you were fixing a single error (see the previous procedure).

To remove keys from curves automatically:

1. On the Select panel, highlight the biped part in the list.

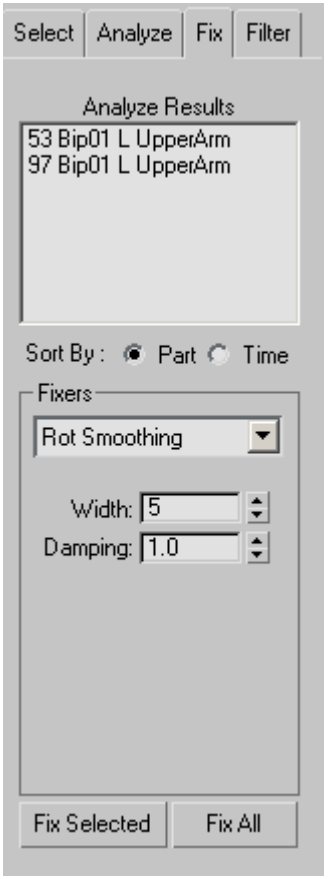
   Tip: If you don't know the name of the part, you can click it in a viewport.
2. On the Analyze panel, choose the appropriate Analyzer type and properties.
3. Click Analyze.

   The errors are listed in the error result list at the bottom of the panel.
4. On the Fix panel, choose the errors you want to fix.
5. In the Fixers drop-down list, choose Remove Keys.
6. Click Fix Selected.

   The removed keys will disappear from the list and the curves.

   Tip: You can use the Delete Keys option and Interval Width to create a threshold for deleting keys. If you turn on Delete Keys, the Interval Width will disallow the deletion of keys based on that value. Use this if too many keys are being deleted when you use Remove Keys.

Interface



Analyze Results error list—Displays all the errors the analyzer found for the selected body parts. The name of the Biped object is preceded by the frame number of the error. Clicking the error once displays a yellow line at that frame in the Curve View. Clicking it twice moves the current frame to match the error frame number: the viewports display that frame as well.

Sort by:—Lets you change the display of the error analysis results.

- **Part:—** Displays all the errors for each part listed together.
- **Time—**Displays all the errors sequentially, by frame.

Fixers—Determines the method used to attempt to correct the errors found by the analyzer. In general, try to match the fixer to the analyzer property. For example, if you have selected a Noise detector with Rot Speed, then choose a Fixer with Rot in the name as well. The exception to this would be when it is obvious that you'd prefer to remove the keys rather than modify them. In that case, choose a Remove Keys fixer.

Fixer options include:

- **Rot Smoothing—**Corrects rotational errors by smoothing the curve. Can be repeated on the same track for greater effect. Slower than blurring, but keeps more detail of the curve.
- **Rot Blurring—**Corrects rotational errors by blurring the curve. Good if you just want to keep the general shape of the animation, but don't want to keep the details. May remove too much detail in the animation.
- **Adv Rot Smoothing—**Works even better than the normal smoothing filter when it comes to only modifying large changes. Can be repeated on the same track for greater effect. It is the slowest filter, great at keeping minute details, which can be very important, especially with motion-capture data when you want to keep nuances.
- **Pos Smoothing—** Corrects move transform problems by smoothing the curve. Can be repeated on the same track for greater effect. Slower than blurring, but keeps more detail of the curve.
- **Pos Blurring—**Corrects positional problems by blurring the curve. Good if you just want to keep the general shape of the animation, but don't want to keep the details. May remove too much detail in the animation.
- **Remove Keys—**Fixes errors by deleting keys based on an interval setting.
- **Knee Wobble —**Fixes knees that wobble or shake when a foot is planted. Uses a Knee Angle value as criteria for error status.
- **Knee Extension—**Fixes knees that overextend when a foot is planted. Uses a Fluctuation value as criteria for error status.

Fixer parameters

Individual fixers display different parameters. These include:

- **Width—**Determines how much of the curve is affected around the keyframe.

  All the Fixers except for Remove Keys display this parameter.
- **Damping—**Determines how the fix is blended into the existing curve. Changes how effective the fixer should be. A higher value, the more dramatic the change. A damping value of 1.0 is normal behavior

  All the Fixers except for Remove Keys display this parameter.
- **Delete Keys—**Allows for selective key removal based on Interval Width. When Delete Keys is used, keys are removed only within the Interval Width. When Delete Keys is not turned on, Remove Keys will remove the selected keys without considering the Interval Width.

  This parameter is available only when Remove Keys is the fixer.
- **Interval Width—**Determines the number of frames around the error to consider when deleting keys.

  This parameter is available only when Remove Keys is the fixer.
- **Knee Angle—** Sets the rotation angle to determine what constitutes knee extension errors.

  If the knee is planted and its angle is less than or equal to the knee limit angle, then the

center-of-mass object is adjusted so that the knee angle doesn't exceed the limit.

This parameter is available only when Knee Extension is the fixer.

- **Frames—** Determines the interval to consider for fluctuation evaluation.

  This parameter is available only when Knee Wobble is the fixer.

- **Fluctuation—** Establishes the amount of change allowed before knee wobble is identified.

  This parameter is available only when Knee Wobble is the fixer.

  If the knee angle fluctuates greater than the fluctuation threshold for 2 knee angles within the frame interval threshold, then the knee's angle is changed, so that the wobble is eliminated.

Fix Selected—Applies the Fixer operation to the selection in the Analyze Results error list.

Fix All—Applies the Fixer operation to all the errors, whether they are selected in the Analyze Results error list, or not.

Filters Panel

Select a biped part. > Motion panel > Biped Apps > Workbench > Filters panel

Filters are simply operations that can be performed to different biped body parts. They are similar to fixers, but they operate over an interval of time rather than the results of the last analysis. Thus you can use a filter without analyzing. Certain filters share parameter types and functionality as fixers, though the values are not shared and are distinct for each filter. Use a filter whenever you want to perform a general operation over a body parts motion, like smooth out or boost up some noise."

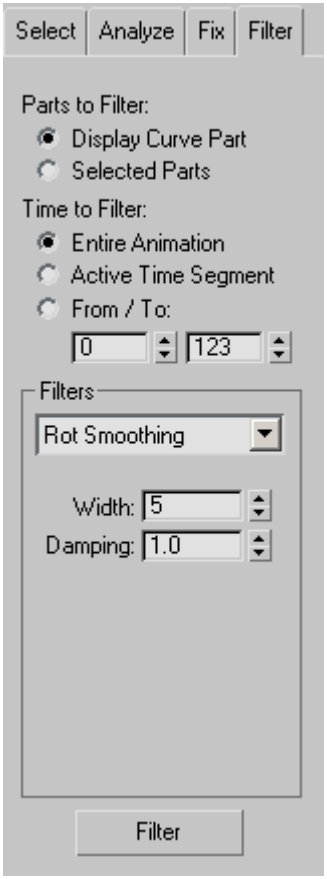
Procedure

To filter a track:

1. Choose the parts you want to filter: either just a particular curve, or an entire selection.
2. Choose the time range you want to filter: either the entire biped animation, the active time segment, or a custom range.
3. Choose the filter you want to apply. Adjust its individual parameter as required.
4. Click Filter at the bottom of the panel.
5. Observe the effect on the graph in the Curve view.
6. Play the animation in the viewport and observe the difference, if any.

   Click Undo if the results are unexpected or incorrect.

Interface



Parts to Filter—Sets the filter to act on either the display object or the entire selection of objects.

- **Display Curve Part —**Sets the filter to act on the displayed object curves.
- **Selected Parts—**Sets the filter to act on the active selection. Especially useful when working on the whole biped.

Time to Filter—Sets the range to be filtered to either the entire biped animation, the active time segment, or a custom range of frames.

- **Entire Animation—**Sets the time to filter to be the complete biped animation (disregarding the active time segment).
- **Active Time Segment—**Sets the time to filter to be the scene active time segment. You set the active time segment in the 3ds Max Time Configuration dialog.
- **From / To:—**Sets the time to filter to a custom range of frames.

Filters group

Filters list—Choose the filter to use from this drop-down list.

The available filter types include smoothing, blurring, boosting, key reduction, and subanims.

- **Rot** filters work in quaternion rotation space, and modify the quaternions of the specified animation tracks.
- **Pos** filters work in the specified positional coordinate system, and modify the positions of the animation tracks.
- **Blurring—**Uses basic Gaussian filters that take a weighted average over the width. Blurring filters are good for smoothing out noise, but can also over-smooth areas that aren't noisy.

  See *Blurring, Smoothing, and Boosting parameters (page 2–1025)*, below.
- **Smoothing** filters, on the other hand, are much better than blurring filters for keeping the general shape of the track or curve. They affect only areas that have big changes. The drawback of the smoothing filters is that they don't change the curve as dramatically as the blurring filters do, so sometimes you'll need to run the smoothing filter multiple times to smooth out a particularly noisy area.
- **Adv Rot Smoothing** works even better than the normal angular smoothing filter when it comes to modifying only large changes and not small ones.

  See *Blurring, Smoothing, and Boosting parameters (page 2–1025)*.

(There is no comparable "advanced position" fixer or filter.)

- **Boosting—**Boosting filters are the opposite of blurring filters. They increase, rather than decrease, changes in the track. Boosting filters are helpful if you want to make a track have more exaggeration or a little more jerkiness.

  See *Blurring, Smoothing, and Boosting parameters (page 2–1025)*.

- **Sub Anims—**The Sub Anims filter manages sub-animation of biped objects. Clicking a button that corresponds to a transform (Position, Rotation, or Scale) displays a dialog that lets you assign a controller to that transform. Assigning a controller in this way assigns it only to the selected biped parts. This is very handy for adding controllers to multiple biped parts all at once. Also includes the ability to enable or disable the three list controllers selectively, and avoids creating SubAnims, where they already exist.

  See *Sub Anim parameters (page 2–1025)*.

- **Key Reducer —**Creates tracks with fewer keys, by removing certain keys based on tolerance and key spacing parameters. This works similarly to how key reduction works during motion-capture import.

  See *Key Reducer parameters (page 2–1026)*.

- **Knee Wobble —**Corrects knees that wobble or shake when a foot is planted. If the knee is determined to wobble by the filter, the center of mass object is adjusted to eliminate the wobble.

- **Knee Extension—**Corrects knees that overextend when a foot is planted. If the knee extension is identified by the filter, the center of mass object is adjusted to eliminate the extension.

Blurring, Smoothing, and Boosting parameters

Width— Width is the size of the filter kernel width in frames. It tells you how much of the animation is taken into effect when filtering a specified keyframe.

Damping—The Damping value changes how effective the filter should be. A higher value, the more dramatic the change. A damping value of 1.0 is normal filter behavior

Sub Anim parameters



Enable—Turns Position, Rotation, and Scale subanims on and off for the selected biped parts.

Collapse—Adds the subanim animation of the selected biped part to that part's Transform controller.

Note: When you choose the Sub Anims filter, the Filter button at the bottom of the panel disappears. That's because the operations of this filter are immediate.

Don't Delete—Maintains the subanim in the list as it's collapsed onto the biped.

Per Frame—Sets keys at every frame of the collapsed controller.

The Sub Anims filter tools are also available from the Motion panel. You can assign controllers

to a biped subanim, and then collapse it by right-clicking and choosing Properties. The difference between using this in the Motion panel and in the Workbench is that you can apply subanims to multiple biped objects in one step when using the Workbench filter.

Key Reducer parameters

Create Key Per Frame—When on, the filter creates a key per frame for every selected track. Default=off.

Tolerance—Sets the maximum angular or positional deviation for a track (other than a COM track). Default=3.0.

Values are in units of translation for position tracks, and in degrees for rotation tracks.

Key Spacing—Sets the minimum number of frames between keys (other than COM keys).

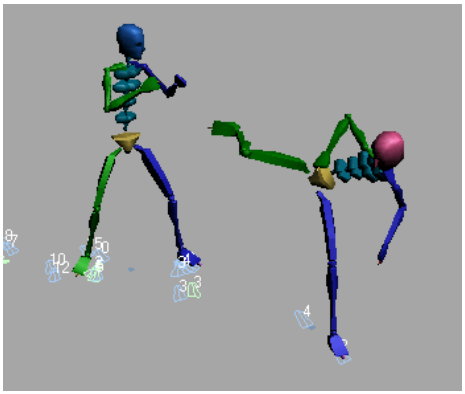
Tolerance is computed first, then Minimum Key Spacing computes further key reduction.

COM Parameters—These Tolerance and Key Spacing settings affect key reduction only on specified center-of-mass (COM) tracks (horizontal, vertical, or rotational).

Filter——Click to apply the active filter to the selected tracks.

Using Motion Flow

Select a biped > Motion panel > Biped rollout > Motion Flow Mode



Motion flow mode provides an area to graphically arrange clips into a network and tools to create and edit transitions between clips. Motion flow mode is used to organize clips into a network to animate one or more bipeds. The network of clips are joined together by transitions. A motion flow script is used to associate the network of clips with the biped. To animate one biped, you create a single motion flow script that uses a list of clips to animate the biped. To animate multiple bipeds or a crowd of bipeds, you can either use the random method of clip selection or a delegate–driven approach.

The random method simply picks clips at random and creates random scripts for each biped. This approach works well if the bipeds are standing still or are far apart and don't require collision detection: a crowd of cheering fans at a ball game, for example. Clip and transition percentages are set with the *Create Random Motion (page 2–1055)* command during motion synthesis.

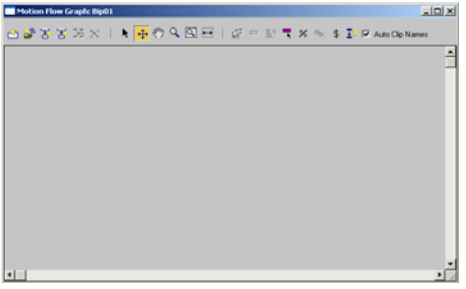
When dealing with a crowd, or multiple bipeds that are close together, the delegate driven approach is the best solution. This approach uses

many parameters to simulate moving crowds and incorporates collision detection, surface follow and other parameters. The delegate–driven method uses a network of clips, but instead of random selection. it bases clip selection on a delegate's speed and heading. In a delegate–driven crowd simulation, clips are arranged to follow a logical sequence. For example, the first clip could be a start walk clip, then a walk loop, then a branch to a turn right and turn left clip, then a slow to stop clip and so on. During motion synthesis, this arrangement is used to pick clips. If the software senses a collision ahead, the slow to stop clip is selected, or a veer to avoid clip is chosen. For a detailed breakdown of delegate driven crowd behavior, see *Crowd Animation (page 2–1154)*.

Placing Motions on the Motion Flow Graph

The *Motion Flow Graph (page 2–1045)* contains the tools you'll use to add clips to the dialog window, calculate optimized transitions, set random script transition values, move and delete clips, and display clip dependencies. Clips display as icons and transitions appear as arrows in the Motion Flow Graph.

The Motion Flow Graph displays when you click Show Graph on the *Motion Flow rollout (page 2–1045)* on the Motion panel.



Motion Flow Graph

The first step in Motion Flow mode is to add clips in the Motion Flow Graph for use in scripts. Clips represent all or part of a *.bip* file. Transitions represent different paths through the clips in the Motion Flow Graph. A transition looping back to the same clip represents a cycle or loop.

If a biped is using a shared motion flow, then the title of the graph dialog says "Shared Motion Flow Graph", followed by the name of the shared motion flow. Shared Motion Flows are used to control multiple bipeds with one shared motion flow.

There are two ways to create clips in the *Motion Flow Graph (page 2–1045)*. The first method lets you create a blank clip to which you later assign a *.bip* file. The second method lets you create multiple clips at once by choosing a several *.bip* files at once.

Procedures

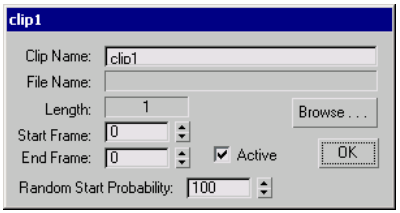
To create clips in the Motion Flow Graph:

1. Select a biped
2.  Open the Motion panel.
3.  Click Motion Flow Mode on the Biped rollout.
4.  On the Motion Flow rollout, click the Show Graph button to open the Motion Flow Graph.
5.  On the Motion Flow Graph toolbar, click Create Clip
6. Click a few times inside the Motion Flow Graph dialog.

   Clip icons will appear in the *Motion Flow Graph (page 2–1045)* that are named *clip1*, *clip2* and so on.
7.  On the Motion Flow Graph toolbar, click Select Clip/Transition.

8. Right-click over a clip in the window.

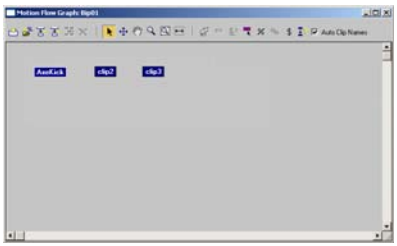
   A clip dialog is displayed.



9. Click Browse in the clip dialog and choose a *.bip* file.

10. Optionally, set the start and end frame and then click OK.

    The selected icon displays the clip name in the *Motion Flow Graph (page 2–1045)*; this clip can now be used for script creation.



11. Click Select Clip/Transition again and load *.bip* files for the other clips you created.

To create multiple clips in the Motion Flow Graph:

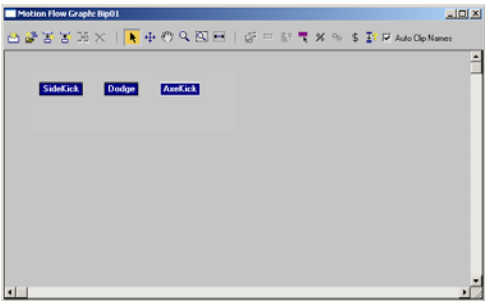
1. On the Motion Flow rollout, click the Show Graph button to open the Motion Flow Graph.

2. On the Motion Flow Graph toolbar, click Create Multiple Clips.

   The Open dialog appears. Use it to choose the location of your *.bip* files.

3. While holding down the Ctrl key, choose multiple clips in the Open dialog.

   Once processed, multiple clips appear in the *Motion Flow Graph (page 2–1045)* window.

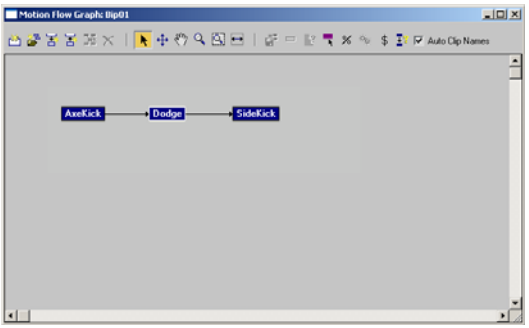


Note: The number of clips you can load depends on the amount of memory you have in your system.

Creating Transitions

Select a biped. > Motion panel > Biped rollout > Motion Flow Mode > Motion Flow rollout > Motion Flow Graph

Transitions, shown as arrows, link motion files (clips) together to create longer character animation and crowd simulations. Transitions can be created manually with the *Transition Editor (page 2–1051)* or automatically by the software.



Transitions can also be optimized for better flow. To create optimized transitions, use the Optimize Selected Transition on the Motion Flow Graph toolbar or Optimize Transition in the Transition Editor. By default, minimum motion loss is used to compute transitions. Optimized transitions use an

algorithm that minimizes foot sliding. Optimized transitions take longer to compute but yield very fluid results.

Crowd simulations can use dozens of motion clips so automatic creation of optimized transitions can be a big time saver. Whether you plan on animating one or many bipeds using Motion Flow mode you'll need transitions between the clips in the Motion Flow Graph.

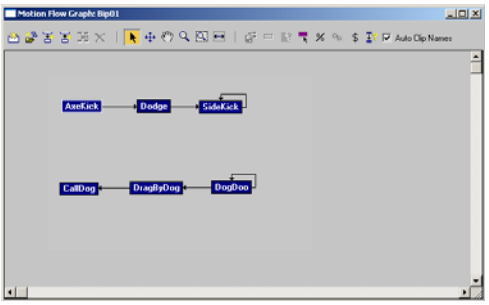
Procedures

To create transitions using From > To or To < From method:

1.  On the Motion Flow rollout, click the Show Graph button to open the Motion Flow Graph.
2.  On the Motion Flow Graph toolbar, click Create Multiple Clips.

   The Open dialog appears. Use it to choose the location of your *.bip* files.
3. While holding down the Ctrl key, choose multiple clips in the Open dialog.

   Once processed, multiple clips appear in *Motion Flow Graph (page 2–1045)*.
4.  Turn on either Create Transition From -> To or Create Transition To <- From on the Motion Flow Graph toolbar and then drag from one clip to another clip.



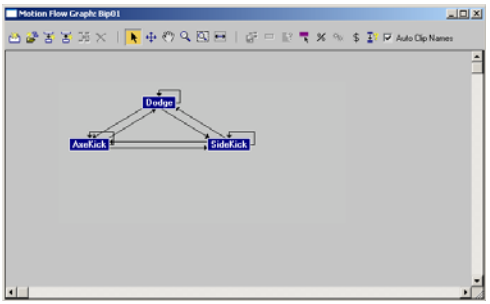
The top set of clips are linked using Create Transition From->To. The bottom set of clips are linked using Create Transition To<-From.

To create transitions using Create All Transitions:

1.  On the Motion Flow rollout, click the Show Graph button to open the Motion Flow Graph.
2.  On the Motion Flow Graph toolbar, click Create Multiple Clips

   The Open dialog appears. Use it to choose the location of your *.bip* files.
3. While holding down the Ctrl key, choose multiple clips in the Open dialog.

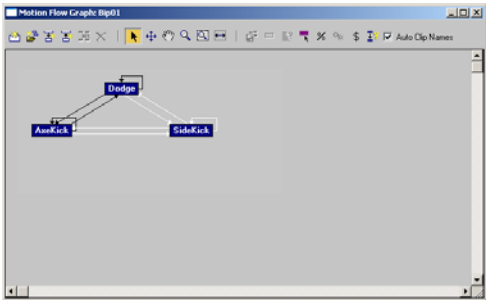
   Once processed, multiple clips appear in the Motion Flow Graph dialog.
4.  Click Select Clip/Transition and select all the clips you added.
5.  To create transitions between all the clips, turn on Create All Transitions on the Motion Flow Graph toolbar.
6. At the Biped dialog, click Yes to create transitions from each selected clip to itself.



Possible transitions from one clip to the next are automatically generated.

To optimize transitions in the Motion Flow Graph:

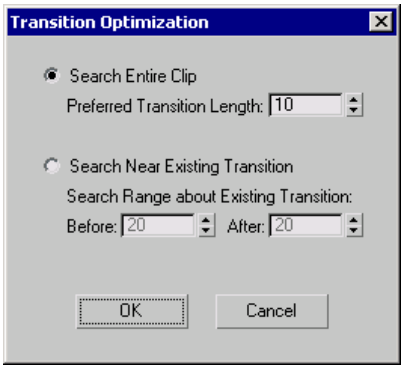
1. After adding several clips and transitions, click Select Clip/Transition and select one or more transitions.



Selected transitions display in white

2. On the Motion Flow toolbar, click Optimize Selected Transitions .

   The Transition Optimization dialog is displayed.



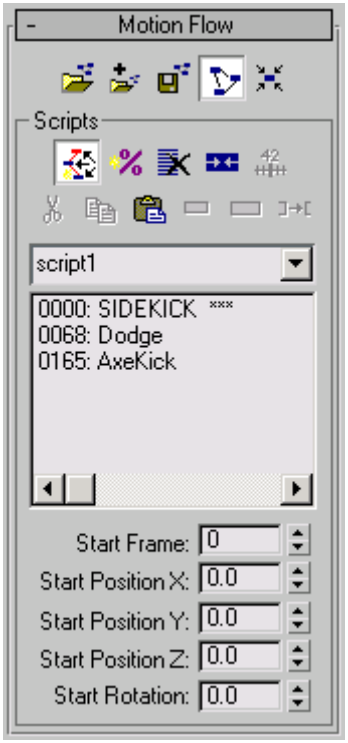
3. Click OK.

   A processing bar scrolls across the top of the *Motion Flow Graph (page 2–1045)*.

Creating a Motion Flow Script

Select a biped. > Motion panel > Biped rollout > Motion Flow Mode

A script is a list of clips (*.bip* files) that controls the character you are animating. To create a script, add clips to the *Motion Flow Graph (page 2–1045)*, then click Define Script in the Scripts section of the *Motion Flow rollout (page 2–1045)*, and click a sequence of clips in the Motion Flow Graph. Default transitions are assigned if no transitions exist between the clips. The clip names and starting frame numbers display in the list on the Motion Flow Script list.



A sample script as it appears in the Scripts group on the Motion Flow rollout

Random scripts vary the order of the clips found in the Motion Flow Graph. Scripts are run in a top-down order to animate the character. Click Play to view scripted motions on the biped.

Tip: As the length of a script varies, use **Alt+R** to set the Active Time in 3ds Max to the length of the script.

Once a script is put together, you have additional controls that let you adjust the starting frame, the starting X, Y, or Z locations and the starting rotation of the animation.

Procedure

To create a Motion Flow script:

1. Create a motion flow graph with three or more motion files.

   For information on how to set up the graph, see *To create multiple clips in the Motion Flow Graph (page 2–1028)*.

2.  In the Scripts section of the Motion Flow rollout, click Define Script, and then select a sequence of clips from the *Motion Flow Graph dialog (page 2–1045)*.

   As you select the clips in the Motion Flow Graph dialog, they are added to the script list.

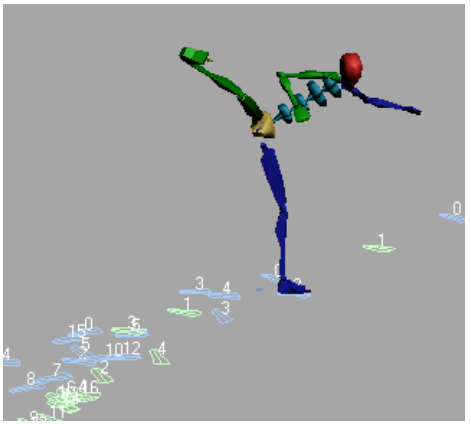


   Note: Transitions are automatically made in the Motion Flow Graph as you add clips to the script.

3. Give the script a new, descriptive name.

   Adding a descriptive name is less confusing if you ever append Motion Flow Editor files together.

4. Click the Play Animation button to see the biped move to the scripted clips.



5. Click Stop Animation. Change the Start Frame setting to **25** and move the frame slider.

   Now the biped will hold its start position until frame 25, then start moving.

6. Change the Start Position X:, Y,: or Z: setting.

   These settings let you control the position of the footsteps that the character follows.

7. Change the Start Rotation setting to **90**.

   Often, the footsteps in a *.bip* file are not oriented in the direction you want. This setting lets you adjust which direction the character faces.

Note: Each Start setting affects the entire animation. You cannot select a single entry in the script and make a change to affect only that clip.

Saving, Loading, and Appending Motion Flow Graphs

Select a biped. > Motion Flow mode > Motion Flow rollout

Scripts, transitions, and clip references are saved in a Motion Flow Editor.*mfe*, file, for later editing. A *.mfe* file allows you to save a Motion Flow and script from one biped and load it onto another biped.

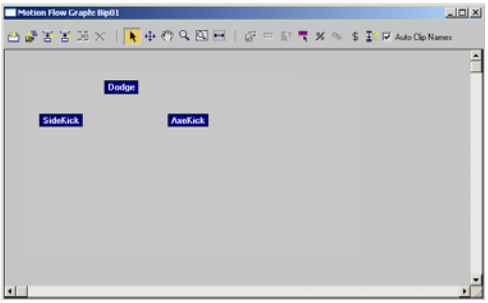
Saving, loading, and appending *.mfe* files are done from the Motion Flow rollout.

Procedure

To save Motion Flow Editor files:

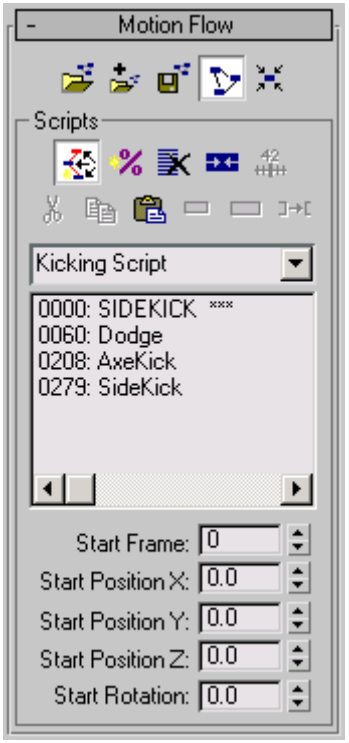
1. Create a Motion Flow Graph with three or more motion files.

   For information on how to set up the graph, see *To create multiple clips in the Motion Flow Graph (page 2–1028)*.



2. Create a script using the clips in the Motion Flow Graph.

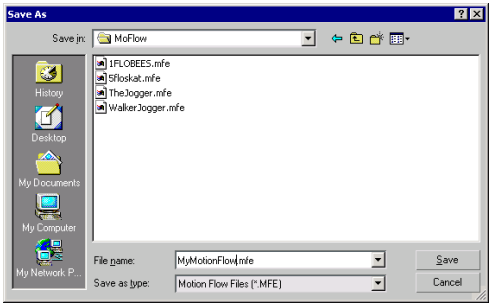
   For information on how to create a script, see *To create Motion Flow Scripts (page 2–1031)*.



A sample script called Kicking Script containing four clips

3. Click the Save File button to save a, Motion Flow Editor *.mfe*file.

   The Save As dialog is displayed.

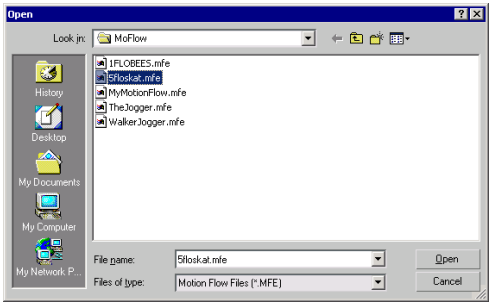


4. Specify a folder on your hard drive where you plan to store your Motion Flow Editor files.

To load Motion Flow Editor files:

1. Select a biped.
2. Open the Motion panel.
3. On the Biped rollout, click Motion Flow mode .
4. On the Motion Flow rollout, click Load File.

   The Open dialog is displayed.

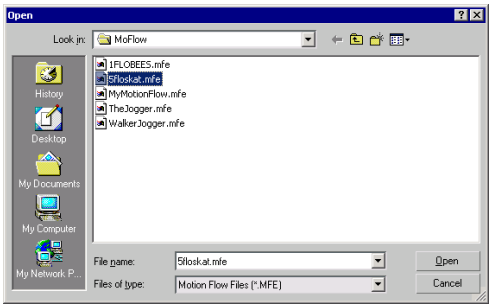


5. Navigate to the folder where your *.mfe* files are stored.
6. Choose the file to load.
7. Click the Play Animation button to see the biped move to the scripted clips.

To append Motion Flow Editor files:

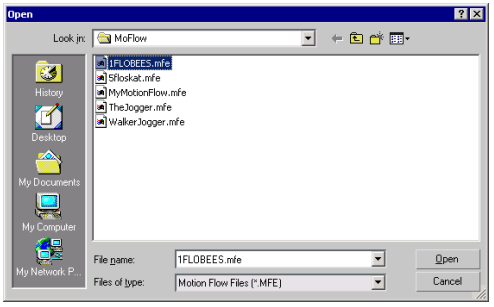
1. Select a biped.
2. Open the Motion panel.
3. On the Biped rollout, click Motion Flow mode .
4. On the Motion Flow rollout, click Load File .

   The Open dialog is displayed.



5. Select a *.mfe* file from the folder where your Motion Flow Editor files are stored.
6.  Click Append File.

   The Open dialog is displayed again.



7. Choose a different *.mfe* file.

   A new script appears in the script drop–down list and the appended sequence of clips appears directly below the original sequence of clips in the Motion Flow Graph dialog, so you may have to scroll down to see it.

Customizing Transitions

Select a biped. > Motion panel > Biped rollout > Motion Flow Mode > Motion Flow rollout

A good transition links two clips together seamlessly, the motion through the transition should appear natural, as though the motion was captured as one long motion sequence. If clips do not transition smoothly, you may find it necessary to edit or customize your transitions. There are two ways to manually edit transitions.

- In the Motion Flow Graph, you can right click a transition arrow.
- If you have defined a script, highlight a clip and click the Edit Transition button.

Transitions can be edited automatically by using the Optimize Transition features. When you create a script, default transitions are set between the clips. Default transitions use minimum motion loss and are quick to compute. However, the best quality transitions are the optimized transitions. Optimized transitions use a minimum foot sliding algorithm to compute the transition and yield very good results.

Procedures

To manually customize transitions between two clips:

1.  On the Biped rollout, select a biped and turn on Motion Flow mode.
2. Create or load a script with at least two clips.

   For information on how to create a script, see *To create Motion Flow Scripts (page 2–1031)*
3. Choose the first clip in the clip list in the Scripts section on the Motion Flow rollout.
4.  Click Edit Transition to display the Transition Editor for the selected clip and the clip following it in the list.
5. Scrub the Frame spinners in the Ghost areas of the Source Clip and Destination Clip, to find a place in both clips to start the transition.

   Two stick figures appear to help you find an appropriate start frame. The yellow figure represents the Source Clip and the red figure represents the Destination Clip.

The stick figures can be positioned far apart during this process. Look for body motions in the two clips that will transition well.

6. Click Set Start Frame in both the Ghost areas when a good Start Frame is located.

   Character studio repositions the destination clip for a best body fit between the two clips.

7. Set the transition duration in the Length field.

   A transition duration of 10 to 25 frames is normal.

8. Click OK.

9. Click Play Animation or scrub the time slider to view the transition.

To automatically customize transitions between two clips:

1. On the Biped rollout, select a biped and turn on Motion Flow mode .

2. Create or load a script with at least two clips.

   For information on how to create a script, see *To create Motion Flow Scripts (page 2–1031)*

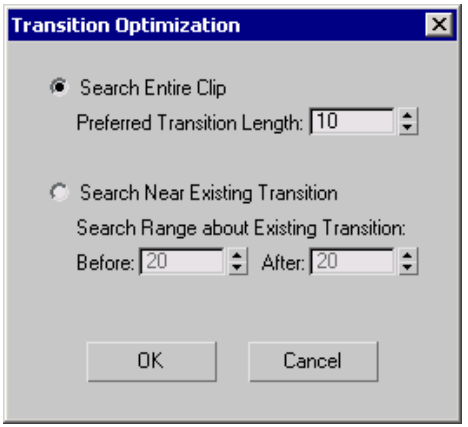
3. Choose the first clip in the clip list in the Scripts section on the Motion Flow rollout.

4. Click Edit Transition to display the Transition Editor for the selected clip and the clip following it in the list.

5. At the upper-left corner of the Transition Editor, click Next Transition in Script.

6. Click Optimize Transition.

   The Transition Optimization dialog is displayed.



7. Choose either Search Entire Clip or Search Near Existing Transition, and click OK.

8. Click OK on the Transition Editor dialog.

9. Click Play Animation or scrub the time slider to view the transition.

Creating Random Motion

You can randomly traverse clips in a motion flow graph to animate one or more bipeds using controls in the Create Random Motion dialog. You set parameters for random motion in the Motion Flow Graph, in the clip and transition dialogs, as well as in the Create Random Motion dialog.

You create random motion by first adding clips to the Motion Flow Graph and adding transitions between the clips. Clips and transitions are then given percentages, which are used to create random motion for one or more bipeds. You can manually control the “weighting” for a possible start clip, transitions, and frame range. This allows you to animate multiple bipeds in a crowd scene for example. A different script is created for each biped.



Note: If you want to create random motion for multiple bipeds, they must be sharing a motion flow.

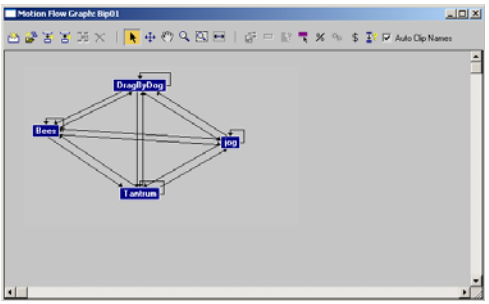
Procedure

To create a random script for one biped:

1. On the Biped rollout, select a biped and turn on Motion Flow mode.
2. On the Motion Flow rollout, click Show Graph

   This displays the Motion Flow Graph.
3. On the Motion Flow Graph toolbar, click Create Multiple Clips.

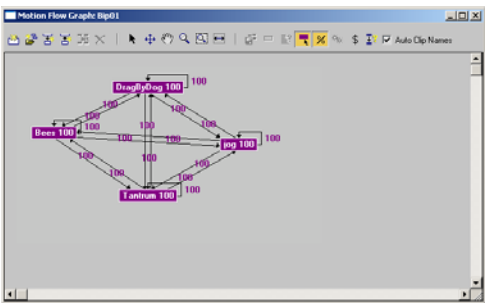
   The Open dialog appears where you can choose the location of your *.bip* files.
4. After adding several clips and transitions, click Select Clip/Transition and select all the clips.
5. On the Motion Flow Graph, click Synthesize Motion Flow Graph and click Yes when asked if you want to create transitions from each selected clip to itself.



6. On the Motion Flow Graph toolbar, click Select Random Start Clips.

   All the transitions linking clips are weighted at 100, meaning all the transitions have a even chance of starting first.
7. Region select all the clips in the Motion Flow Graph.

   The clips all turn purple and are assigned a weighting of 100. All the clips have an even chance of starting first.



8. In the Scripts section on the Motion Flow rollout, click Create Random Motion.
9. On the Create Random Motion dialog set a value for Minimum Animation Length.

   This should be long enough to include the clips you want in the animation.
10. Click Create.

A random script is created based on clips in the Motion Flow Graph. You can vary clip and transition percentages in the clip dialog or Transition Editor to favor a clip or transition if you like.

11.  Click Play Animation or scrub the time slider to view the transition.

Details of Random Motions

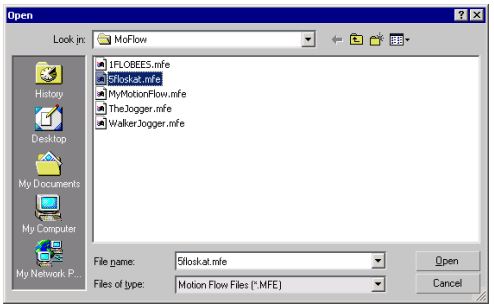
When choreographing a random motion for your biped, the default "weighting" for all clips and transitions is set to 100. This ensures that all clips and transitions have an equal chance of being randomly chosen. You can adjust the weighting of transitions in the Motion Flow Graph or Transition Editor to give transitions a higher probability of being selected. You can also set a specific start clip for your biped to use while the remainder of the clips in the script are randomly chosen from the Motion Flow Graph.

To customize transition "weighting" and setting a start clip:

1.  Oon the Biped rollout, select a biped and turn on Motion Flow Mode.
2.  On the Motion Flow rollout, click Show Graph.

   This displays the Motion Flow Graph.
3.  On the Motion Flow rollout, click Load File

   The Open dialog is displayed.



4. Choose a *.mfe* file from the folder where your Motion Flow Editor files are stored.

   The clips and transitions appear in the Motion Flow Graph. Each transition displays its "weighting" value, 100 by default.
5.  Click Select Clip/Transition and right-click a transition that you want to adjust.

   The Transtion Editor dialog is displayed.
6. Adjust the Probability setting to give the transition a higher or lower probability of being randomly chosen for the motion flow script.

   Note: The Probability range is 0 to 100.
7.  On the Motion Flow Graph toolbar, click Select Random Start Clips.
8. Select the clip you want used as the start clip.

   The clip turns purple and is assigned a weighting of 100. The selected clip will be used first by any biped using the random motion script.
9.  In the Scripts section on the Motion Flow rollout, click Create Random Motion.
10. On the Create Random Motion dialog, set a value for Minimum Animation Length.

    This should be long enough to include the clips you want in the animation.
11. Click Create.

A random script is created based on clips in the Motion Flow Graph and the Minimum Animation Length.

12. Click Play Animation or scrub the time slider to view the transition.

Unifying Motion

The Create Unified Motion command allows you to create one motion from a script. The entire unified motion is then available when you exit Motion Flow mode. Create Unified Motion converts a script into a freeform unified motion. The created motion replaces animation present when the Motion Flow Mode button is turned off. This is useful if you want to make changes to the biped like raising its arms or turning its head.

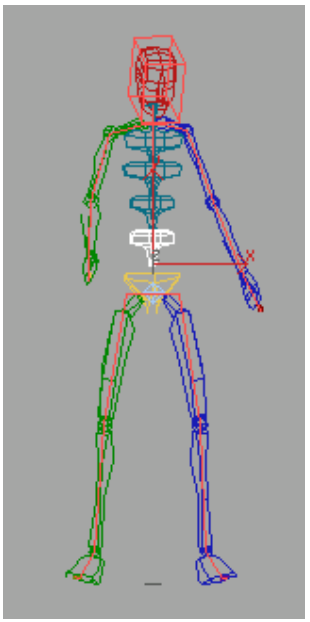
Procedure

Example: To create a unified motion:

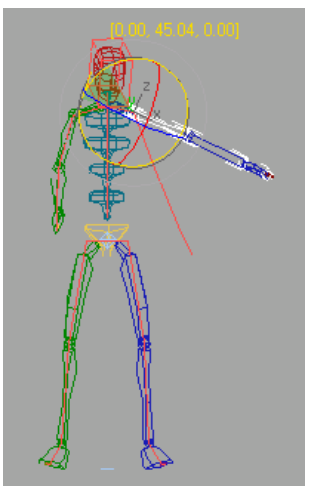
1. Create a biped.
2. On the Biped rollout, turn on Motion Flow Mode.
3. On the Motion Flow rollout, click Show Graph.

   This displays the Motion Flow Graph.
4. On the Motion Flow rollout, click Load File and choose *1flobees.mfe* from the Open dialog.
5. Click Create Unified Motion.
6. On the Biped rollout, turn off Motion Flow mode.
7. Open the Layers rollout and click Create Layer.

   A new layer is created and you see a red skeleton appear on your biped.



8. Auto Key Activate the Front viewport and turn on Auto Key.
9. At frame 0, select and rotate the biped's upper left arm about the Y axis.



10.  Turn off Auto Key and click Play Animation.

Sharing Motion Flow

Shared Motion Flow (page 2–1056) allows you to assign one motion flow to multiple bipeds or crowds. Rather than building a motion flow network of clips for each biped, you can create a motion flow with all the clips and transitions to animate multiple bipeds. Random motion creation will use each biped's own motion flow. If a biped's motion flow happens to be a shared motion flow, then the shared motion flow will be used to compute random motion.

A biped that shares a motion flow shares only the clips you have added to the Motion Flow Graph. Scripts are unique to each biped, although the scripts point to the clips of the shared motion flow. You can manipulate each biped's motion flow and scripts in the usual ways. You can create random motion on a biped that shares a motion flow, or create a motion flow script using the crowd system.

There are a few indicators in the user interface that show if a biped is using a shared motion flow. If it is, a white circle surrounds the Shared Motion Flow button when you edit that biped. If you edit that biped's motion flow graph, the title of the graph dialog will say "Shared Motion Flow Graph", followed by the name of the shared motion flow.

Note: All bipeds in a shared motion flow should have the same leg scale and structure. Adaptation for differently sized bipeds does not occur in a shared motion flow. If you want differently sized bipeds in a crowd, then create a shared motion flow for *each* size.

Tip: The only way to save a shared motion flow, along with the bipeds sharing it, and keep everything hooked up correctly is to save it all in a *.max* file.

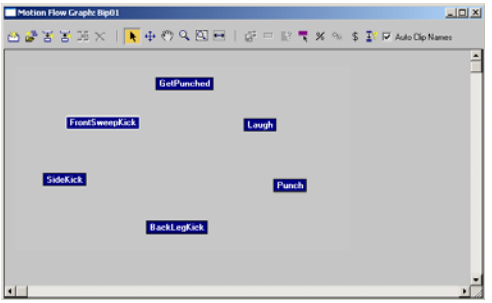
Procedure

Example: To share a random motion flow among multiple bipeds:

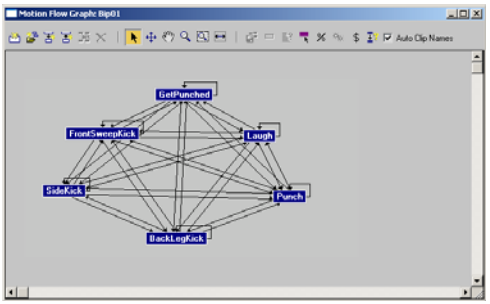
1. Create three bipeds.
2.  On the Biped rollout, turn on Motion Flow Mode .
3.  On the Motion Flow rollout, click Show Graph.

   This displays the Motion Flow Graph.
4.  Click Create Multiple Clips and add at least four clips to the Motion Flow Graph.

   The Open dialog appears where you can choose the location of your *.bip* files.
5.  After adding the clips, click Move Clips and arrange the clips so they are not so close to one another. This will make it easier to see the transitions later.



6.  Click Select Clip/Transition and region select all the clips.
7.  On the Motion Flow Graph, click Synthesize Motion Flow Graph and click Yes when asked if you want to create transitions from each selected clip to itself.

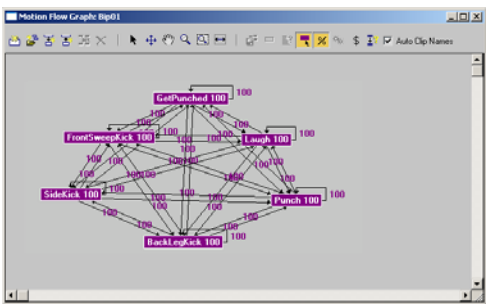


8.  On the Motion Flow Graph toolbar, click Select Random Start Clips.

   All the transitions linking clips are weighted at 100, meaning all the transitions have a even chance of being randomly used.

9. Region select all the clips in the Motion Flow Graph dialog.

   The clips all turn purple and are assigned a weighting of 100. All the clips have an even chance of starting first.

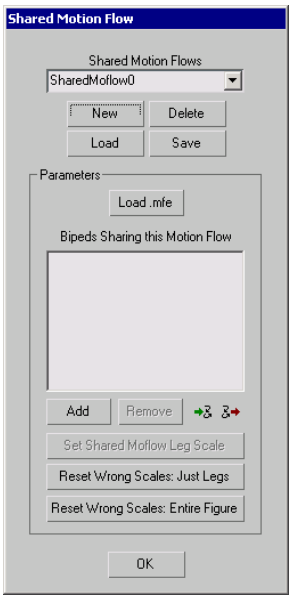


10.  On the Motion Flow rollout, click Shared Motion Flow .

    The Shared Motion Flow dialog is displayed.

11. On the Shared Motion Flows dialog, click New.

    A new shared motion flow is created.



12. On the Shared Motion Flows dialog in the Parameters group, click Add.

13. On the Select dialog, choose all the bipeds.

    The bipeds are added to the list. If you add clips to the motion flow graph they will be shared by the bipeds.

14. Click OK to close the Shared Motion Flow dialog.

15.  In the Scripts section of the Motion Flow rollout, click Create Random Motion.

    The Create Random Motion dialog is displayed.

16. On the Create Random Motion dialog, turn on Create Motion for all bipeds sharing this motion flow.

    Failure to do this will result in the motion flow being added only to the selected biped.

17. Click Create.

    A random script is created for each of the bipeds that shared the motion flow.

18.  Click Play Animation or scrub the time slider to view the transition.

To apply a shared motion flow to a different group of bipeds:

New in **character studio** is the ability to save and load shared motion flow setups. With this facility, you can create and refine a crowd simulation using bipeds that don't have meshes attached, thus avoiding the computational overhead that Physique requires. Once the simulation is ready, you can save the shared motion flow setup, apply it to a similar set of bipeds with Physique applied, and then render the final animation with no further crowd setup required.

For best results, the two sets of bipeds should have the same number of members and be of comparable structure and size, and the biped structures should have the same root names.

1. Create a biped/crowd simulation using a shared motion flow.

   For details, see *Using bipeds in a crowd simulation (page 2–1175)*.

2. On the Shared Motion Flow dialog, click Save and save the shared motion flow in *.smf* format.

3. Set up a new scene containing the same number of bipeds, each with a character mesh with Physique applied. The easiest way to do this is to set up one such character and then clone it.

   Make sure the bipeds have the same root names (e.g., Bip01, Bip02) as those in the original setup.

4.  Select one of the bipeds, and on the Motion panel, turn on Motion Flow mode, and then click Shared Motion Flow.

5. On the Shared Motion Flow dialog, click New, and then click Load and load the shared motion flow from step 2.

   The shared motion flow is loaded and applied to the new set of bipeds. The biped names appear in the dialog list. If "(wrong scale)" appears after any biped names, use the Reset Wrong Scales buttons to correct the scales.

6.  Click the Put Multiple Bipeds In Motion Flow button to place all of the bipeds in the list in Motion Flow mode, and then click OK to exit the dialog.

7. Play the animation to ensure that it's correct, and then render.

Refer to *Working with Crowd Animation (page 2–1154)* for information regarding crowd animation.

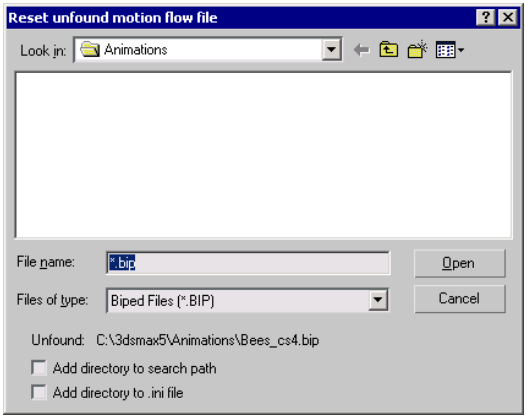
Setting Up Paths for Motion Flow Files

On a Motion Flow Graph, the location of the referenced *.bip* files is saved in the *.mfe* file. If a *.bip* file cannot be found, the program looks to the motion flow directory specified in *\plugcfg\biped.ini*. By default, this directory is *MoFlowDir=<maxdir>\cstudio\moflow*.

If a *.bip* file is not found in the motion flow directory specified in the *biped.ini* file, then 3ds Max looks for the file in the same location as the currently loaded *.max* file. If it is still not found, the program appends the local path of the referenced file to the path to the current *.max* file. The file is saved back to the same location in which it is found.

If a referenced *.bip* file cannot be found in its current location, there are two ways to resolve this.

- From the Reset Unfound Motion Flow Files dialog, direct 3ds Max to the folder where the missing *.bip* file resides.



- Find and move the missing *.bip* file to the specified Motion Flow folder.

Note: When you find the *.bip* file, you can turn on the Add Directory To .INI File option and the directory is automatically added to your *biped.ini* file. Since the *biped.ini* file accepts multiple search paths, the new directory is added as another *MoFloDir=* line.

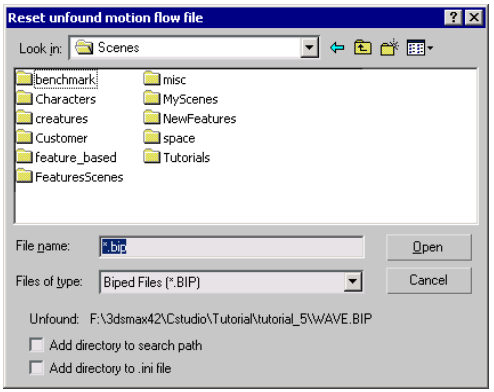
The program will search the directories in the order they appear and will use the first instance of the file that it finds. When network rendering, the file names need to be UNC–compatible.

Procedure

To load a Motion Flow scene from another system:

1. Open a Motion Flow scene that was created on a different system.

   If referenced *.bip* files cannot be found, the Reset Unfound Motion Flow File dialog is displayed.



2. If any of the referenced *.bip* files are on your local drive, navigate to the folder where they reside and select the *.bip* file.
3. Turn on the Add Directory To .ini File option, and click Open.

   Turning on the Add Directory To .ini File option automatically adds another *MoFlowDir=* line to your *biped.ini* file.

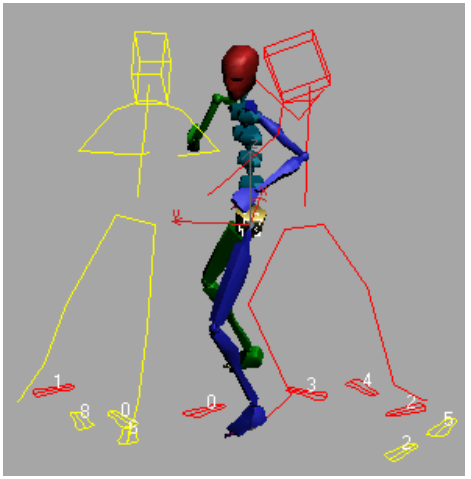


   Note: Turn on Add Directory to Search if you want to search for referenced *.bip* files during your current session, but do not want to add the directories to the *biped.ini* file.

4.  Click Play Animation after paths for all the referenced *.bip* files are properly set.

Motion Flow User Interface

Select a biped. > Motion panel > Biped rollout > Motion Flow Mode



*The Transition Editor uses **ghosts** to represent the source (yellow) and destination (red) clips. Use these stick-figure ghosts to judge body position and set a likely start frame in both clips for the transition.*

In Motion Flow mode, you combine BIP files, using either velocity-interpolated transitions or optimized transitions that compute minimum foot sliding, to create longer character animations. First, clips are added and referenced to BIP files in the Motion Flow Graph. These are then selected to create a script in the Motion Flow Script list. The Transition Editor is used to adjust transitions between BIP files. Scripts, transitions, and clip references are saved in a Motion Flow Editor file (MFE) for later editing.

Random Motion and Crowds

Random motions can be created for one or more bipeds using the Create Random Motion command. You can animate a crowd of bipeds for example. For a crowd, you must share one motion flow script among many bipeds.

If you are driving a crowd using delegates and behaviors, then rather than a completely random motion, the software picks appropriate clips based on the delegate's speed and direction. If the delegate slows to a stop, the software will find and use a clip that slows to a stop, if one exists.

In all crowd simulations you must load clips and create transitions before synthesizing the crowd motion. Often, many clips are used to synthesize crowds. Automatic transitions relieve you of having to create transitions between clips manually.

One additional rollout displays in Motion Flow mode:

- *Motion Flow (page 2–1045)*: Display the Motion Flow Graph; load and save motion flow files (MFE). The Motion Flow rollout also contains a Scripts section where you can create scripts, edit transitions, create a unified motion, and create random motions for the bipeds.

Note: In Motion Flow mode, some **character studio** controls are disabled.

Workflow: Getting Started with Clips and Transitions in Motion Flow Mode

Create your own library of BIP files from imported, and converted, motion capture data, and from your own character animation. With a biped selected, turn on Motion Flow mode. Select Show Graph on the Motion Flow rollout to open the Motion Flow Graph. On the Motion Flow Graph, click Create Clip, and add clips in the Motion Flow dialog. Associate the clips to BIP files by first turning on Select Clips-Transitions on the Motion Flow Graph toolbar and then right-clicking a clip icon; a clip dialog appears, allowing you to browse for a clip.

Click Define Script on the Motion Flow rollout, then select a series of clips on the Motion Flow Graph. By left-clicking on a clip it will appear in

the Scripts list. The clips are connected visually in the Motion Flow Graph with red arrows (active script) representing default transitions (Minimum Motion Loss). The default duration for a transition is 25 frames, which provides good results in many cases. Transitions, by default, use velocity interpolation between clips.

Use Edit Transition in the Scripts section to fine-tune the transitions between clips. Select a clip in the Scripts list, then from the Scripts commands, click Edit Transition; the *Transition Editor (page 2–1051)* displays, and the source and destination clip names are displayed at the top of the Transition Editor.

In the Source Clip area, Start Frame represents the frame in the source file where the transition starts. For example, if Source Clip Start Frame is 60, the transition *from* the source clip starts at frame 60. The *duration* of the transition is set in the Length field at the upper left; Length values are in frames. In this example, if Length is set to 10, the transition to the destination clip takes 10 frames. In the Destination Clip area, Start Frame represents the frame in the *destination* clip that the transition starts; a value of 80, for example, starts the transition at frame 80 of the *destination* clip. In this example, the source clip plays from 0 to 60, there is a 10 frame transition from frame 60 of the source clip to frame 90 of the destination clip (frames 80 to 90 cover the destination clip transition period), then the rest of the destination clip plays.

When the Transition Editor displays, the first things to try, before manually editing the transition, are the Optimized Transitions. Optimized transitions use a minimum foot sliding algorithm. The top right-hand corner of the Transition Editor dialog is the icon for optimized transitions. If the optimized transition is not satisfactory, try editing the transition manually.

Manual transition editing offers the most control; the Frame spinners in the Ghost areas of the Transition Editor allow you to scrub the source and destination clips while viewing two stick figures. Find a good start frame in both clips using the Frame spinners. Things to look for in both clips are similar supporting feet, body momentum that will appear natural, and arm motion similarities. If velocity changes between the clips are too abrupt, use the Length field to adjust the duration of the transition.

On the Motion Flow rollout, click Save File to save your work as an MFE file; transitions and scripts are saved. These BIP files contain no footsteps. Biped foot keys are saved with an IK Blend value of 1 for footsteps. To extract footsteps, exit Motion Flow mode, use Load File on the Biped rollout, and then click Convert. Convert looks at foot IK Blend values of 1 to extract footsteps.

Tip: The location of any referenced BIP file is saved in the MFE file. If the BIP file cannot be found, a Missing Motion Flow Files dialog opens up respectively, in which the missing files are listed. To rectify this, open the *Asset Tracking dialog (page 3–487)* and set a new path of the missing files.

See also

The additional topics that discuss Motion Flow mode in this section are:

Motion Flow Rollout (page 2–1045)

Motion Flow Graph Dialog (page 2–1045)

Motion Flow Scripts Group (page 2–1048)

Transition Editor (page 2–1051)

Create Random Motion Dialog (page 2–1055)

Shared Motion Flow Dialog (page 2–1056)

Transition Optimization Dialog (page 2–1058)

Clip Properties Dialog (page 2–1059)

Motion Flow Rollout

Select a biped. > Motion panel > Biped rollout > Motion Flow Mode > Motion Flow rollout

The Motion Flow rollout displays when Motion Flow mode is turned on in the Biped rollout. Load, append, and save motion flow editor files (*.mfe*), display the *Motion Flow Graph (page 2–1045)* and display the Shared Motion Flow dialog using controls on this rollout.

To load, append, or save a Motion Flow Editor, *.mfe*, file refer to the following procedures:

- *To load Motion Flow Editor files (page 2–1033)*
- *To append Motion Flow Editor files (page 2–1033)*
- *To save Motion Flow Editor files (page 2–1032)*

Interface



Load File—Load a Motion Flow Editor file (*.mfe*). Motion Flow Editor files include:

- **Clips—**References to biped animation files.
- **Transitions—**Names, attributes, and connections between clips.
- **Scripts—**Different paths through a set of connected clips and transitions.

Tip: The location of the referenced *.bip* files is saved in the *.mfe* file. If the *.bip* file cannot be found, the program looks to the Motion Flow directory specified in *plugcfg\biped.ini*. By default, this setting is:

MoFlowDir=<maxdir>\cstudio\moflow

If you load an *.mfe* file onto a biped using a shared motion flow, you will get a warning and the biped will be removed from the shared motion flow. The biped will get the newly loaded motion flow and all its scripts.

Append File— Append a Motion Flow Editor (*.mfe*) file to the *.mfe* that is already loaded. Displays a load file dialog.

The appended graph will appear directly below the bottom of the existing graph in the graph window, so you may have to scroll down to see it.

Save File— Save a Motion Flow Editor (*.mfe*) file.

Saving a *.mfe* file from a biped having a shared motion flow will save the motion flow and its scripts as if it were not shared. It will be a normal *.mfe* file.

Note: To save a script as a *.bip* file, use Unified motion to have the scripted motion available when you exit Motion Flow mode.

Show Graph— Opens the *Motion Flow Graph (page 2–1045)*. The first step in script creation is to add clips to the Motion Flow Graph.

Shared Motion Flow— Displays the *Shared Motion Flows dialog (page 2–1056)*. Allows you to create, delete, and modify shared motion flows. Shared motion flows are used to animate one or more bipeds to simulate a crowd. If the selected biped is using a shared motion flow, then the icon has a white circle around it.

Motion Flow Graph Dialog

Select a biped. > Motion panel > Biped rollout > Motion Flow Mode > Motion Flow rollout > Motion Flow Graph > Motion Flow Graph dialog

Use tools in the Motion Flow Graph to add clips to the graph, calculate optimized transitions, set

random script transition values, move and delete clips, and display clip dependencies. Clips and transitions display as icons in the Motion Flow Graph dialog. The Motion Flow Graph displays when you click Show Graph on the *Motion Flow rollout (page 2–1045)* on the Motion panel.

The first step in Motion Flow mode is to add clips in the Motion Flow Graph for use in scripts. Clips represent all or part of a *.bip* file. Scripts represent different paths through the clips in the Motion Flow Graph. The first clip in the current script is red. Transitions are shown as arrows between clips, red arrows represent the path through the active script. Black transition arrows indicate unloaded scripts. A transition looping back to the same clip represents a cycle or loop.

If the biped is using a shared motion flow, then the title of the graph window will say "*SHARED* Motion Flow Graph", followed by the name of the shared motion flow. Shared Motion Flows are used to control multiple bipeds with one shared motion flow.

By default, minimum motion loss is used to compute transitions. Optimized transitions use an algorithm that uses minimum foot sliding. Optimized transitions take longer to compute but yield very high quality results.

Random Scripts for One or More Bipeds

You can create random scripts by using the Create Random Motion command in the Script group on the Motion Flow rollout. Random scripts are created by randomly traversing clips in a motion flow graph. To use Create Random Motion, each biped must be in the same shared motion flow.

To use a shared motion flow to create random scripts for multiple bipeds, the first step is to click Shared Motion Flow on the Motion Flow rollout, and then add bipeds that will share one shared motion flow. On the Motion Flow Graph, clips are added and transitions are created between all the clips. Then Create Random Motion is used to compute a random motion for all the bipeds.

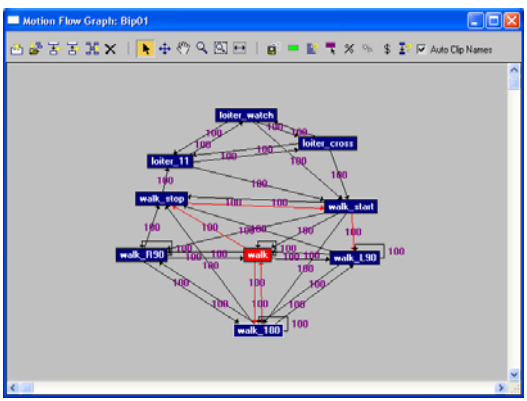
Clips and transitions are given percentages that are used by the software to generate random scripts. Percentages for clips and transitions are set in the clip dialogs and in the Transition Editor dialog.

You can load many clips and use Create All Transitions to create all possible transitions between the clips. You can then use Optimize Selected Transitions to create optimized transitions. Once the transitions are created, you can quickly create scripts with optimized transitions or generate random motions for a crowd of bipeds.

A *.mfe* file stores pointers to the clips, transition parameters, and scripts.

See *Placing Motions on the Motion Flow Graph (page 2–1027)* to add clips to the graph, or *Creating Random Motion (page 2–1035)* to create a random script and transitions.

Interface



Create Clip—Select and click in the dialog window to create clips.

The clips are empty. Right-click a clip using the Select Clip tool to display the *Clip Properties dialog (page 2–1059)*; then select a *.bip* file and

set its duration in the Clip Properties dialog. You can also set the Random Start Probability here. Random Start Probability is used when multiple clips are selected as possible start clips when you generate a random motion flow.

Note: Setting clip duration is not critical for transitions; the *Transition Editor (page 2–1051)* allows you to start and end a transition on any frame of a clip.

 **Create Multiple Clips—**Load multiple motion files.

Displays an open file dialog. Select multiple files and click OK; multiple files are loaded into the Motion Flow Graph window.

Set lowest starting foot height to Z=0 (.bip files only)—Sets the lowest starting foot height to Z=0. This is an option in the Load File dialog. Default=On.

In Biped, the height of a motion clip can be retained. This is important if you want to retain the height of a motion clip for motions adapted to characters of different sizes. If, for example, the character is jumping off a rock and you want to retain the Z position of the character, you would turn this option off. Leave this option off if Motion Flow motions must be blended that begin and end at different heights, such as three clips that have the character mounting a bicycle, riding the bicycle, and dismounting the bicycle.

Turning off this option can, however, cause a jump in the motion during motion flow transitions. Turn this on for smooth transitions in Motion Flow mode. If adaptation takes place, the height is set so that the lowest foot at frame 0 starts at the Z=0 height. This lines up clips along the Z axis and creates smooth transitions.

 **Create Transition From –> To—**Create a transition between two clips.

Click+drag from one clip to another in the Motion Flow Graph dialog. Click on a single clip and then mouse up to create a loop transition. It is necessary to have this capability in order to create random scripts. You can create transitions which are not included in a script.

 **Create Transition To –> From—**Create a transition between two clips.

Click+drag to one clip from another in the Motion Flow Graph dialog. Click on a single clip and then mouse up to create a loop transition. It is necessary to have this capability in order to create random scripts. You can create transitions which are not included in a script.

 **Create All Transitions—**Creates transitions between every clip, including loop transitions. Select all the clips that require transitions, then click Create All Transitions. The transitions are not optimized. Use Optimize Selected Transitions to optimize the transitions. Optimized transitions take time to compute but are high quality.

 **Delete Clip/Transition—**Deletes a clip or transition.

If a script is dependent on the clip, a dialog displays a warning; clicking OK on the dialog deletes the clip and the script that is dependent on it.

If you delete the selected clips and transitions from a shared motion flow, it will delete all the scripts from all the bipeds sharing that motion flow which are dependent on those clips and transitions.

 **Select Clip/Transition—**Selects a motion clip or transition.

Right-click a clip to display the *Clip Properties dialog (page 2–1059)*. Right-click a transition to display the *Transition Editor dialog (page 2–1051)*.


Move Clip—Moves clips within the Motion Flow Graph. This does not affect the animation.


Pan—Pans the layout of the clips.


Zoom—Adjusts the view magnification of the Motion Flow window. Drag up to increase magnification. Drag down to decrease magnification.


Zoom Region—Click Zoom Region mode to drag a rectangular region and magnify that region to fill the Motion Flow Graph window.


Fit to Window—Re-sizes the contents to fit the size of the Motion Flow Graph window.


Save Clip Files—Lets you set a path where selected clip files can be stored.


Clip Mode—Edit biped footsteps and limbs for the selected clip. Use Set Key on the Keyframing rollout to set biped limb keys. Clips turn green in the Motion Flow Graph window in Clip Mode.


Show Script Dependencies—Displays the scripts that use the selected clip.

If you push the Show Script Dependencies button on a shared motion flow graph, it will check all the bipeds sharing that motion flow for scripts dependent on the selected clips and transitions.


Select Random Start Clips—Turn on and select clips in the Motion Flow Graph window.

Press Ctrl+click to add clips. Selected clips are used by Create Random Motion in the Scripts group to start on one of the selected clips based on percentage. If three clips are selected using the default weighting of 100, then each clip has an equal chance of being the start clip.


Show Random Percentages—Displays clip and transition percentages in the Motion Flow Graph window.

Random start clips display in purple and display their probability of starting a random script. This also shows the probability (0 - 100) that each transition will be chosen. Create Random Motion in the Scripts group uses clip and transition percentages to generate random scripts.


Optimize Selected Transitions—Select one or more transitions and then click Optimize Selected Transition to optimize them. Displays the *Transition Optimization dialog (page 2–1058)* to set the location of the transition.

Optimized transitions take time to compute. A progress bar is displayed when you use this feature. Minimum foot sliding is the method used to compute an optimized transition.


Show Optimal Transition Costs—Displays costs in the Motion Flow Graph window. The lower the number the better the transition.


Check All Transitions—Checks the graph for overlapping transitions and transitions whose length is too long for the clip. It informs you of any problems, or tells you that none have been found.



Auto Clip Names—Names the clip based on the name of the motion file.

Turn off to name a clip yourself.

Motion Flow Scripts Group

Select a biped. > Motion panel > Biped rollout > Motion Flow Mode > Motion Flow rollout > Scripts group

Create and delete scripts, name scripts, edit transitions, edit clips, and position the entire

animation using tools in the Scripts group. Create random motion for one or more bipeds using controls in this rollout. The Scripts group on the Motion Flow rollout is only available when *Motion Flow mode (page 2–1043)* is active.

Scripts

A Script is a list of clips (*.bip* files) that you constructed and are executed as you designed to animate a character. To create a script, add clips to the Motion Flow Graph, then select Define Script in the Scripts group and click a sequence of clips from the Motion Flow Graph window. Default transitions are assigned if no transitions exist between the clips. The clip names and starting frame numbers display in the list in the Scripts group.

Different scripts vary the order of the clips found in the Motion Flow Graph. Scripts are run in a top-down order to animate the character. Click Play to view script motion on the biped.

Tip: As the length of a script varies, use **Alt+R** to set the Active Time in 3ds Max to the length of the script. Turn on the Plug-in Keyboard Shortcut Toggle near the 3ds Max status line to enable **character studio** shortcuts.

Transitions Between Clips

By default, when a script is created, Minimum Motion Loss is used to find likely start frames for the source and destination clips.

Editing transitions using the *Transition Editor (page 2–1051)* allows you to determine where a transition occurs in the source and destination clip. Transition duration and the orientation of the destination clip can also be adjusted in the Transition Editor.

Random Motion

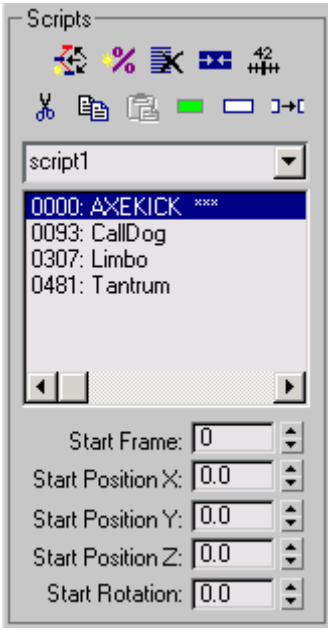
The Create Random Motion command traverses clips in the Motion Flow Graph based on transition percentages. Transitions are given percentages, and the software creates random scripts based on the transition percentages. This is a quick way to create crowd scenes or to try out different scripts on a single biped. For example, if five clips are present in the Motion Flow Graph and transitions exist between all the clips and each clip has a percentage or probability of being used, you can use Create Random Motion to create a script that is comprised of the five clips that are selected at random.

Position the Entire Animation

Use the Position, Rotation, and Start Frame controls to position the entire animation. If you are editing the script for a character in a scene with other objects or characters, use these controls to position the animation relative to the rest of the scene.

For more information regarding scripts, refer to *Creating a Motion Flow Script (page 2–1030)*.

Interface



Define Script—Displays the Biped Motion Flow Script dialog (no dialog displays if there are no scripts; in this case, simply select clips in the Motion Flow Graph).

The Biped Motion Flow Script dialog has the following options:

- **Create New Script**—Names a new script. Select clips in Motion Flow Graph to create the clip list for the new script.
- **Redefine Script**—Keeps the script name and removes clips in the list. Select clips in Motion Flow Graph to create a new clip list.
- **Insert Above Selected Clip item**—Inserts a clip above the selected clip in the list.

  First select a clip in the list, choose Insert Above Selected Clip item, and then click a clip in the Motion Flow Graph.
- **Insert Below Selected Clip item**—Inserts a clip below the selected clip in the list.
- **Append to End of Script**—Appends a clip to the end of the clip list.

Create Random Motion—Displays the *Create Random Motion (page 2–1055)* dialog.

Controls in the Create Random Motion dialog allow you to create random scripts to animate one or more bipeds. Random motion on multiple bipeds can be used to create a crowd scene.

Delete Script—Deletes the current script; displays the previous script if one is present.

Create Unified Motion—Converts a script into a Freeform unified motion. The created motion replaces animation present when Motion Flow Mode is turned off.

Go to Frame—Make the first frame of the selected clip the current frame.

Cut—Removes the selected clip from the script list and creates a default transition to the next clip on the list.

Copy—Copy the selected clip to the clipboard.

Paste—Paste a clip from the clipboard.

Clip Mode—Edit biped footsteps and limbs for the selected clip. Use Set Key on the Keyframing rollout to set biped limb keys.

Edit Clip—Displays the *Clip Properties dialog (page 2–1059)*. Change the start and end frame for the current clip, or replace the current clip with another one. You can also set the Random Start Probability here. Random Start Probability is used when multiple clips are selected as possible start clips when you generate a random motion flow.

The clip name in the script list and the icon in the graph window are updated if the clip is replaced.

 **Edit Transition—**Displays the *Transition Editor (page 2–1051)* for the selected clip.

Edit the transition for the selected clip and the clip following it. By default, when a script is created, Minimum Motion Loss is used to find start frames for the source and destination clips. Use Edit Transition to select your own start frames or to try out optimized transitions.

Note: Right-clicking a transition "arrow" in the Motion Flow Graph also displays the Transition Editor, but it will only give you the basic transition editing tools. To edit the transition's Source Clip-Frame Start and the Destination Clip, use the Edit Transition button.

Move the Entire Animation

Start Frame—Set the start frame for the first clip in the script.

Start Position X—Move the entire script along a world X-axis.

Start Position Y—Move the entire script along a world Y-axis.

Start Position Z—Move the entire script along a world Z-axis.

Start Rotation—Rotate the entire script around the world Z-axis.

All transformation and rotation is based on the original position and affect the entire script.

Transition Editor

Select a biped. > Motion panel > Biped rollout > Motion Flow Mode > Scripts group > Select a clip in the script list. > Transition Editor

Display the Transition Editor by selecting a clip in the list of the Scripts group and clicking Edit Transition in the same group or right-clicking a transition arrow in the Motion Flow Graph.

Transitions

A good transition links two clips together seamlessly; the motion through the transition should appear natural, as though the motion was captured as one long motion sequence. Like an AB roll transition in video editing, an appropriate section in both clips is selected for the transition (dissolve) from the source clip to the destination clip. Velocity differences between the source and destination clips are matched during the period of transition producing a seamless result. By default, Minimum Motion Loss is used to find likely start frames in the source and destination clips when clips are appended to a script. Optimized transitions can be computed by using Optimize Transition in the upper right-hand corner of the Transition Editor dialog. Optimized transitions use a minimum foot sliding method to compute transitions. Optimized transitions take longer to compute, but produces very smooth motions.

Although only one arrow is used to represent a transition between two clips in the Motion Flow Graph, any number of transitions can be named and stored in the Transition Editor representing that transition. If, for example, you create 5 different transitions between two clips for one script, all of these transitions are available in a new script that uses the same two clips. Think of Motion Flow Graph as a data storage area; if all of the scripts are deleted, the transitions are preserved and can be stored in a *.mfe* file.

Automatic Transitions

When you create a script, default transitions are used between the clips. Default transitions use minimum motion loss and are quick to compute. However, the best quality transitions are the optimized transitions. Once the Transition Editor is open, the first thing to try, before manual editing, are the optimized transitions (upper-right corner of the dialog).

Length (Transition Duration)

Set the duration of a transition in the Length field. A value of 10, for example, creates a transition of 10 frames between the source and destination clips. During the period of transition, the velocity of the source clip is interpolated to the velocity of the destination clip. If the transition takes place at the last frame of the source clip and the first frame of the destination clip, and Length is set to 10, then the last frame of the source clip is interpolated with the first 10 frames of the destination clip.

Editing Transitions Manually (Ghosts)

Manually setting the Start Frame for the source and destination clips offers the most control. Unwanted motion in either clip can be avoided and judging the best Start Frames for both clips is left to you.

The Ghost area Frame spinners allow you to view and scrub the source and destination clips by displaying stick figures (ghosts); yellow and red stick figures represent the source and destination clips. When a suitable Start Frame is located for both clips, use Set Start Frame in the Ghost area to copy the Frame values to the Start Frame fields in the Source and Destination Clip areas.

Scrubbing the time slider over the transition period enables you to view the biped's transition from the yellow stick figure (source) to the red stick figure (destination).

Other Transition Editor Features

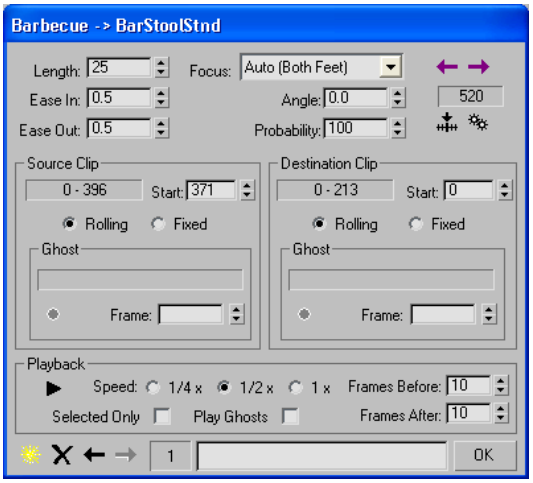
Rolling and Fixed specify whether a clip is rolling (in motion) or fixed (single frame) during the transition. Change the direction of the destination clip using the Angle field.

Other parameters in the Transition Editor allow you to create and name new transitions, scroll through the saved transitions, jump to the transition-starting frame, set automatic transition parameters, and go to the next transition in the script

All transitions and their attributes are saved with a *.mfe* file.

See *Customizing Transitions (page 2–1034)* to manually and automatically customize transitions.

Interface



Probability—Set a probability value for random transitions. This is used by Create Random Motion when a random script is generated.

Length—Sets the number of frames for the duration of the transition.

Transitions are calculated by matching velocities in both clips. Smooth out abrupt velocity changes using longer transitions.

Ease In—Ease in value for the source clip.

Ease Out—Ease out value for the destination clip.

Transition Focus—Lets you specify a focus point on the biped where the transition takes place. Default = Center of Mass.

- **Center Of Mass—**The transition focus is based on the center of mass position of the biped as it transitions from one clip to the next.
- **Left Foot—**The transition focus is based on the left foot position of the biped as it transitions from one clip to the next.
- **Right Foot—**The transition focus is based on the right foot position of the biped as it transitions from one clip to the next.
- **Both Feet—**The transition focus is based on an averaged foot position of both of the biped's feet as it transitions from one clip to the next.

Note: The best way to see the differences between the transition foci is by watching the yellow and red ghosts.

Angle—Sets the direction of the destination clip.

The angle of the destination clip is automatically set for best body fit between the two clips when the Start Frame values change. Use Angle to change the direction of the destination clip.



Previous Transition—Go to the previous transition in the script.

Displays the previous transition in the Transition Editor, moves the time slider to the start frame of the previous transition, and highlights the previous clip in the Scripts list.

Next Transition—Go to the next transition in the script.

Displays the next transition in the Transition Editor, moves the time slider to the start frame of the next transition, and highlights the next clip in the Scripts list.

Optimize Transition—Displays the *Transition Optimization (page 2–1058)* dialog.

Options in the Transition Optimization dialog allow you to search for a location for the transition.

Go To Start Frame—Moves the time slider to the first frame of the transition.

The number field shows the start frame number.

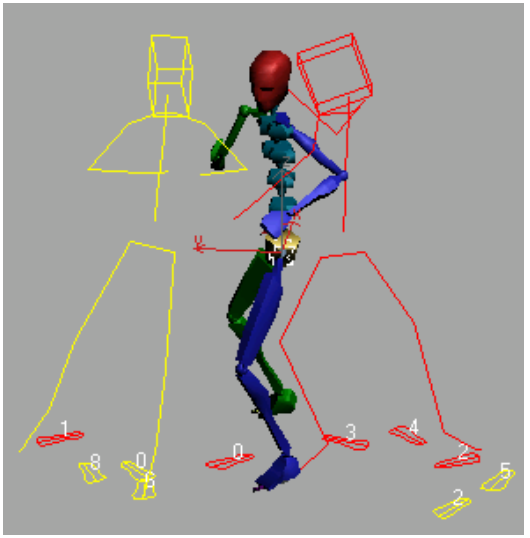


Start Frame—Set the transition start frame for the source and destination clips in their respective fields. Duration for the source and destination clips display above the Start Frame fields.

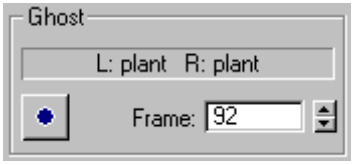
Rolling—Keep the clip in motion during the transition.

Fixed—Freeze the biped at the Start Frame position during the transition.

Ghost group



The Ghost group Frame spinners allow you to view and scrub the source and destination clips by displaying stick figures (ghosts); yellow and red stick figures represent the source and destination clips. The source and destination bipeds might not be near each other during this scrubbing process; the destination clip will be repositioned when you click Set Start Frame is clicked. When you locate a suitable start frame, click Set Start Frame to copy the values in the Frame field to the Start Frame field. Monitor foot position status in the field provided.



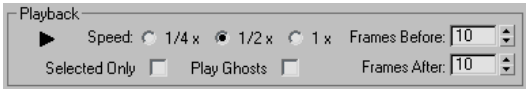
Set Start Frame—Copy the value in the Frame field of the Ghost area to the Start Frame field in the Clip area. The position of the destination clip changes to match the biped body in the destination clip to the biped body in the source clip.

Locate an appropriate start frame for the source and destination clips by using the Frame spinner and viewing the positions of both stick figures, then click Set Start Frame.

The destination clip is rotated and positioned to match both bipeds. Use the Angle spinner to reorient the destination clip.

Frame—Use the Frame spinner to scrub a stick figure back and forth, which allows you to determine a start frame for the source and destination clips. Visual feedback of the stick figures is a good way to judge which start frames are needed for the source and destination clips.

Playback group



These controls let you play back the transition from the Transition Editor dialog.

- **Play Transition Icon**— Toggles the playback of transition.
- **Speed**— Changes the playback speed.
- **Frames Before/Frames After**— Sets the number of frames to play before and after the transition period.
- **Selected Only**—When checked, plays only the selected biped in the scene.
- **Play Ghosts**— When checked, shows transition ghosts during playback.



Create Transition—Click to create a *new* transition. The transition number field increments. Edit and name the new transition. Any number of transitions can be stored.

Note: Clicking OK saves the displayed transition. Create Transition is used only if you want to work on a new transition.

Delete Transition—Click to delete a transition.

Previous Transition—Go to the previous stored transition. This button is grayed if no previous transition exists.

Next Transition—Go to the next stored transition. This button is grayed if no next transition exists.

OK—Store transitions and exit the dialog.

Create Random Motion Dialog

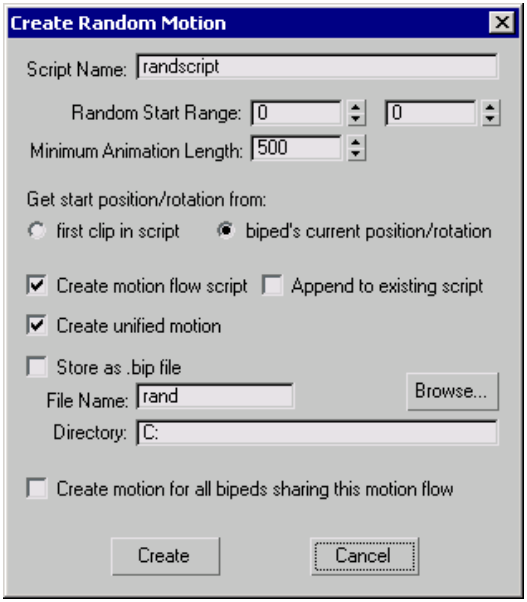
Select a biped. > Motion panel > Motion Flow Mode > Motion Flow rollout > Scripts group > Create Random Motion > Create Random Motion dialog

You can randomly traverse clips in a motion flow graph to animate one or more bipeds using controls in the Create Random Motion dialog. Parameters for random motion are set in the Motion Flow Graph, in the clip and transition dialogs, as well as in the Create Random Motion dialog.

Random motion is created by first adding clips to the Motion Flow Graph window and adding transitions between the clips. Clips and transitions are then given percentages, which are used to create random motion for one or more bipeds. You can manually control the "weighting" for possible start clip, transitions, and frame range. This allows you to animate multiple bipeds in a crowd scene, for example. A different script is created for each biped.

If you want to create random motion for multiple bipeds, they must be sharing a motion flow. For more information about using random motions and transitions, refer to *Creating Random Motion (page 2–1035)*.

Interface



Script Name—Type a name for the script to be generated.

Random Start Range—Set the start and end frame range over which the new script(s) will start.

Minimum Animation Length—Set the minimum animation length.

When a random motion is created, it is done by making a motion flow script which traverses the clips in the Motion Flow Graph, adding clips based on random calculations. It will add clips until the length of the script is greater than or equal to the minimum animation length, specified here.

first clip in script—Gets the start position and rotation from the first clip.

biped's current position/rotation—Uses the biped's current position to start the script.

Create motion flow script—Creates a script after computing the motion.

If this is on, a motion flow script will be created. If not, a script will still be created internally in

order to generate the random motion, but it will be deleted after the random motion is generated.

Append to existing script—Appends random motion to the existing script.

Create unified motion—Creates a unified motion. The generated motion will be available when you exit Motion Flow mode.

If multiple bipeds are in the random calculation, then the motion is unified for each biped.

Store as .bip file—Stores the random motion as a *.bip* file.

If multiple bipeds are in the random calculation, they are saved separately with incrementing numbers.

File Name—Type a name for the *.bip* file.

A *.bip* extension will be added automatically.

Directory—Type a directory path or browse for the path.

Browse—Browse to a directory.

Create motion for all bipeds sharing this motion flow—Turn on to create a random script for each biped sharing the current biped's motion flow.

You can create a shared motion flow by using the Shared Motion Flow command on the Motion Flow rollout and adding bipeds in the Shared Motion Flow dialog.

Create—Creates random motion for the selected biped or all the bipeds in the shared motion flow.

Cancel—Cancel and close the dialog.

Shared Motion Flow Dialog

Select a biped. > Motion panel > Biped rollout > Motion Flow Mode > Motion Flow rollout > Shared Motion Flow > Shared Motion Flow Dialog

Controls in the Shared Motion Flow dialog allow you to assign one motion flow to multiple bipeds. Instead of choreographing a motion flow script for individual bipeds, you can create a motion flow with all the clips and transitions to create a script to animate multiple bipeds. Random motion creation will use each bipeds' own motion flow. If a biped's motion flow happens to be a shared motion flow, then the shared motion flow will be used to compute random motion.

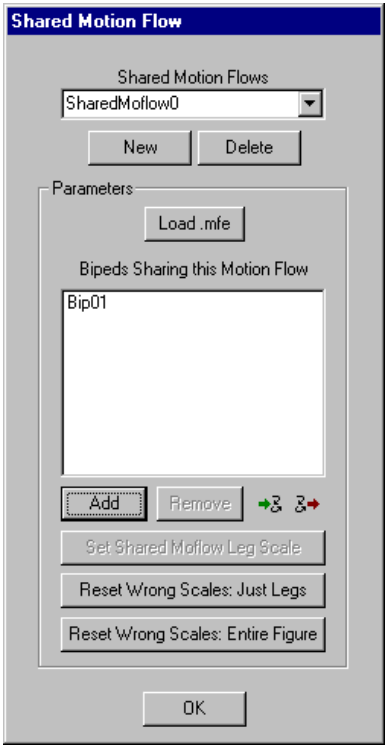
A biped that shares a motion flow shares only the graph. Its scripts are unique to that biped, although the scripts point to the clips of the shared motion flow. You can manipulate that biped's motion flow and scripts. You can create random motion on a biped that shares a motion flow, or create a motion flow script via the crowd system.

A white circle around the Shared Motion Flow icon indicates that a biped is using a shared motion flow. If you edit that biped's motion flow graph, the title of the graph dialog will say "Shared Motion Flow Graph", followed by the name of the shared motion flow.

Bipeds in a shared motion flow should have the same lower body scale and structure. Adaptation for differently sized bipeds does not occur in a shared motion flow. If you want differently sized bipeds in a crowd, then create a shared motion flow for *each* size.

See *Sharing Motion Flow (page 2–1039)* for more details about setting up a shared motion flow.

Interface



Shared Motion Flows list—Lists shared motion flows.

New—Creates a new shared motion flow.

Delete—Deletes the current shared motion flow.

The scripts of the bipeds sharing the deleted motion flow will be deleted. Those bipeds will have an empty motion flow graph and no scripts.

Load—Loads a shared motion flow (*.smf*) file.

This file is created with the Save command (see following). The file contains the shared motion flow graph as well as the scripts for all bipeds sharing the motion flow.

Save—Saves a shared motion flow (*.smf*) file.

The file contains the shared motion flow graph as well as the scripts for all bipeds sharing the motion flow.

Use this facility to speed biped/crowd-simulation setup. Create and refine the crowd simulation with bipeds that don't use Physique, save the shared motion flow, load or create your final crowd members that use Physique, and then load the shared motion flow. For a procedure, see *To apply a shared motion flow to a different group of bipeds: (page 2–1041)*.

Load .mfe—Displays a load file dialog. Load an *.mfe* file into the shared motion flow.

Note: If you load a *.mfe* file, via the Motion Flow rollout, into a biped using a shared motion flow, you will get a warning and the biped will be removed from the shared motion flow. The biped will get the newly loaded motion flow and all its scripts. The shared motion flow will remain the same.

Bipeds Sharing this Motion Flow list—Lists the bipeds that share this motion flow.

Add—Displays a dialog where you can choose bipeds to add to the motion flow list.

Add bipeds to the list of bipeds that share a motion flow. When you add a biped, its current motion flow graph and motion flow scripts will be deleted. It will now have the shared motion flow graph. A biped can only share one motion flow graph. You must remove a biped from its shared motion flow in order to add it to a different shared motion flow.

Remove—Removes the selected bipeds in the list from the current shared motion flow.

These bipeds' scripts will be deleted. They will have an empty motion flow graph.

Put Multiple Bipeds in Motion Flow—Put the bipeds in the list into Motion Flow mode.

The crowd system needs Motion Flow mode to be turned on to perform calculations for motion. This is a convenient way of turning on Motion Flow mode for multiple bipeds.

Take Multiple Bipeds out of Motion Flow—Take the bipeds in the list out of Motion Flow mode.

Set Shared Moflow Leg Scale—Adapts the shared motion flow to the scale of the biped currently selected in the list. After this operation, the selected biped will have the correct leg scale, although other bipeds may not.

Reset Wrong Scales—Just Legs: Reset the leg scale only of the bipeds that have the wrong scale, so that they adapt appropriately to the shared motion flow.

Reset Wrong Scales—Entire Figure: Resets the entire figure structure of the bipeds that have the wrong scale, to match the figure structure of the correctly scaled biped.

When you add the first biped to the shared motion flow, the system adapts to the size of that biped. If the leg scale of a biped you add subsequently doesn't match that of the first biped, it will be marked in the list with "wrong scale" after it.

All the bipeds you plan on using in a shared motion flow must have the same lower body structure and scale.

Transition Optimization Dialog

Motion Flow Mode > Motion Flow Graph > Select transitions in the graph window. > Optimize Selected Transitions > Transition Optimization dialog

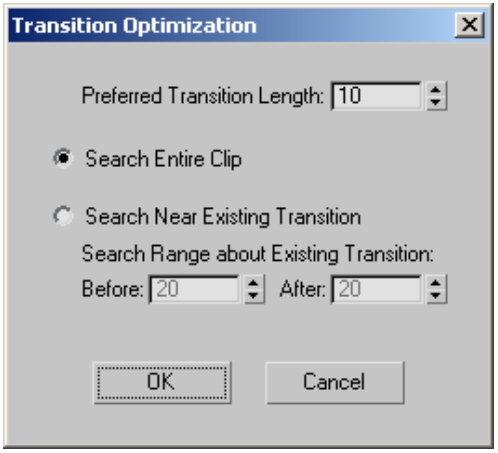
Motion Flow Mode > Motion Flow rollout > Scripts group > Select a clip in the Script list. > Edit Transition > Optimize Transition > Transition Optimization dialog

Options in the Transition Optimization dialog allow you to select the range over which the optimize algorithm will search for the transition. It can search either the whole clip, or it can search near the existing transition. You must specify the preferred length of the optimized transition. The software will try to get as close to that length as possible, still opting to give you the best length.

If it searches about the existing transition, you must specify the number of frames about which it will search before and after the existing transition.

Optimized transitions compute for minimum foot sliding over the range of the transition. This method yields very high quality results.

Interface



Preferred Transition Length—Specify the length of the optimized transition.

Search Entire Clip—Search the entire clip for an optimized transition start frame.

Search Near Existing Transition—Create an optimized transition near the existing transition. You may search the range about the existing transitions by setting *before* and *after* frame values.

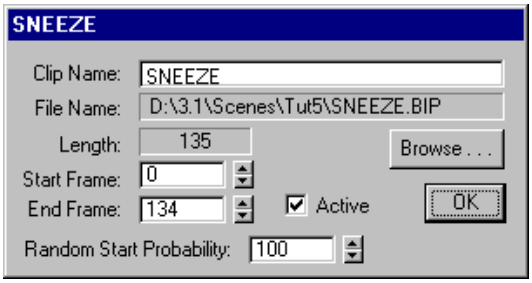
Clip Properties Dialog

Motion Flow Mode > Motion Flow Graph > Right-click a clip in the Motion Flow Graph window. > Clip Properties dialog

Motion Flow Mode > Motion Flow script rollout > Select a clip in the list. > Edit Clip > Clip Properties dialog

Parameters in the Clip Properties dialog allow you to browse for a clip, set a start and end frame for the clip and set a random start probability for the clip. Random start probability is used when you use Create Random Motion to generate a random script for a biped.

Interface



Clip Name—The clip name as it appears in the Motion Flow Graph.

File Name—The path and file name of the motion clip.

Length—The length of the clip in frames.

Browse—Displays a load file dialog. Browse for a motion file.

Set lowest starting foot height to Z=0 (.bip files only)—Sets the lowest starting foot height to Z=0. This is an option in the Load File dialog. Default=On.

In **character studio** the height of a motion clip can be retained. This is important if you want to retain the height of a motion clip for motions adapted to characters of different sizes. If, for example, the character is jumping off a rock and you want to retain the Z position of the character, you would turn this option off. Leave this option off if Motion Flow motions must be blended that begin and end at different heights, such as three clips that have the character mounting a bicycle, riding the bicycle, and dismounting the bicycle.

Turning off this option can, however, cause a jump in the motion during motion flow transitions. Turn this on for smooth transitions in Motion Flow mode. If adaptation takes place, the height is set so that the lowest foot at frame 0 starts at the Z=0 height. This lines up clips along the Z axis and creates smooth transitions.

Start Frame—Sets the start frame for the clip.

End Frame—Sets the end frame for the clip.

Active—Activates the clip. Inactive clips display as a green color in the Motion Flow Graph.

Random Start Probability—Set a percentage for random start probability. This is used when multiple clips are selected as possible starting clips in a random motion flow. The Create Random Motion command allows you to generate random motion for one or more bipeds.

Using Motion Capture

Motion capture is the practice of getting motion data from live actors performing various actions. The motion data is captured (retrieved) via sensors placed at the actors' joints and extremities.

3ds Max does not perform motion capture, but it accepts motion-capture data in the most commonly used formats. This data can be imported to the biped and used as is, or combined with other motions with *Motion Flow (page 2–1026)* or the *Motion Mixer (page 2–604)*.

Motion-capture data typically needs some adjustment before it fits your biped or animation perfectly. When motion-capture data is imported to 3ds Max, it can be filtered to:

- Use fewer keyframes
- Create footstep motion
- Use props in the scene

In addition, some motion-capture files come with a separate marker file, which can be used to match the biped posture to the motion-capture actor.

How Motion-Capture Data is Acquired

Motion-capture data is typically acquired by one of several means:

- Optical sensing technology

  Optical systems have become quite popular over the last couple of years. These systems can offer the performer the most freedom of movement since they do not require any cabling. Optical systems incorporate directionally reflective balls, referred to as *markers (page 3–969)*, that attach to the performer. Optical systems require at least three video cameras, each of which is equipped with a light source that is aligned to illuminate the field of view for that camera. Each camera is in turn connected to a synchronized frame buffer. The computer is presented with each camera view in order to calculate a 3D position of each marker; the resulting data stream therefore consists of 3D position data for each marker. This data is typically applied to an inverse kinematics system, to animate a skeleton.

- Electro-magnetic sensing technology

  This is a popular method used for performance capture. Magnetic capture involves the use of a centrally located transmitter, and a set of receivers that are strapped on to various parts of the performer's body. These receivers are capable of measuring their spatial relationship to the transmitter. Each receiver is connected to an interface that can be synchronized to prevent data skew. The resulting data stream consists of 3D positions and orientations for each receiver. This data is typically applied to an inverse kinematics system to drive an animated skeleton. This magnetic approach shares the same lack of occlusion problems with the audio method, but it also shares the same negative factors, such as the hindrance of cables, lack of sufficient receivers, and the limited capture area. In addition, being magnetic, the system is affected by any sizable areas of metal in the vicinity of the capture area, such as girders, posts, and so on

- Prosthetic sensing technology

  This is one of the early methods for capturing the motion from various parts of human anatomy. These methods include simple "on/off" type of motion-detection systems as well as complex motion-tracking systems. The latter type of prosthetic motion capture would be an ideal approach if it weren't for the complex mechanical requirements and the performance-inhibiting qualities generally associated with such designs. However, the type of data provided can be clean, rotational data collected in real time without any occlusion problems. This method is based on a set of armatures that must be attached all over the performer's body. The armatures are then connected to each other by using a series of rotational and linear encoders. These encoders are then connected to an interface that can simultaneously read all the encoders in order to prevent data skewing. Finally, through a set of trigonometry functions, the performer's motion can be analyzed. These design restrictions seem to be quite difficult to overcome, and will probably limit the use of this type of device for character animation.

- Acoustic sensing technology

Acoustic capture is another method currently used for performance capture. This method involves the use of a triad of audio receivers. An array of audio transmitters are strapped to various parts of the performer's body. The transmitters are sequentially triggered to output a "click" and each receiver measures the time it takes for the sound to travel from each transmitter. The calculated distance of the three receivers is triangulated to provide a point in 3D space. An inherent issue with this approach is the sequential nature of the position data it creates. In general, one would like to see a "snap shot" of the performer's skeletal position rather than a time-skewed data stream. This position data is typically applied to an inverse kinematics system, which in turn drives an animated skeleton.

See also

Importing Motion-Capture Data (page 2–1061)

Filtering Motion-Capture and Marker Data (page 2–1061)

Sliding Footsteps (page 2–1064)

Prop Bone (page 2–1065)

Importing Motion-Capture Data

In **character studio**, you can import both rotation and position type motion-capture files.

- BioVision (*.bvh*) files contain limb and joint rotation data.
- **character studio** marker files (*.csm*) contain raw marker position data generated by a motion-capture device: markers are attached to an actor during a motion-capture performance.

  Marker files typically require some calibration. If necessary, the biped is sized to fit the markers first, then the biped limbs are oriented to align them to the markers. Marker files should be loaded with keys at every frame and no footstep extraction; this is required for calibration, and also enables the calibration controls. Calibration controls are the second row of buttons on the Motion Capture rollout.

  In **character studio**, the *.csm* marker file format supports a prop bone in either or both hands. See *Prop Bone (page 2–1065)*.

Key reduction, track selection, footstep extraction, and clip looping are some of the options available using the Motion Capture import filter. Key reduction makes it easy to manipulate the biped and personalize the imported motion-capture data. Extracting footsteps from motion-capture data prevents inappropriate sliding feet, which are a common problem with motion-capture data.

Typically you do not use an entire motion-capture clip as is, unless you capture motions at a studio for your own production. You should become familiar with the body motion in the files you have, then use Motion Flow mode to cut portions of these files together to create animation. For example, take a stretch motion in one clip, and combine it with the walking motion in another clip. The edited script can then be saved as a *.bip* file using the Save Segment command on the Biped rollout. Load this *.bip* file for standard motion editing. This provides a good starting place for you to edit the result to your liking.

All motion-capture controls are on the *Motion Capture rollout (page 2–1065)*.

Filtering Motion-Capture and Marker Data

Motion-capture and marker data typically have keys at every frame. Filtering motion-capture data reduces keys, simplifying the job of altering or personalizing the motion data. Biped lets you filter the data of each track with its own filtering settings, so you have control over which nuances

of motion you want to pick up without filling the rest of the tracks with unwanted keys. Filtering is done using the Motion Capture Conversion Parameters rollout.

Other filtering options include footstep filtering and extraction, looping the data, and importing a portion of the motion-capture file.

3ds Max ships with a variety of raw (unfiltered) motion-capture data files, in *.bip*, *.csm*, and *.bvh* formats. Some of the same data is available in filtered versions, either with footsteps or freeform. Try your own filtering adjustments on the raw versions of this data. Importing the raw data displays the original motion very accurately when you select Show Buffer on the Motion Capture rollout. Use the Motion Capture buffer as a guide when adjusting and refining the filtered data. Several tools are available in the Motion Capture rollout to aid you in this process.

Create your own library of imported and optimized motion capture data by saving *.bip* files for use with other characters, or as part of a longer script in Motion Flow mode. Use a biped that has no mesh attached with Physique.

Tip: Overall, you import the data, adjust it to your liking, and save it as a *.bip* file. You can also run standard *.bip* files through this filtering process to create loops or to extract footsteps from a freeform animation.

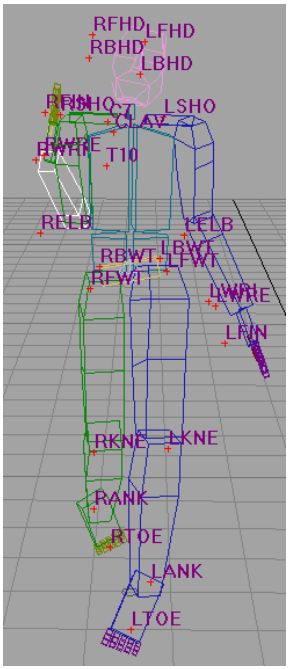
Note: Marker files, such as *.csm*, contain position data. Hierarchical motion-capture files, such as *.bvh*, contain joint rotation data.



Motion Capture rollout

Character Studio Marker Files

The *.csm format (page 3–925)* is an ASCII file used to import positional marker data from motion-capture systems onto a biped.



The Show Markers command displays marker positions and names.

Procedures

To import a motion capture file:

1. Select a biped in the viewports.
2.  On the Motion Capture rollout, click Load Motion Capture File.
3. Choose the file type: *.bvh*, *.bip*, or *.csm*. Search for files in the *cstudio\motions\mocap* directories.

   Tip: *.csm* marker files, loaded for the first time, should be imported with no key reduction and no footstep extraction. This enables the calibration buttons. Marker files typically need some calibration.

4. Select a file and click Open.

   The Motion Capture Conversion Parameters dialog displays (see *Motion Capture Conversion Parameters Dialog (page 2–1070)*).

5. Select the filter options you want and click OK.

   The biped adapts itself to the motion data. If Footstep Extraction is turned on, footsteps appear.

   Tip: Use a biped that does not have a mesh attached with Physique. Import motion-capture data with the idea of then saving a *.bip* file that can be used for any character. If skeletal scale information is loaded from a motion-capture file, a mesh with a Physique modifier might deform unnaturally.

To import a marker file:

Typically, when a marker file is loaded for the first time, it requires scale and position calibration. A raw marker file must be loaded, with no key reduction or footstep extraction, to enable the calibration functions. After calibration is performed, use Convert From Buffer to extract footsteps and reduce keys.

1. Select a biped.
2.  On the Motion Capture rollout, click Load Marker Name File to load a marker name file (*.mnm*).

   This step is not required if the marker names in the marker file adhere to the **character studio** marker naming convention.

3. On the Marker Name File dialog, click Load CSM Marker File, and choose the *.mnm* file from the file open dialog that appears.
4.  On the Motion Capture rollout, click Load Motion Capture File and choose a *.csm* marker file.

   The Motion Capture Conversion Parameters dialog displays (see *Motion Capture Conversion Parameters Dialog (page 2–1070)*).

5. Adjust the filter parameters and click OK.

   Note: Load raw marker data (No Key Reduction, Freeform) to enable the marker calibration buttons.

   The biped adapts itself to the marker data.

6.  On the Motion Capture rollout, click Show Markers. On the Marker Display dialog, turn on Show Recognized Markers and the option for On All Objects.
7.  Now, on the Motion Capture rollout, click Talent Figure Mode. Use Non-Uniform Scale or Rubber Band Mode (on the Biped rollout) to size the biped to the displayed markers.

   Note: This step is optional and should be used if you need to correct for slight differences in limb scale between the original talent who performed the motion and the scale of the biped after the data is imported. For example, if the leg is too short, scale the length of the leg in Talent Figure mode to adjust the knee position.

8.  Click Talent Figure Mode again to exit the mode.
9. Key adaptation takes place when you exit Talent Figure mode. Now biped limb positions relative to the markers can be adjusted.
10. Align the biped limbs to the markers if necessary, then click Adjust Talent Pose to compute the offset for the entire animation.
11. Use Save Talent Figure Structure and Save Talent Pose Adjustment as a *.fig* and *.cal* file.
12. Load these files in the Motion Capture Conversion Parameters dialog when similar marker files are imported in the future.

 At this point, you can use Convert From Buffer to extract footsteps and reduce keyframes. Both scale and position adjustments will be incorporated. Save the motion as an optimized *.bip* file.

Sliding Footsteps

Motion-capture and marker data typically have keys at every frame. Filtering motion-capture data reduces keys, making the job of altering or personalizing the motion data much simpler. You can create your own library of imported and optimized motion-capture data by saving *.bip* files for use with other characters or as part of a longer script in Motion Flow mode. Use a biped that has no mesh attached to it with Physique. You import the data, adjust it to your liking, and save it as a *.bip* file. You can also run standard *.bip* files through this filtering process to create loops or to extract footsteps from a freeform animation.

Motion Capture Import

Footstep motion capture is enabled automatically when importing motion capture data using the Sliding Tolerance and Sliding Angle controls in the *Motion Capture Conversion Parameters dialog (page 2–1070)*.

Footstep Extraction group

The following options are active when Footstep Extraction is on during motion capture conversion.

Extraction Tolerance—Sets the sensitivity of footstep extraction. **character studio** determines if the footstep is there by checking that the foot does not move beyond the distance determined by the Extraction Tolerance value. Smaller numbers are more sensitive and extract more footsteps. The value is a percentage of foot length. The default value is 0.15. Increase this value to 0.2 or 0.25 if too many footsteps are generated.

Sliding Distance—Creates a *sliding footstep (page 2–1064)* when positional tolerance is reached. This value is a percentage of foot length. By default the foot must slide its own distance (100), before a sliding footstep is created.

Use this with motion-capture files that contain sliding feet. A sliding footstep can be created manually by setting IK Blend > 0 for a biped foot at a "touch" state key (biped foot first touches a footstep).

Note: Sliding footsteps display as a footstep with a line through the center.

Sliding Angle—Creates a *sliding footstep (page 2–1064)* when rotational tolerance is reached. This value is in degrees. The default is set high (360 degrees), the foot must make a complete turn before a sliding footstep is created.

Use this with motion-capture files that contain feet that pivot, as in a dance motion.

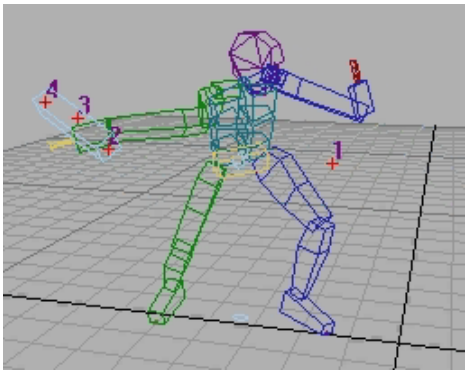
Note: Sliding footsteps display as a footstep with a line through the center.

Only Extract Footsteps Within Tolerance—Turns on Z -axis Tolerance. These controls filter out footsteps that do not fall within a given range of the ground plane. Use this when filtering motions, such as hopping or pitching a baseball, in which a foot might come off the ground and remain stationary, but its position is not intended as a footstep.

- **Tolerance—**Value is a percentage of leg length.
- **From Z Level—**Set a Z value (ground).

Flatten Footsteps to Z=0—Moves extracted footsteps to Z=0. Use this to flatten out minor differences in the height of the extracted footsteps.

Prop Bone



In 3ds Max, the CSM marker file format supports a prop bone in either or both hands (left, right, or middle). There are nine additional markers for the top, bottom, and middle of the three prop types. If the software detects these tracks, it creates a *dummy helper (page 2–16)* object.

The length of the prop is the average distance between the top and bottom prop marker during animation. The prop will be oriented in the plane of the three prop markers, and its origin will be at the bottom prop marker.

The dummies are named *bip01prop, bip02prop*, and so on. If a prop of that name already exists, its animation and size is reset upon loading the CSM file. The Marker Display dialog contains a check box for toggling display of the prop markers. The prop track names are LPRPB, LPRPM, LPRPT, MPRPB, MPRPM, MPRPT, RPRPB, RPRPM, and RPRPT. These stand, respectively, for left prop bottom, left prop middle, left prop top, middle prop bottom, middle prop middle, middle prop top, right prop bottom, right prop middle, and right prop top. You can change the track names by editing the MNM file.

The props don't necessarily have a relation to the left and right hands. For example, a prop could be a hat.

Motion Capture User Interface

Create or select a biped. > Motion panel > Motion Capture rollout

The tools on the Motion Capture rollout on the Motion panel are typically used for working with raw motion-capture data. You can also load standard BIP files using Load Motion Capture File. For example, you might do this if you want to loop the motion. This rollout includes tools for:

- Batch conversion of motion-capture files.
- Converting the motion-capture file stored in the motion capture buffer.
- Pasting one frame of motion-capture data from memory to selected biped limbs.
- Displaying raw motion-capture data as a stick figure.
- Displaying raw motion-capture trajectories.

The buttons in the bottom row are used mainly with marker files, although the calibration controls also work with raw BVH files. Import, calibrate, and filter marker files (CSM) using tools on the Motion Capture rollout. Markers are placed on an actor during motion capture to identify joints; calibration lets you adjust the biped relative to the original marker positions if necessary. Load only raw marker files with no key reduction or footsteps to enable marker calibration controls.

Note: For BVH and CSM file specifications, see the *BVH.rtf* and *CSM.rtf* documents on the program disc.

Motion-Capture Buffer

Raw motion data is automatically stored in the motion-capture buffer when files (CSM, BVH, and BIP) are loaded using *Load Motion Capture File (page 2–1065)*. This buffered raw motion data is

independent of the biped motion in your scene, and can be used in various ways:

- Use *Convert From Buffer* to try alternate filter settings quickly; this saves you from having to browse for the same file.
- After importing a BVH or BIP file using footstep extraction and key reduction, you can use *Paste From Buffer* on selected biped limbs (and COM) to paste keys from the raw motion-capture data to the filtered data; do this if critical motion has been lost in the filtering process.
- If you specify Load Buffer Only in the *Motion Capture Conversion Parameters dialog (page 2–1070)*, the motion file is loaded into the motion-capture buffer without altering the biped animation. Use this to paste posture and limb keys from any file onto the biped animation in your scene.

When you load a motion-capture file, the motion-capture buffer is loaded with motion-capture data from that file. This buffer is altered during calibration. It is also used to show the motion-capture markers and trajectories.

Internally, there is only one motion-capture buffer. It is often large, so its contents are not saved before a file load or a calibration. Therefore, if you undo a motion-capture file load, the contents of the motion-capture buffer do not change. That's why you'll still see old markers and trajectories. It is not possible to undo calibration.

Marker Files

Unlike a BIP or BVH file that contains limb rotation data, a *CSM marker file (page 3–969)* contains only marker position data. When a raw marker file is imported, only marker position data is buffered in the motion-capture buffer. 3ds Max uses the marker data to extract limb rotation data to position the biped. After using the calibration controls on the Motion Capture rollout to correct biped scale and posture relative to the markers, use *Convert From Buffer* to filter the raw marker data to key reduce and extract footsteps.

Note: A CSM marker file is an ASCII file.

See also

Importing Motion-Capture Data (page 2–1061)

Procedures

To use Convert From Buffer:

A motion-capture file should already be in memory. Use Load Motion Capture File on the Motion Capture rollout to import a motion-capture file if one is not already be in memory.

1. Select a biped.
2. On the Motion Capture rollout, click Convert From Buffer to display the Motion Capture Conversion Parameters dialog.
3. Adjust parameters, then click OK.

To compare raw and filtered trajectories:

1. Select a biped and turn on Show Buffer Trajectory on the Motion Capture rollout.
2. Click Trajectories on the Display rollout.

   As you select various biped parts, two trajectories are displayed. The yellow trajectory represents raw motion capture data in the motion capture buffer; the purple trajectory represents the filtered data.

To use Show Buffer:

1. Create a biped.
2. Use Load Motion Capture File on the Motion Capture rollout to import a motion-capture file.
3. Turn on Show Buffer on the Motion Capture rollout.

   A red stick figure appears, representing the raw motion-capture data.

Play the animation; the animation of the biped representing filtered motion-capture data and the red stick figure play back together.

Tip: For a very accurate visual comparison between raw motion capture data and filtered data, toggle Show/Hide Objects on the Display rollout to hide the biped. Toggle Show/Hide Bones in the same rollout to display only biped bones (the yellow stick figure), then play the animation with Show Buffer turned on. The two stick figures move together, and any discrepancies are easily spotted.

To learn how to use Show Buffer with the Fit To Existing parameter in Motion Capture Conversion Parameters rollout, see *To use Fit To Existing to import a motion capture file (page 2–1071)*.

To import a motion-capture file:

1. Select a biped.
2.  If necessary, click Load Marker Name File on the Motion Capture rollout (see *MNM Files (page 3–973)*).

   Note: This step is not required if the marker or joint names in the motion-capture file adhere to the **character studio** marker naming convention.
3.  Click Load Motion Capture File.

   The Open dialog opens.
4. Choose a file type: BVH, BIP, or CSM.
5. Select a file and click Open.

   The *Motion Capture Conversion Parameters dialog (page 2–1070)* displays.
6. Select the filter options you want, and click OK.

   The biped adapts itself to the motion data. If Footstep Extraction is turned on, footsteps appear.

   Note: Load raw marker data (No Key Reduction, Freeform) to enable the marker calibration buttons.

   Tip: Use a biped that does not have a mesh attached to it with Physique. Import motion-capture data with the idea of then saving a BIP file that can be used for any character. If skeletal scale information is loaded from a motion-capture file, a mesh with the Physique modifier applied might deform unnaturally.
7. If a marker file was loaded, turn on Show Markers as a visual aid for biped scale and limb correction.

   If correction is necessary, adjust biped scale first. Keyframe adaptation takes place in order to accommodate a biped scale change.

   The remaining steps in this procedure are optional, unless you need to calibrate motion-capture files.
8. Click Talent Figure Mode and use Non-Uniform Scale or Rubber Band Mode (on the Biped rollout) to size the biped to the displayed markers.
9. Click Talent Figure Mode again to exit the mode.

   Key Adaptation takes place when you exit Talent Figure mode. Now, biped limb positions relative to the markers can be adjusted.
10. Align the biped limbs to the markers, if necessary, and then click Adjust Talent Pose to compute the offset for the entire animation.
11. Use Save Talent Figure Structure and Save Talent Pose Adjustment to save a size and position offset to a FIG and a CAL file, respectively.

    Load these files in the Motion Capture Conversion Parameters dialog when you import similar marker files in the future.

At this point, you can use Convert From Buffer to extract footsteps and reduce keyframes. Both scale and position adjustments will be incorporated. Save the motion as an optimized BIP file.

Interface




Load Motion Capture File—Key reduce and extract footsteps from raw motion-capture data. Load BIP, CSM, or BVH files. After you open a file, the *Motion Capture Conversion Parameters dialog (page 2–1070)* opens. Following are descriptions of settings in the Open dialog.

Files of type—

- **BIP—**Filters the raw version of the motion-capture data that ships with 3ds Max. These are in a BIP format. Filter standard BIP files to convert footstep animation to a freeform animation; extract footsteps from a freeform animation; and to loop a BIP file.
- **BVH—**BioVision motion-capture data file. Contains the "actor's" skeletal and motion information. Once the motion-capture data is filtered and adjusted, save it as a BIP file for later use.
- **CSM—**Imports a **character studio** marker file (ASCII format file). Optionally, load a Marker Name file (MNM), a Talent Structure file (FIG), and a Talent Pose file (CAL) before loading a CSM file. Marker files should be loaded with no key reduction and no footstep extraction to enable the calibration controls.

Restructure biped to match file (BIP files only)—Turn on to change the biped structure to match the structure stored in the BIP file. The file loads with the stored biped structure. This is an Open dialog option.

This option is unavailable when you load a BIP file into a clip or onto a biped that is in Edit Clip mode, because all the clips in the motion flow would have to be adapted.

Note: CSM and BVH files always load with the biped structure stored in the file.

Set lowest starting foot height to Z=0 (BIP files only)—Sets the lowest starting foot height to Z=0. This is an Open dialog option. Default=on.

In 3ds Max, the height of a motion clip can be retained. This is important if you want to retain the height of a motion clip for motions adapted to characters of different sizes. If, for example, the character is jumping off a rock, and you want to retain the Z position of the character, you would turn this option off. Leave this option off if Motion Flow motions must be blended that begin and end at different heights, such as three clips that have the character mounting a bicycle, riding the bicycle, and dismounting the bicycle.

Turning off this option can, however, cause a jump in the motion during motion flow transitions. Turn this on for smooth transitions in Motion Flow mode. If adaptation takes place, the height is set so that the lowest foot at frame 0 starts at the Z=0 height. This lines up clips along the Z-axis and creates smooth transitions.

First load raw marker data and turn on Show Markers to help you decide if calibration is necessary. If both scale and position calibration are necessary, calibrate scale first (Talent Figure mode) and then calibrate limb position. 3ds Max adapts biped keys after biped scale is changed in Talent Figure mode. Orient the biped limbs relative to the markers and click Adjust Talent Pose to apply this offset to the entire animation. Save Talent Figure Structure and Talent Pose to a FIG and a CAL file, respectively. Load a FIG and a CAL

before loading a marker file that requires the scale and position offsets contained in these files.

Note: Calibration files can be loaded in the Motion Capture Conversion Parameters dialog before filtering marker files.

 **Convert from Buffer—**Filters the most recently loaded motion-capture data. This data is stored in the motion-capture buffer. Displays the Motion Capture Conversion Parameters dialog.

The most recently imported motion-capture file is stored in its raw form in the motion-capture buffer. Convert From Buffer provides a quick way to try new conversion parameters in the Motion Capture Conversion Parameters dialog.

 **Paste from Buffer—**Pastes a frame of raw motion-capture data to the selected parts of the biped.

After importing a motion-capture file, you might discover a subtle movement has been lost in the process of reducing keyframes. Paste From Buffer can add a keyframe from the raw motion-capture data to a selected biped body part to restore this motion. Turn on Auto Key before using Paste From Buffer, or click Set Key after using Paste From Buffer to store the new position in a key.

 **Show Buffer—**Displays raw motion-capture data as a red stick figure.

Use Show Buffer to compare raw and filtered motion-capture data. Ideally, the motion of the biped and the red stick figure are very similar. If this is not the case, alter the filter parameters and import the motion-capture file again, or select a biped object and use Paste From Buffer at selected frames to restore the lost motion.

Raw motion-capture data is buffered for the currently loaded or most recently imported motion-capture file, allowing easy comparison of the raw and filtered motion data. Show Buffer displays a red stick figure representing the raw buffered data; compare this to the filtered motion of the biped during playback.

 **Show Buffer Trajectory—**Displays buffered raw motion-capture data as yellow trajectories for the selected biped body parts.

Use Show Buffer Trajectory to display a trajectory based on the buffered raw motion-capture data for any biped body part. Use this in combination with Show/Hide Trajectories on the Display rollout to see how closely the raw and filtered data match.

This assumes a motion-capture file has been imported.

 **Batch File Conversion—**Converts one or more CSM or BVH motion-capture files to filtered BIP format. Displays the *Motion Capture Batch File Conversion dialog (page 2–1075)*.

 **Talent Figure Mode—**After loading a raw marker file, turn on Talent Figure mode to scale the biped relative to the markers. Calibration for the entire marker file takes place when you exit Talent Figure mode.

Keyframe adaptation takes place in order to accommodate the new biped scale; because of this, you should adjust the biped scale before adjusting the biped position relative to the markers.

Use Rubber Band mode on the Biped rollout and Non-Uniform Scale to size the biped in Talent Figure mode.

Ideally, you will not need to use this feature. When loading a motion-capture file, 3ds Max attempts to extract the appropriate figure scale from the given data. Use Talent Figure mode only if the extracted scale of the biped doesn't match the scale of the original talent. Even minor differences in scale will alter the motion.

Note: Calibration controls are enabled only when a marker or BVH file is imported in its raw form. Do not use key reduction or extract footsteps when you import a marker file for the first time.


Save Talent Figure Structure—After changing the biped scale in Talent Figure mode, you can store the changes into a FIG file. Use this file in the Motion Capture Conversion Parameters dialog to adjust marker files created by the same actor.


Adjust Talent Pose—After loading a marker file, use Adjust Talent Pose to correct the biped position relative to the markers. Align the biped limbs to the markers, then click Adjust Talent Pose to compute this offset for all the loaded marker data.

Note: Calibration controls are enabled only when a marker or BVH file is imported in its raw form. Do not use key reduction or extract footsteps when you import a marker file for the first time.


Save Talent Pose Adjustment—Saves a Talent Pose adjustment as a CAL file.

Save a CAL file after adjusting the biped relative to the markers. A CAL file is used for processing marker files that require the same adjustment. A CAL file can be loaded in the Motion Capture Conversion Parameters dialog during marker file importation.


Load Marker Name File—Loads a Marker Name (MNM) file to map incoming marker names in motion-capture files (BVH or CSM) to the **character studio** marker naming convention. Displays the Marker Name File dialog.

Load a CSM marker file—Browses for a marker file for use with a CSM file.

Load a BVH marker file—Browses for a marker file for use with a BVH file.

When a BVH file is loaded, checks for and reports unknown track names, but loads the file anyway. Reports if any required tracks were not in the file and if so aborts the file load

Use—Uses the marker name file when importing motion capture files

If necessary, load a Marker Name File before loading a BVH or CSM file. Marker Name files are bundled with Character Studio to map marker names in popular third-party marker files. Edit these ASCII files if the marker files you have use unique names for markers.

Note: For BVH and CSM file specifications, see the *BVH.rtf* and *CSM.rtf* documents on the program disc.


Show Markers—Opens the *Marker Display dialog (page 2–1075)*, with settings for specifying how markers are displayed.

Marker and marker names are displayed around the biped. You can use these to spot and adjust discrepancies; for example, the biped elbow position relative to the elbow marker. For information on how to correct these discrepancies, see *Talent Figure mode* and *Adjust Talent Pose*.

See also

Marker Files (page 3–969)

Motion Capture Conversion Parameters Dialog

Create or select a biped. > Motion panel > Motion Capture rollout > Load Motion Capture File button > Open a file. Load motion capture data. > Motion panel > Motion Capture rollout > Convert from Buffer button

Motion-capture and marker data typically have keys at every frame. Filtering motion-capture data reduces keys, making the job of altering or personalizing the motion data much simpler.

Other filtering options include footstep extraction, applying the skeletal structure stored in the motion-capture file to the biped, looping the data, importing a portion of the motion-capture file, and selecting tracks to load.

Motion-Capture Buffer

Any file imported using Load Motion Capture File is stored in its raw (nonfiltered) form in the motion-capture buffer. This buffer is used to try new filtering options with the Convert From Buffer command, and to paste keys from the raw motion-capture data to the biped using Paste From Buffer on the Motion Capture rollout. The Show Buffer command displays a stick figure that represents the buffered data.

You can create your own library of imported and optimized motion-capture data by saving *.bip* files for use with other characters or as part of a longer script in Motion Flow mode. Use a biped that has no mesh attached to it with Physique. You import the data, adjust it to your liking, and save it as a *.bip* file. You can also run standard *.bip* files through this filtering process to create loops or to extract footsteps from a freeform animation.

Note: Marker files contain *position* data. Regular motion-capture files contain joint *rotation* data.

biped.ini Parameters

The *biped.ini* text file in the *\plugcfg* directory has parameters for smoothing the values of motion-capture joint angles for the body's center of mass, the spine, and the head during import. The lines appear as below:

```
MocapHeightSmoothing=0
MocapBodyHorzSmoothing=0
MocapBodyRotSmoothing=0
MocapSpineSmoothing=0
MocapHeadSmoothing=0
```

Higher values of smoothing will cause the importer to filter and "smooth out" the data using a gaussian filter, taking out spikes and jerks in the motion. These values must be integers, and can range from 0 to any high number, but practical values would probably always be less than 10. For most cases, smoothing is not needed, so the default values are set to zero.

Procedure

To use Fit To Existing to import a motion-capture file:

Use a motion-capture file that contains footsteps and other motion. A handspring motion would be ideal for this example.

1. Select a biped.
2. Click Motion panel > Motion Capture rollout > Load Motion Capture File.
3. Select a motion-capture file.

   The Motion Capture Conversion Parameters dialog displays.
4. Set Footstep Extraction to On, then click OK.

   The motion-capture file loads.
5. Select Motion Capture rollout > Show Buffer to display raw motion-capture data as a red stick figure, then click Play.

   During the "hand spring" period of the playback, the hands on the red stick figure representing the raw motion data touch the ground. The biped, using the filtered data, is positioned higher and cannot reach the ground.

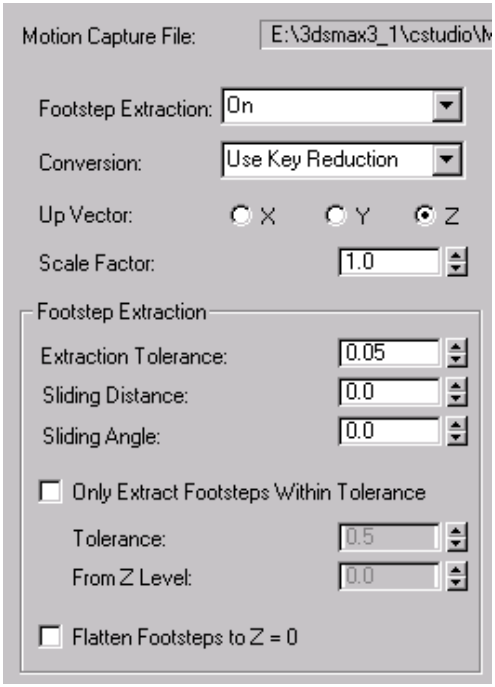
   Dynamics calculates the biped higher than the raw motion data because the footsteps before and after the handspring are so far apart in time. Creating a freeform period in Track View — Dope Sheet and reloading the same file using the Fit To Existing option will match the biped's position to the motion-capture position during the handspring.
6. Find the Footsteps track for the biped in Track View.

7. Right-click the footsteps area of the Track View Edit window, then select Edit Free Form (No Physics) in the Footstep Mode dialog.
8. Click the "handspring" area between the footsteps. It turns to a solid yellow.

   Motion-capture data will replace the keys in this freeform period.
9. Reload the same motion-capture file using the Fit To Existing option on the Motion Capture Conversion Parameters dialog.

   The freeform area created in Track View is replaced with motion-capture data. The biped closely matches the red stick figure during the handspring part of the playback.

Interface



Motion Capture File—Displays the file to be imported.

Footstep Extraction—Motion capture data can be applied to the biped in one of three ways:

- **None: Freeform—**No footsteps are extracted.

  For swimming or flying motion data, footstep extraction is not necessary. For a traditional approach to character keyframing, use this option to keyframe the biped without footsteps or Biped Dynamics; this is essentially a freeform animation.
- **On—**Extracts footsteps.

  Direction and style of the motion-capture data are easily edited.

  Allows changes to the toe structure of the biped after import; footsteps will readjust the character's motion to maintain correct foot-toe-ground contact at all times, a common problem associated with motion-capture import.

  Inappropriate "sliding feet" in the motion data are corrected.
- **Fit to Existing—**Fits to existing footsteps.

  Use with motion data that has both footstep motion and flying, swimming, falling, or tumbling motions.

  First load the motion data using Perform Footstep Extraction. Create a freeform period for the flying, swimming, or falling portion of the data in Track View, and then reload the same motion capture file using the Fit to Existing option. The freeform area is loaded with data from the motion capture file without the influence of biped dynamics.

Conversion—Chooses the type of key processing.

- **Use Key Reduction—**Reduces keys for simpler key editing.
- **No Key Reduction—**Does not reduce keys. Use this on files that are already key reduced or if you want to work with all the data in a raw motion-capture file.

  Note: Marker files imported for the first time should be loaded with no key reduction or

footstep extraction to enable the calibration controls on the Motion Capture rollout.

- **Load Buffer Only—**Does not apply the data to the biped, but loads the data to the motion-capture buffer only. Use this either to compare your edited version with the original or to paste postures from the motion-capture buffer to the biped in the scene.

  Tip: To compare the filtered data with raw motion-capture data, use Show Buffer on the *Motion Capture rollout (page 2–1065)* rather than turning off key reduction.

Up Vector—Sets the vertical axis used in the motion-capture data.

Scale Factor—Multiplies the stored talent size by this value and size the biped accordingly.

Footstep Extraction group

Options here are active when Footstep Extraction is on.

Extraction Tolerance—Sets the sensitivity of footstep extraction. **character studio** determines if the footstep is there by checking that the foot does not move beyond the distance determined by the Extraction Tolerance value. Smaller numbers are more sensitive and extract more footsteps. The value is a percentage of foot length.

The default value is 0.15. Increase this value to 0.2 or 0.25 if too many footsteps are generated.

Sliding Distance—Creates a *sliding footstep (page 2–1064)* when positional tolerance is reached. This value is a percentage of foot length. By default, the foot must slide its own distance (100), before a sliding footstep is created.

Use this with motion-capture files that contain sliding feet. A sliding footstep can be created manually by setting IK Blend > 0 for a biped foot at a "touch" state key (biped foot first touches a footstep).

Note: Sliding footsteps display as a footstep with a line through the center.

Sliding Angle—Creates a *sliding footstep (page 2–1064)* when rotational tolerance is reached. This value is in degrees; the default is set high (360 degrees). The foot must make a complete turn before a sliding footstep is created.

Use this with motion-capture files that contain feet that pivot, as in a dance motion.

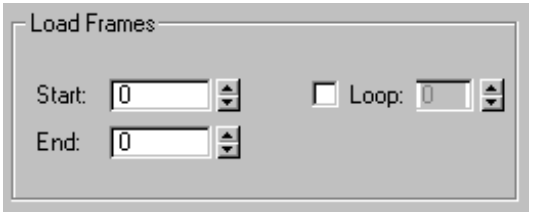
Note: Sliding footsteps display as a footstep with a line through the center.

Only Extract Footsteps Within Tolerance—Turns on Z -axis Tolerance. These controls filter out footsteps that do not fall within a given range of the ground plane. Use this when filtering motions, such as hopping or pitching a baseball, in which a foot might come off the ground and remain stationary, but its position is not intended as a footstep.

- **Tolerance—**Value is a percentage of leg length.
- **From Z Level—**Set a Z value (ground).

Flatten Footsteps to Z=0—Moves extracted footsteps to Z=0. Use this to flatten out minor differences in the height of the extracted footsteps.

Load Frames group



Start—Start importing at this frame. Default is frame 0, the first frame.

End—Stop importing at this frame. Default is the last frame of the clip.

Loop—Loop the data by the value set here.

This is relative. Succeeding loops start where the previous loop left off. The clips are not blended and may require editing unless the original clip was designed to loop.

Use this for clips designed to loop.

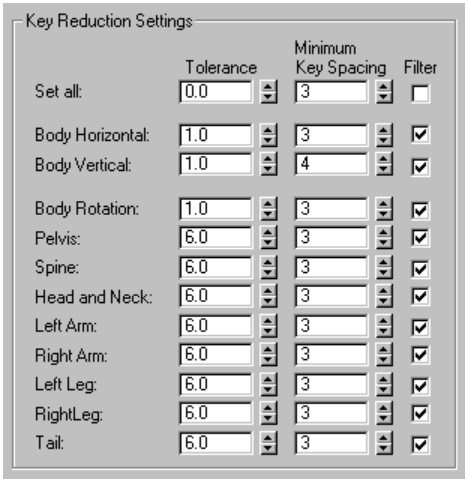
Note: This often works best if Footstep Extraction is tuned off.

Key Reduction Settings group

Key reduction keeps the original motion intact and intelligently filters out more than 80 percent of the keys in the motion-capture file, making the process of altering the biped animation much simpler.

Use the Tolerance and Minimum Key Spacing settings to fine-tune key reduction for a specific track (body part). Specify whether or not to filter a track in the Filter column.

These settings are available only if Conversion is set to Use Key Reduction.



Tolerance—Sets the maximum angular or positional deviation for a track.

Values are in units of translation for position tracks, and in degrees for rotation tracks.

Minimum Key Spacing—Sets the minimum number of frames between keys.

Tolerance is computed first, then Minimum Key Spacing computes further key reduction.

A Minimum Key Spacing value of 10 for the head track ensures that no two keys are closer than 10 frames for this track.

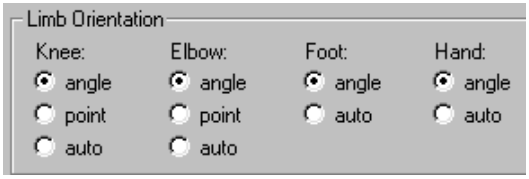
Filter—Turn off to prevent filtering of the motion capture data into a track. When this is off, there is no key reduction for the track.

Set All—Forces all tracks to the values set in these fields.

Higher values here can determine how much key reduction is possible while preserving the original motion.

Limb Orientation group

The biped elbow and knee hinge joints are perpendicular to the triangles formed by the shoulder-elbow-wrist and hip-knee-ankle respectively. Resolve errors in the motion-capture data that break this rule by using either the angle or point method.



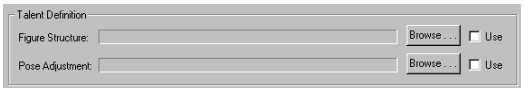
Angle—Moves the knee or elbow position to create the biped joint key.

Point—Rotates the shoulder-elbow-wrist or hip-knee-ankle to create the biped joint key.

Auto—Auto reads exact hand and foot positions from the motion-capture data; **character studio** then places the knees and elbows in a natural position. For marker files involving running and walking, this option can clean up the data nearly

instantly, regardless of how many markers were used and where they were placed.

Talent Definition group



Loads a Figure Structure File (*.fig*) and a Pose Adjustment file (*.cal*) prior to importing a marker file. Typically you correct a marker file by importing it and adjusting the biped scale and limb positions relative to the markers, then saving a *.fig* and a *.cal* file using Save Talent Figure Stucture and Save Talent Pose Adjustment on the Motion Capture rollout. These files can then be loaded in the Talent Definition area when importing marker files created by the same actor in a motion capture session.

Figure Structure—Loads a *.fig* file.

Pose Adjustment—Loads a *.cal* file.

Browse—Browses for a *.fig* or *.cal* file.

Use—Use either or both the *.fig* and *.cal* files to adjust marker files during a marker file import procedure.

Load Parameters—Loads a motion capture parameter file (*.moc*).

Save Parameters—Saves a motion capture parameter file (*.moc*).

Motion Capture Batch File Conversion Dialog

Create or select a biped. > Motion panel > Motion Capture rollout > Batch File Conversion button

Converts one or more CSM or BVH motion-capture files to filtered BIP format.

Interface

Source File Selection—Opens a dialog where you can specify motion-capture files to convert. Use Shift+click to select all files between two that you click. Use Ctrl+click to add individual files to the selection.

Destination Directory—Specifies the directory in which to store the filtered files.

Specify Conversion Parameters ...—Lets you choose how to specify conversion parameters:

- **Specify Conversion Parameters Once—**Imported files use a single set of conversion parameters.
- **Specify Parameters For Each File—**You specify conversion parameters for each file.

Marker Display Dialog

Create or select a biped. > Motion panel > Motion Capture rollout > Show Markers button

The Marker Display dialog lets you specify how markers from *.csm* files are displayed in the viewports. For further information on markers and *.csm* files, see *Character Studio Marker Files (page 3–920)*.

Interface



Show Recognized Markers—When on, displays the markers that **character studio** recognizes.

- **On Selected Objects—**Displays the markers on selected objects only.
- **On All Objects—**Displays the markers on all objects.

Show Prop Markers—Enables the display of markers on *prop bones (page 2–1065)*.

Show Unrecognized Markers—Displays the markers that **character studio** does not recognize.

Physique



Use the Physique modifier to attach a skin to a skeleton structure such as a biped. The skin is a 3ds Max object: it can be any deformable, vertex-based object such as a mesh, a patch, or a shape. When you animate the skeleton with skin attached, Physique deforms the skin to match the skeleton's movement.

Animating the underlying skeleton enables you to animate a single contiguous model of a character that bends, creases, and bulges about an arbitrary number of joints within the attached skeleton.

With Physique, you can define how the skin behaves when it deforms. For example:

- You can make portions of the skin solid, excluding them from Physique's deformation, though solid portions still move along with the root node of the skeleton they are attached to. These solid portions are said to be root vertices.
- You can make portions of the skin deformable. They move with the deformation spline, the smooth curve running through the links of the skeleton they are attached to.
- You can make portions of the skin rigid, directly moving along with the skeleton they're attached to.
- You can add bulges to simulate bulging muscles. Bulges are controlled by editable cross sections of the skin, and by bulge angles that you set.
- You can add tendons to distribute the effect of one bone's motion to areas of the skin other than those around the bone itself.
- You can save Physique data to a Physique (*.phy*) file, preserving data common to all objects sharing a given Physique modifier. Later, you can reload the data file, either to restore the data that belongs to a particular skin or portion of skin.

Physique works with bipeds created and animated using the Biped plug-in, and with 3ds Max hierarchies, including the Bones systems. Physique also works with bones that are not in a hierarchy and splines.

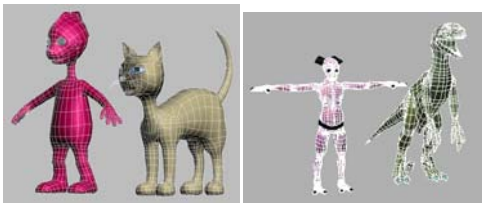
See also

Posing the Biped (page 2–847)

Using Physique

Creating a Skin

A mesh deformed by a skeletal structure is called a *skin*. In **character studio**, Physique is a modifier you apply to a skin to make it deformable by the biped, or by another skeletal structure. The images below show meshes for different skeletons.



Meshes for Physique

A skin used with Physique can be any 3ds Max object that has vertices or control points. Specifically, a skin can be:

- An editable mesh or editable poly object. This is the most commonly used type of object for Physique. Often, it has been collapsed from an object with modifiers, or a compound object.
- An uncollapsed object with modifiers or a compound object.
- A parametric geometry primitive such as a cylinder.

  Geometric primitives are useful mainly for simple applications of Physique; for example, a cylinder with two bone links to depict an arm.
- A patch object.
- A spline or text shape.
- A NURBS object.
- A Free From Deformation (FFD) modifier.
- A mesh object you import from another application such as AutoCAD®.

Tip: Although you can apply Physique to a compound object or an object with modifiers, you should collapse the stack before applying Physique, if possible. This will maximize performance and reduce the amount of work you'll need to do to get Physique working properly. An exception is the Optimize modifier, which can be useful for improving performance when applied below Physique on the stack.

Warning: After you collapse a compound or modified object, you can no longer edit it parametrically. If you work extensively with complex meshes of this sort, you can save two *.max* files: one to contain the original, editable objects and modifiers, and the other to contain only the collapsed mesh.

You can create a figure's skin out of several objects. For example, you could have separate objects for the torso, legs and arms. In this case, select all the objects and apply Physique to all of them at once.

Posing the Skin

When you create a skin to use with a biped figure, you should pose the arms and legs of the skin in a standard *reference pose*.



Mesh in reference pose for use with bipeds

Use the following positions when you create the reference pose:

- Spread the legs somewhat apart, at a *parade rest* position.
- Spread the arms wide, level with shoulder height. The hands should be level with the arms, not dangling: palms facing down, fingers straight and slightly spread apart.
- Position the head so it will face in the correct direction when you load the biped's at-rest

standing pose. If the skin and biped are for a figure that stands erect, position the head normally. If the character stoops forward, for example, a chimpanzee, make the head face upward so that it will face forward after the spine is bent.

As a general rule, create a reference pose that has the limbs outstretched, but otherwise represents the character's natural at-rest posture.

Skin Simplicity

The degree of detail on the model makes a difference in how well it works for skinning. On one hand, your skin must have a sufficient number of vertices so Physique can deform it smoothly. On the other hand, the fewer vertices the mesh has, the easier it will be for you to adjust Physique for the mesh. In addition, a highly complex mesh can slow your system's performance when working with Physique.

If you plan to use an Editable Poly mesh, create the mesh with evenly-sized, rectangular polygons. Avoid using long, triangular polygons, as these do not deform smoothly with Physique. These attributes are particularly important around the hip and shoulder areas.



Models with evenly-sized rectangular polygons

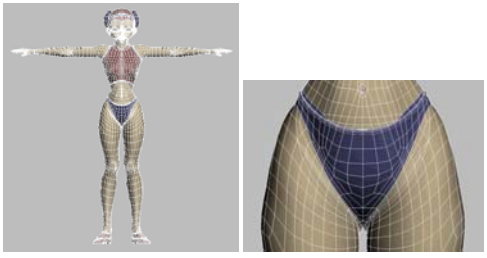
A common workflow is to create the simplest possible version of the mesh, apply Physique to it, and apply the MeshSmooth modifier above Physique on the stack. This makes Physique as easy as possible to work with, but retains mesh smoothness for rendering.



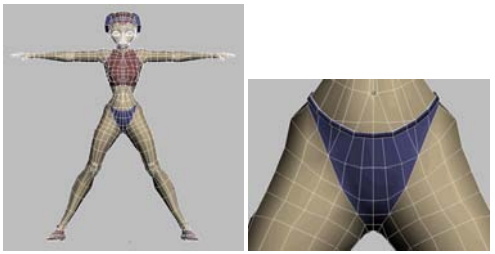
Same models rendered with MeshSmooth applied

If you want to use a particular mesh but its polygons violate these guidelines, use 3ds Max to add or delete edges and polygons as needed before applying Physique to the mesh. Usually, these rules are not so important for the facial area of the model, where little deformation will take place with Physique.

For example, the character shown below was originally modeled as designed, but with no regard for the best polygon distribution for Physique. The numerous odd-shaped polygons (especially around the hips), and the pose with the legs too close together, would have made this model difficult to work with after Physique was applied.



The model was altered to work better with Physique by deleting edges and vertices, and reshaping many of the polygons in the hip area. The model was also altered to make the character assume the reference pose.



After MeshSmooth is applied to the simplified mesh, renderings of the two models look identical. However, the second model works much better with Physique.

Using Physique with a Biped

The skeleton to which you attach a skin using Physique can be a 3ds Max hierarchy, bones in a hierarchy, bones not in a hierarchy and splines. Physique deforms the skin based on the relative position of the *bone* or links in the hierarchy. Specifically, it uses the length of each link and the angle between two connected links; it can also use the scale of a link.

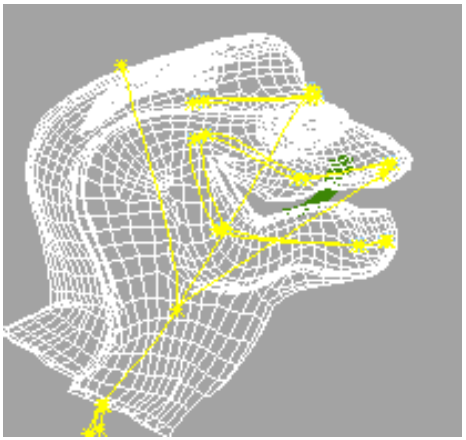
The skeleton hierarchy can also be a 3ds Max system object that defines a behavior as well as a hierarchy. Three kinds of objects that are especially useful with Physique:

- Bipeds are provided by **character studio**.
- Bones are a standard Systems object provided with 3ds Max.
- Splines can be used rather than a "bones" hierarchy.

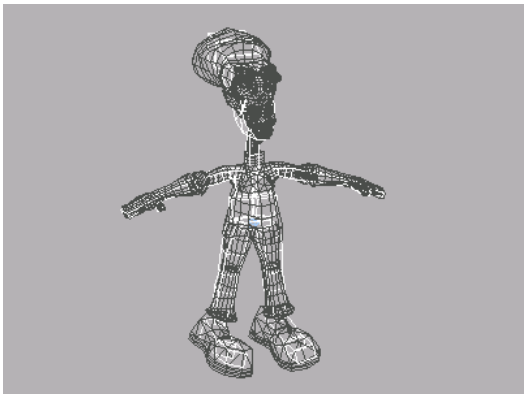
 You create both bones and bipeds using the Systems object category on the Create panel.

Bones are useful for facial animation, a face with moving lips for example, or for non-bipedal characters. Bipeds are the system of choice for humanoid and other bipedal characters.

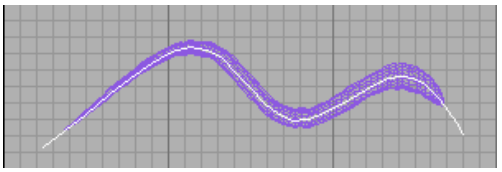
Usually you create the Physique skin before you create the skeleton, because you must adapt the skeleton's dimensions to the dimensions of the skin, in order to optimize vertex assignment to the links in the hierarchy.



This Physique hierarchy is created with dummy objects linked to each other.



Physique mesh with a biped skeleton.

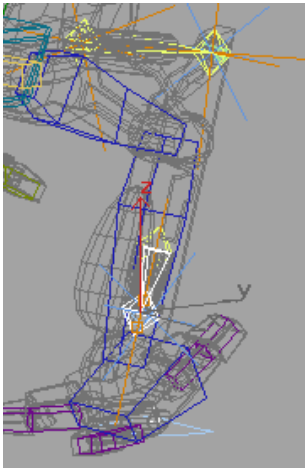


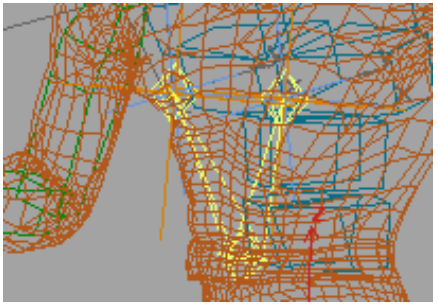
Physique can use a spline to defrom the mesh.

Using Physique with 3ds Max Bones

3ds Max bones with or without the IK Controller can be used with the biped for various effects. They can be used with Physique to add extra links and envelopes for any character, or to animate assemblies on a robot or mechanical character. Bones can also be added to animate extra appendages, a hat, a jaw and so on.

Note: If 3ds Max bones using the IK Controller are linked to the biped, the Auto Key button must be left on while the biped is positioned.





Two examples of compressible bones.

In the top image, 3ds Max bones are used to animate the linked piston assemblies. The bottom image shows 3ds Max bones used for added control when Physique is applied. In both cases the bones compress automatically as the biped is positioned.

Bones That Compress

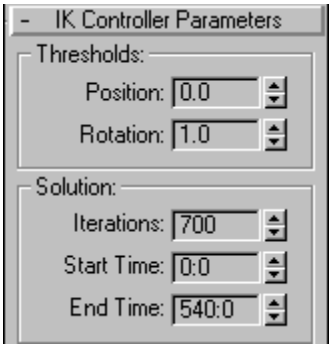


One way to use bones with the biped is to use the Select and Link tool on the 3ds Max toolbar to link the root of the bone to one part of the biped and have the bone End Effector follow another part of the biped. As the biped moves these bones compress and expand.

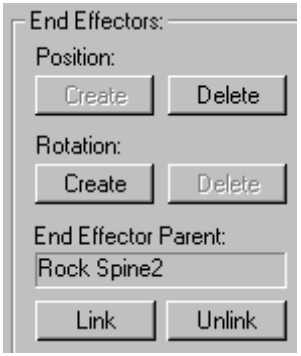
In the image of the mechanical leg, piston objects are linked to the bones. As the bones compress and expand, they also animate the linked piston assembly. As the bones compress and expand in the abdominal area in the bottom image, the envelopes created by Physique for the bones compress and expand to animate this area on the mesh.

IK Solution

In the IK Controller Parameters rollout, you specify the accuracy of the IK solution and the frame duration of the solution. For bones that compress, you generally want an accurate solution to ensure that the bone follows the end effector perfectly, without any drift. In the image of the rollout Position has a value of 0, the least amount of allowable distance (Rotation is at its default value of 1, the Rotation End Effector is not used in this example). Iterations is set high, ensuring that that the solution is accurate, even if the character is scaled down a great deal. Start and End Time should span the animation.



There are many ways to use bones with the biped. You could automatically animate the bending of a hose that is attached to a character's shoulder and mouth when the biped head rotates, for example. Bones attached to the biped create links and envelopes when Physique is applied; these extra Physique links offer localized skin control if necessary.



Procedures

To make compressible bones (bones with end effectors):

1. Create a bones system with end effectors.
2. Link the root of the bone to the appropriate biped object.
3. On the Motion Panel specify an End Effector Parent.
4. Specify another part of the biped.
5. On the Motion Panel set the Position Threshold to 0.
6. On the Hierarchy panel > IK > Sliding Joints rollout, turn on all the sliding parameters.
7. Animate the character.

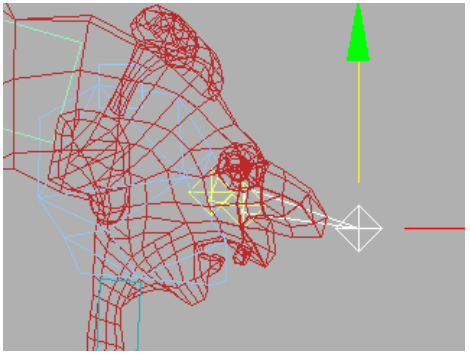
   The bones expand and compress with the motion of the character.

To add a bone after Physique is applied using Reinitialize:

1.  Turn on Figure mode.
2. Add a bone where it is needed.
3. Link the root node of the bone to the biped.
4.  On the Physique rollout, click Reinitialize.
5. On the Physique Initialization dialog, click Initial Skeleton Pose and then click Include New Bones.

   Vertex Link Assignment turns on also.
6. Click Initialize.
7. Adjust envelopes in Sub-Object Envelope.
8. Turn off Figure mode.

   If you want the end of the bone to follow the biped, select the bone and delete the bone end effector in the Motion Panel



This bone is used to animate the character's nose.

To add a bone after Physique is applied using Add (Add Bone):

1.  Turn on Figure mode.
2. Add a bone where it is needed.
3. Link the root node of the bone to the biped.
4. On the Physique Bones rollout, click Add.
5. In the viewports, click a bone.

   Repeat until all bones are added.
6. Adjust envelopes.
7. Turn off Figure mode.

Using Physique with 3ds Max Objects

The skeleton to which you attach a skin using Physique can be a 3ds Max hierarchy, bones in a hierarchy, bones not in a hierarchy, and splines. Physique deforms the skin based on the relative position of the *bone* or links in the hierarchy. Specifically, it uses the length of each link and the angle between two connected links; it can also use the scale of a link.

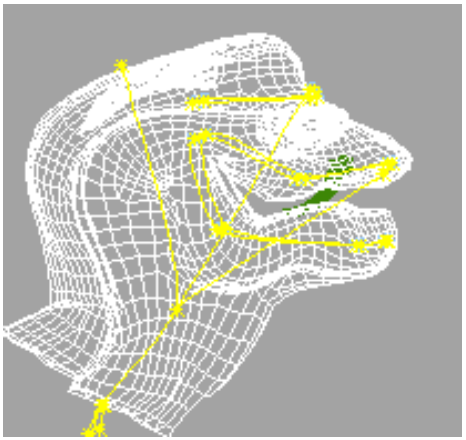
The skeleton hierarchy can also be a 3ds Max system object that defines a behavior as well as a hierarchy. There are two kinds of objects that are especially useful with Physique:

- Bones are a standard Systems object provided with 3ds Max. Bones can either be hierarchically linked or *floating*.
- Splines can be used rather than a "bones" hierarchy.

 You create bones using the Systems object category on the Create panel.

Bones are useful for facial animation, a face with moving lips for example, or for non-bipedal characters.

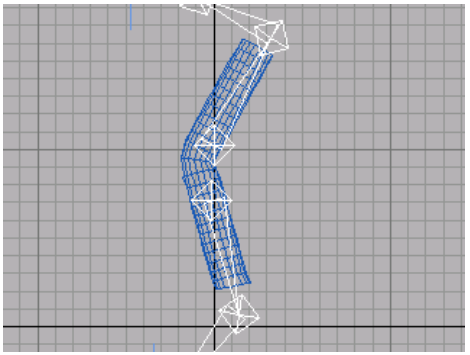
Usually you create the Physique skin before you create the skeleton, because you must adapt the skeleton's dimensions to the dimensions of the skin, in order to optimize vertex assignment to the links in the hierarchy.



This Physique hierarchy is created with dummy objects linked to each other.

Floating Bones

Floating bones are bones that are not linked together and know nothing about each other. By adding floating bones to Physique, you can deform the mesh by animating the bones. This is in contrast to using Attach To Node and clicking the root of a hierarchy, like the biped pelvis. For Attach To Node to work all the bones should be linked together.

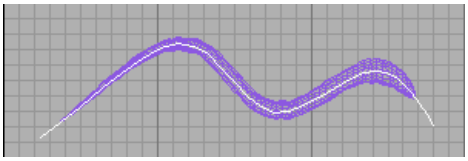


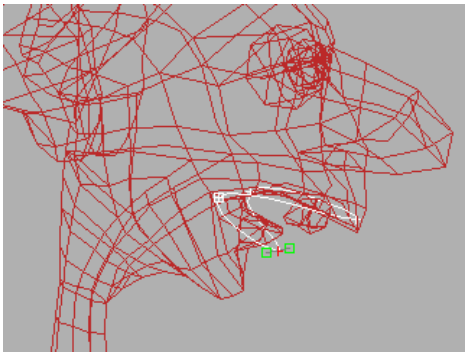
Spline-Based Physique Deformation

In addition to the biped and 3ds Max bones, Physique supports spline and NURBS curves for mesh deformation. By animating vertices on a spline, you can animate the mesh. You can use this technique for facial animation or to deform any mesh.

To control a mesh with a spline:

1. Place a spline inside a mesh.
2. Select the mesh and add the Physique modifier.
3. In the Physique Floating Bones rollout, click Add and select the spline in the viewports.
4. Select the spline, turn on Vertex Sub-Object, and move the spline vertices to animate the mesh.





See also

Using Physique with a Biped (page 2–1079)

Applying and Initializing Physique

After you have created a skin and a skeleton, and fitted the skeleton to the skin, you apply the Physique modifier to the skin. The process entails these steps:

- Selecting the mesh
- Turning on Figure mode (if a biped is used)
- Adding the Physique modifier to the stack
- Attaching the Physique skin to the skeleton
- Initializing Physique

Procedure

To attach a mesh to a bones hierarchy using Physique:

1. Select or create a bones hierarchy.
2.  Position the bones hierarchy inside the mesh.

   Tip: When you use a bones hierarchy with Physique, use frame 0 as your "figure mode." Don't include frame 0 in your animation. Use it as the place where you position the bones and fit them to the mesh.
3.  Select the mesh.
4.  Go to the Modify panel. Use the Modifier List to apply Physique to the mesh.
5.  Click to turn on Attach To Node. In a viewport, select the root bone in the bones hierarchy.

   The *Physique Initialization dialog (page 2–1111)* is displayed.
6. Click Initialize, accepting the default settings.

   When you use Attach To Node to attach a Physique skin to a hierarchy, the Physique Initialization dialog appears.

   Note: This dialog also appears when you want to reset Physique settings by clicking Reinitialize in the Physique rollout.

Physique initialization settings affect how envelopes are created and blending is handled. The Link Settings, Joint Intersections, and Cross Sections rollouts are used later to change default settings globally. For this reason, the Vertex-Link rollout is open when the dialog appears. This is where you determine the following default settings:

- Whether envelopes should be used to manage the vertex-to-link assignments
- Whether Physique uses deformable or rigid envelopes to manage the vertex-to-link assignments
- Number of links considered for blending

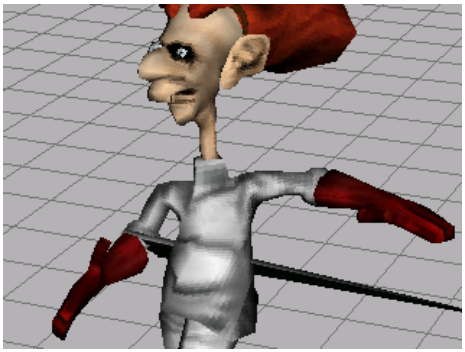
See also

Physique Rollout (page 2–1106)

Previewing Motion

After you have attached the skin to the skeleton, Physique deforms the skin when the skeleton is animated. From this point on, you can preview animation to see how the Physique skin deforms and whether you can use it in a finished animation or need to correct and refine it further.

Very likely, you'll need to adjust some of the default envelope settings to ensure all vertices are being properly handled. It's by moving the skinned character in the viewports that you'll see vertex assignment problems.



A stray vertex on the right elbow not encompassed by an envelope.

See the *Envelope Sub-Object (page 2–1130)* topic for procedures and interface.

Tip: Physique skins are usually too large for 3ds Max to play them back in real time at 30 frames per second. However, speed improvements in **character studio** will allow fluid motions of many skinned meshes. When you attempt to play back at 30 fps, 3ds Max may drop frames during playback, which makes it hard to see how Physique has animated the skin. Use the Time Configuration dialog to turn off real time playback. The animation plays back more slowly than usual, but it plays every frame.

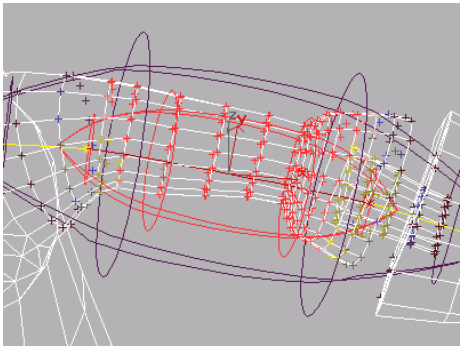
Tip: You can also specify which elements of the characters not to display in the viewports, and thereby speed up redraw. See *Physique Level of Detail Rollout (page 2–1108)*.

For a skin attached to a complicated skeleton such as a biped, you will almost always need to correct some vertex assignments before the skin animates correctly. Most often, you'll do so simply by adjusting the envelope's shape. See .

Envelopes and Vertex Assignments

Envelopes

Envelopes are Physique's primary tool to control skin deformation. Envelopes define an area of influence about a single link in the hierarchy and can be set to overlap adjacent links. Vertices that fall in the overlap area of the envelopes are weighted to produce smooth blending at joint intersections. Each envelope comprises a pair of inner and outer bounds, each with four cross sections.

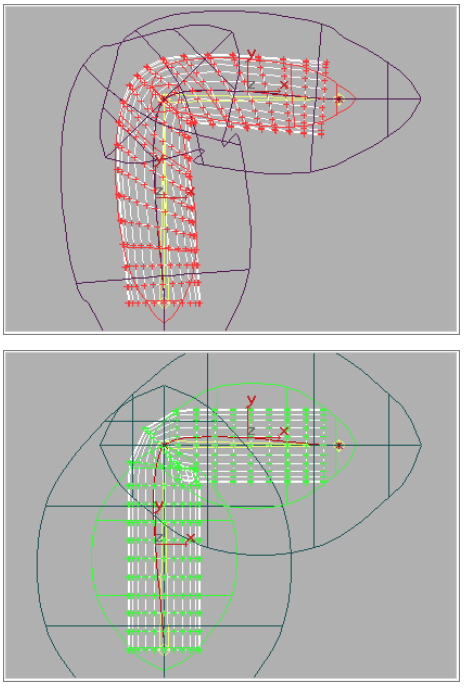


Envelopes

Deformable and Rigid Envelopes

There are two envelope types: deformable and rigid.

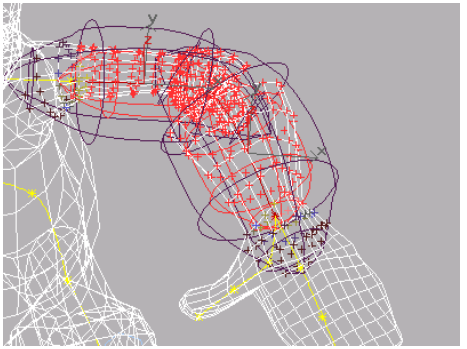
- Deformable envelopes influence vertices they encompass to follow the deformation spline created through the hierarchy. Only the vertices encompassed by deformable envelopes can be affected by bulge angles or tendons.
- Vertices in a rigid envelope are linked to the node (the bone) and move in an immobile relationship to the link. Vertices in a rigid envelope, however, are deformed (blended) in the overlap area of other envelopes. There is a twist parameter in Link Sub-Object that can be enabled in a rigid envelope. This allows the rigid envelope to twist along the length of the link.



Deformable envelopes (above); Rigid envelopes (below)

Typically, you use deformable envelopes when you attach a mesh to the biped pelvis to produce a soft, flexible skin. Later you reassign certain links, such as the character's head, to the rigid envelope to minimize the deformation. For special cases, you can turn on both deformable and rigid envelopes for the same link. This advanced feature allows you to average the effect of the two types of skin deformation for additional firmness in the skin. The forearms and legs are sometimes good candidates for this.

Blending Between Links

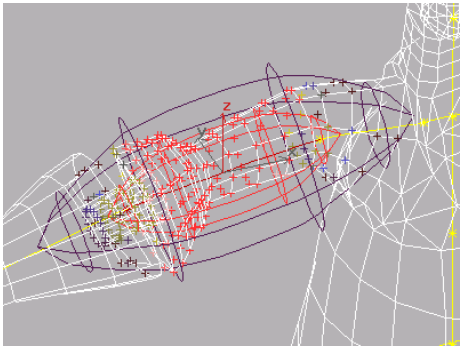


Blending between envelopes controls deformation.

Blending controls specify the influence of overlapping envelopes on vertices contained within the overlap area. By the actual shape of adjacent envelopes, you can control the degree of influence each has on blending at the overlap area.

You can further control the number of envelopes that participate in the blending effect, or whether no blending takes place at all. Where you specify no blending, a vertex in an overlap area is influenced by a single link only.

Inner and Outer Bounds



Inner and outer bounds

Where a vertex falls within the inner and outer bounds determines the percent of influence of the attached link(s). Vertices that fall within the inner bound have a weight of 1. Vertices that fall between the inner and outer bound have a weighted value that falls off to 0 at the outer bound.

Vertices are color—coded in the viewports according to their weight; the color is based on parameters you specify in the *Blending Envelope Display Options Dialog (page 2–1125)*.

See also

Envelope Sub-Object (page 2–1130)

Adjusting Default Envelope Shape

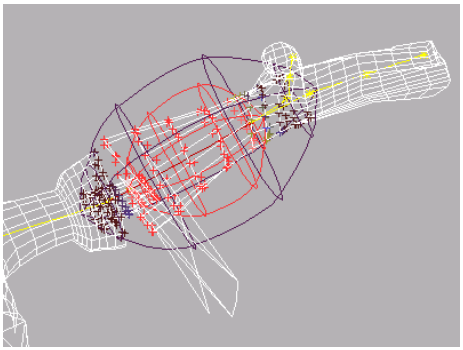
When you want to address vertex-linking anomalies, the first step is to resize or reposition the envelope about the problem link. Most often, the problem is an envelope that istoo small to surround, and thereby influence, certain vertices.

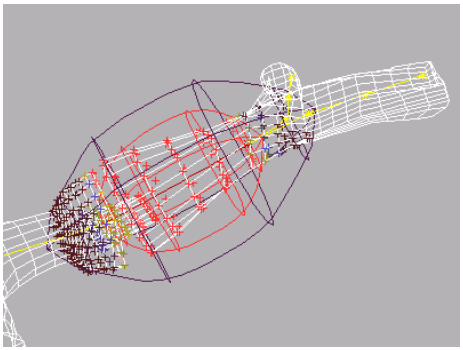
This can happen because your biped character misaligns slightly with the mesh, or the link lengths used to create envelopes with unbounded bones systems links were insufficient to surround all vertices attached to the link.

Another potential source of problems is overlapping inner bounds. This can sometimes create too strong a deformation at the joint.

You'll note problems with vertices as you preview motion. It sometimes looks like vertices got left behind when the link they were supposed to be attached to moves in 3D space. In fact, that's just what is happening: they weren't assigned to any link at all, and remain where they were at the initial skeletal pose: the pose of the mesh and its skeleton at the time Physique was applied and initialized.

The other case would be vertices being moved too much, resulting in a dent in the mesh. You'd fix this by decreasing the overlap of the envelopes affecting a joint.





Stray vertices (above) reassigned to correct link(s) by modifying envelope shape (below).

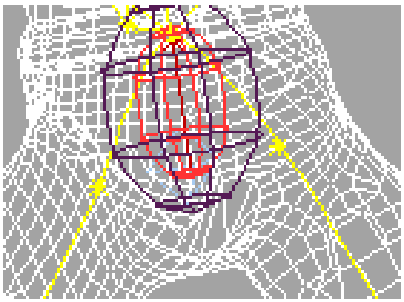
Procedures

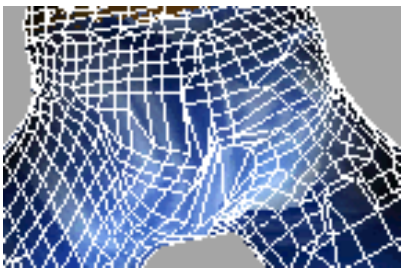
To adjust default envelope shape:

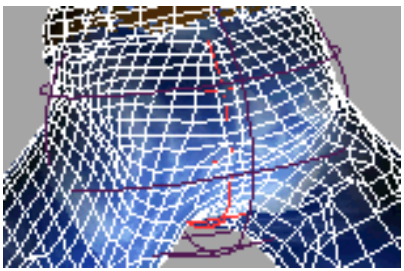
1. In the Selection Level group of the Blending Envelopes rollout, turn on Link. In a viewport, select any link.
2. In the Parameters group, turn on the Inner, Outer or Both option and then increase the value of Radial Scale until the outer envelope just surrounds the character mesh.

   Note: With Both turned on, both the inner and outer bounds scale together.
3. Increase the Parent Overlap value until the inner (red) envelope appears through the mesh.
4. If necessary, decrease the Falloff value to strengthen this envelope.

To adjust envelopes around the biped's pelvis:







The Biped Triangle Pelvis option was used in this model. Triangle Pelvis creates additional links from the legs to the lower spine.

Top: The default envelope from the pelvis to the lower spine object is too small.

Middle: Mesh appears "scrambled" because the pelvis doesn't control all the vertices that surround it.

Bottom: Increasing the envelope size corrects the mesh problem in the pelvis area.

1. In the Selection Level group, turn on Link. In a viewport, select the middle link.
2. In the Parameters group, turn on Both, and then increase the value of Radial Scale until the outer envelope just surrounds the mesh in pelvis area.

Note: With Both turned on, both the inner and outer bounds scale together.

3. Increase the Parent Overlap value until the inner (red) envelope appears through the mesh.
4. If necessary, decrease the Falloff value to strengthen this envelope.
5. To finish correcting the pelvis area, adjust the Overlap of the thigh and buttock envelopes.

To select and edit cross sections:

1.  On the Selection Level rollout, turn on Cross Section.
2. In the Envelope Parameters group, turn on Inner or Outer if you want to work with cross sections solely in that part of the envelope. Turn on Both to be able to work on both Inner and Outer cross sections at the same time.
3.  Turn on Select Object on the Main toolbar, and in a viewport, click the cross section you want.

   As you move the cursor over a selectable cross section, it changes to a plus shape, letting you know the object at that location is selectable.

   By default, selected cross sections are yellow. (You can customize this color with the *Blending Envelope Display Options dialog (page 2–1125)*.)

   You can use Ctrl +click to add other cross sections to the selection set.
4. Use Radial Scale to resize the selected cross section, and the transform tools to move, rotate, and scale it.
5.  To move to the adjacent cross section within the current bound, click Next or Previous.

   Tip:  You can also use the Control Point level to individually transform cross section control points.

To copy an envelope and its settings to a mirrored link (for example, from one thigh to the other):

1. Select the source link.
2. In the Edit Commands group, click Copy.

   The envelope settings are saved in a buffer.
3. Select the mirror link, and then in the Edit Commands group, click Paste.

   The envelope is applied to the opposite link, but without mirroring. It is now the currently selected envelope.
4. In the Edit Commands group, click Mirror.

   The selection flips about the mirror axis.
5. Use the Rotate transform on the Main toolbar to rotate the selection to its final orientation.

See also

Envelope Sub-Object (page 2–1130)

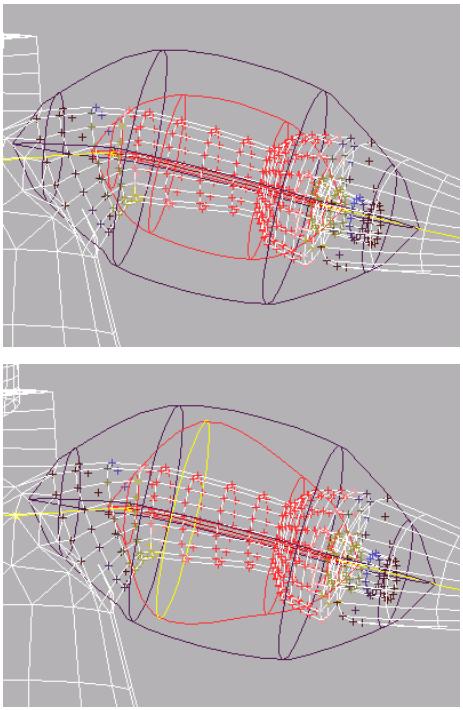
Fine-Tuning Envelopes (page 2–1088)

Fine-Tuning Envelopes

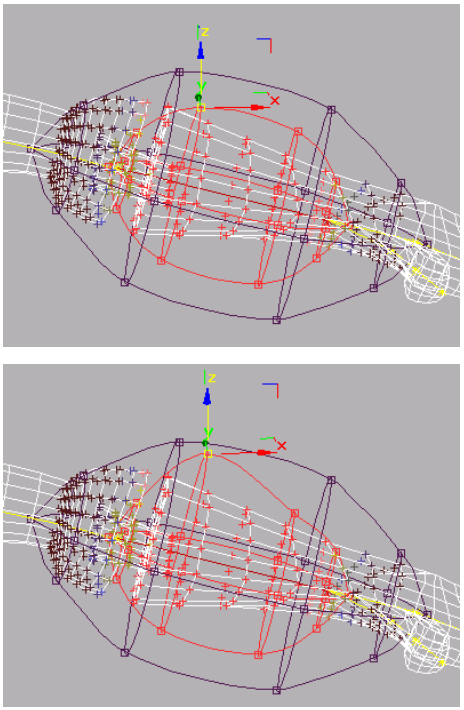
 Once you have adjusted inner and outer bounds of envelopes at the Envelope Sub-Object level, you may find you need still finer control: choose either the Cross Section or Control Point selection level in the Blending Envelopes rollout.

Using Cross Sections

Both inner and outer envelope bounds are like *hulls* about *ribs*, the cross sections. For each envelope, there are by default four cross sections you can use to alter the envelope's shape.



Envelope cross sections can be scaled and moved.



By moving a control point the envelope is also scaled.

Using Control Points

By default, each cross section consists of four control points. You can use these control points to alter the shape of a single cross section. The control points define the shape (perimeter) of the section. Moving one point affects the shape between itself and its neighbors on each side.

Rotating and Scaling

Control points are internally positioned using cylindrical coordinates, that is a distance from the link at some angle around the link. To move a control point, you must scale it to the correct distance from the link and then rotate it into the orientation you desire. The move tool actually serves the dual function of rotating and scaling at the same time.

See also

Envelope Sub-Object (page 2–1130)

Customizing Vertex Assignments

Use Vertex sub-object controls to:

- Override vertex-link assignments currently in effect.
- Change the blending type and weight.
- Assign vertices to their links yourself, bypassing Physique's envelope method for vertex-link assignment.

You might want to reassign link assignments in effect about adjacent fingers, for example, even after working to reshape envelope size and position. You may find that you need greater control in areas where adjacent envelopes overlap:

perhaps movement in the ring finger is causing unwanted movement in some vertices for the middle finger. You'd use Vertex sub-object controls to remove ring-finger vertices from the middle finger links, in effect reassigning them solely to the ring-finger links.

When you choose Vertex as the sub-object to edit in the Physique modifier, the *Physique Selection Status and Vertex-Link Assignment (page 2–1150)* rollouts appear.

Envelope Display Options

Interactive Redraw

By default Interactive Redraw, an option in Envelope and Bulge sub-object levels, is turned off. After changing an envelope the mesh is refreshed on mouse up. Turn this option on for interactive viewing of how envelopes and bulges are affecting vertices. As you adjust the envelopes, the vertices move to show the effects of the resized envelopes

Changing Display Options

You can change the display settings for envelopes and their component parts. Click the associated color sample to open the Color Selector and change color setting for:

- Inner and outer bounds on both deformable and rigid envelope types.
- Differently weighted vertices, as they fall within the inner bound, or between it and the outer bound.
- Selected cross sections and control points.
- Turn on or off the associated check box to determine whether any of the above envelope components are displayed.
- Number of sides for inner and outer bounds. The default is 4.

Choose *Initial Skeletal Pose (page 2–1099)* if you want to work in the position of the mesh and its skeleton at the time Physique was applied. Returning to this pose can be useful if the display becomes confused with the effects of unassigned vertices when the character moves through an animation.

See also

Fine-Tuning Envelopes (page 2–1088)

Working with Deformable Envelopes

Once you've got the deformable envelopes working the way you want to control overall skin deformation, you may want to adjust the finer aspects of skin control:

- Skin bending, twisting, sliding, and scaling about single links are controlled at the Link sub-object level. See *Adjusting Link Parameters (page 2–1091)*.
- Link sub-object primarily controls the *deformation spline (page 3–927)*. You use Link sub-object options to affect the shape of the spline and the smoothness of the skin. See *Adjusting Link Parameters (page 2–1091)*.
- Crease behavior where links meet and bend is controlled at the Link sub-object level with *Joint Intersection (page 2–1140)* parameters.
- Muscle bulges as your character's limbs move through a range of motion are controlled at the Bulge sub-object level. See *Creating Bulges (page 2–1094)*.
- Skin stretching, generally between non-contiguous links, is controlled at the Tendon sub-object level. Tendons provide a secondary movement (a pulling or stretching) based on links farther up or down the skeleton. See *Creating Tendons (page 2–1096)*.

See also

Link Sub-Object (page 2–1135)

Bulge Sub-Object (page 2–1141)

Tendons Sub-Object (page 2–1147)

Working with Rigid Envelopes

Vertices influenced by a rigid envelope follow the link itself rather than the deformation spline, like deformable vertices do. Low vertex-count models used for games are typically connected to the skeletons with rigid envelopes to simplify export to game engines. In such a case, each vertex can be simply described as assigned to a specific link. It can be located by giving a length along the link, a distance from the center of the link, and an angle around the link. As the link is transformed, all the vertices assigned to the vertex are uniformly transformed. This makes it simple for game engines to move the character's skin.

The vertices move equally with the controlling link, so the skin appears rigid. For instance, in the case of a flexing arm, the skin between the wrist and elbow tightly follows the movement of the forearm, with no change in the spacing between vertices. The skin between the elbow and shoulder tightly follows the upper arm. Vertices around the elbow influenced by both links, smoothly blend to an average position between the two changed links.

Rigid and deformable envelopes can be mixed on a given character. You might, for instance, use deformable envelopes for most of the character, but choose rigid for the head and for a suit of armor around the torso.

Note: In Link sub-object, you can turn on twisting for rigid envelopes. As the wrist twists the vertices on the forearm will twist also.

See also

Envelope Sub-Object (page 2–1130)

Working with Both Deformable and Rigid Envelopes

You can assign both Deformable and Rigid to a given link. For example, you might assign both to the shins, and size the envelopes differently to gain the benefits of both.

First, size the deformable envelope to affect the whole lower leg area and smoothly transition into the thigh and ankle areas. Next, you define a smaller rigid envelope to encompass the stiff unbendable section of the shin. Increase the weight of the rigid envelope to balance its effect against that of the deformable envelope.

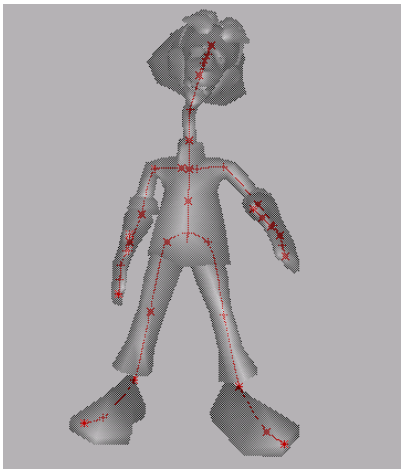
Some users might use both envelopes as a way to have a bone affect two completely separate regions of a mesh. The envelopes are scaled so that the deformable envelope influences certain vertices on one side of the link and the rigid envelope affects the other side.

See also

Envelope Sub-Object (page 2–1130)

Adjusting Link Parameters

Physique creates a continuous spline through each of the joints in the hierarchy. The spline is represented by an orange line within the object. It maintains continuity through each joint: as a joint angle changes, the spline passing through the joint remains a smooth curve. Physique relies on the deformation spline to obtain smooth bends of the skin object. Physique offers bias and tension controls to adjust the shape of the deformation spline.



Links depicted by deformable spline.

You can use controls at the *Link sub-object level (page 2–1135)* to adjust the skin behavior about a given link. These link parameters define how the skin deformation behaves, relative to the motion of the underlying skeleton.

There are four kinds of *link parameters (page 2–1136)*:

- Bend parameters affect the curvature of the deformation spline through the joint at a given link. As a result the skin deformation can range from more angular at the joint to more like a stiff rubber tube.
- Twist parameters control how the skin deforms when a joint rotates along its length, as in turning a doorknob. Consider, for example, winding up the rubber band on a toy airplane propeller; Twist parameters determine how the twist is distributed along the link and across the joint.

  Twist parameters apply to link rotation whose axis is the same as the axis along the length of the link.
- Sliding parameters control how the skin moves along the length of a link as a joint rotates. On the outside, the side where the angle is greater than 180 degrees, the skin moves toward the joint. On the inside, the side where the joint angle is less than 180 degrees, the skin moves away from the joint.

  The sliding effect tightens the outside to keep detail at the joint and prevent the facets of the mesh from moving apart. It relaxes the inside to prevent the mesh from bunching up at the joint.
- Radial Scale parameters affect the influence on the skin caused by scaling links of the skeleton. Certain settings, such as Breathe and Stretch, react to scaling of the skeleton, but only with standard 3ds Max bones or a 3ds Max hierarchy.

  Note: Bipeds and 3ds Max IK controllers don't support scaling.

  Link Scale provides a way to globally scale all the vertices influenced by the link. CS Amplitude scales the effect of bulge cross sections across the entire link. It can be used for animated scaling of the bones or for general mesh adjustments. The intent is not to edit the character itself, but rather the effect the bones have on it.

You set Link parameters at the sub-object level, in the *Physique Link Settings rollout (page 2–1136)*.

Partial Blending and Weight Assignments

Most often, you will want to use envelopes to correct the way skin behaves as the biped moves. However, you can override envelopes by manually assigning vertex properties. For example, you can remove the influence of inappropriate links from selected vertices. You can also change the weight distribution between links for a single vertex by using type-in weights.

In Cases where No Envelopes Use Partial Blending (the Default)

Vertex *v* is assigned to links *l1, l2,* and *l3*, and the weights for these links are: *w1* = 0.2, *w2* = 0.3, and *w3* = 0.4.

In a non-partial blended case, the sum of these vertex weights is *w1 + w2 + w3 = 0.9* (less than 1.0).

With partial blending off, Physique will normalize the weights so they sum to 1.0. For example:

```
w1′ = w1 / w1+w2+w3 = 0.2/0.9 = .2222222...
w2′ = w2 / w1+w2+w3 = 0.3/0.9 = .3333333...
w3′ = w3 / w1+w2+w3 = 0.4/0.9 = .4444444...
```

so

```
w1′+w2′+w3′ = 1.0
```

In Cases where All Envelopes Use Partial Blending

The weight fill-in for vertices with a weight less than 1 will always fill with the weight of the root. Whenever a vertex is assigned to a group of links, all of which are partially blended, then the remaining weight will be assigned and blended with the root link.

This works well for a case where you want partial deformation falling off to no deformation. An example would be a static head, where you are deforming the head for facial expressions, but the head itself remains in place.

For example: vertex *v* is assigned to links *l1, l2,* and *l3.* The weights for these links are: *w1* = 0.2, *w2* = 0.3, and *w3* = 0.4.

The sum of these vertex weights is *w1 + w2 + w3* = 0.9 (less than 1.0). We need an additional fill-in weight (*wf).* The fill-in weight is determined by *1.0 - (w1+w2+w3) = 0.1,* or 10%. Physique fills in with weight *wf* from the root (leaving the vertex partially undeformed).

The resulting deformation will be *w1*l1 + w2*l2 + w3*l3 + wf*root*. The root portion of this deformation is essentially an undeformed portion that simply follows the root of the skeleton.

In Cases where Some Envelopes Use Partial Blending and Some Do Not

The vertex weight fill-in in overlap areas will be based on the percentage of partial and non partial weights. For example, If the total weight of non partial links is 80% of the total summed partial and non partial weight, then 80% of the fill-in will be more of the non partial deformation. The remaining 20% fill-in will come from the root.

Example: If Vertex *v* is assigned to links *l1, l2,* and *l3,* the weights for these links are: *w1* = 0.2, *w2* = 0.3, and *w3* = 0.4.

Let's assume *l1* and *l2* are nonpartial, and *l3* is partial. The nonpartial weight is *w1 + w2* = 0.5; the partial weight is *w3* = 0.4; the nonpartial weight is 0.5/(0.5+0.4) = .555555, or 56%; and the fill-in weight is still (1.0 - .9) or *wf* = 0.1.

The program fills in with 56% of *wf* with more of the nonpartial blended links. The remaining 44% of *wf* is filled in with the root as in the partial blended case. This provides a smooth transition between the partial and nonpartial links.

Bulges

For some animations, simply attaching the skin and correcting its vertex assignments results in an animated skin you can use in final renderings. For other animations, you might need to give the skin more realistic movement, for example, muscles that bulge.

Physique lets you simulate an underlying musculature for the skin by adding *tendons (page 2–1096)* and bulges:

- Bulges change the skin's profile to simulate bulging muscles. You create the bulge by establishing bulge angles, relationships between cross-sectional slices of the skin and specific poses of the skeleton joint. Imagine a cross section to be a slice through the skin's mesh, perpendicular to the link. By making changes to cross sections, you in turn distort the shape of the mesh. Bulges in your character can be constructed by associating certain poses with related changes to the cross sections, in other words by defining *bulge angles*.

  At any joint angle, you can define a bulge angle, and you may define as many bulge angles as needed. The bulge angle consists of the current orientation of the joint together with any defined cross sections. In addition, you can adjust the influence of a bulge angle. Physique considers all the bulge angles as the character moves. The resulting bulge is created by interpolating the effects of the various bulge angles having some influence at the current joint angle.

  For example, to create a bulging biceps muscle, in Bulge sub-object level, on a selected link, insert a cross section near the center of the upper arm. Pose the arm into a flexed position, with the angle between upper and lower arms at 90 degrees or less. Insert a bulge angle and adjust the cross section so that it distorts the mesh appropriately. In the viewports and in the *Bulge Editor (page 2–1114)*, you can edit the shape of the bulge to look like a flexed biceps muscle: higher and wider above the bone than below it. Now as the elbow bends from a straight orientation up and toward the shoulder, Physique bulges the biceps appropriately.

  see *Creating Bulges (page 2–1094)* for more information about creating bulges.

Because bulges are optional, you can approach Physique animation in a couple of ways:

- Apply only as much detail as you need to get the effect you want for a particular scene in your animation. This is probably the best approach when the Physique animation is meant to be used only once, or is not the main focus of the animation.
- Define a fully deformable character, with bulges for its entire range of motion. This is probably the best approach when you intend to reuse the Physique character in an ongoing series of animations, for example, or in a video game that has a variety of character action.

See also

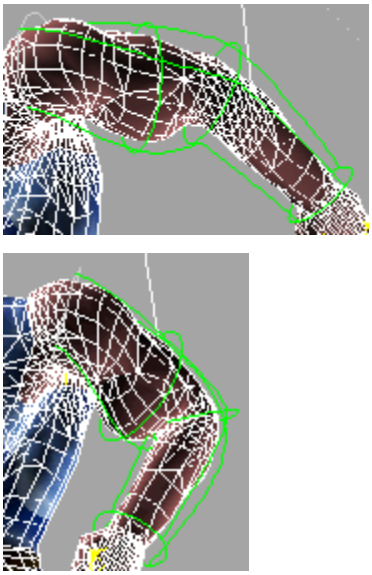
Bulge Sub-Object (page 2–1141)

Creating Bulges

Bulges simulate bulging muscles. Physique creates bulges based on bulge angles and cross section shapes you specify, not on keyframe settings. You create a bulge by:

- Reposing the character to a position where the bulge will have its greatest effect. This can just be a matter of using the time slider to scrub to that place in a loaded motion file.
- Setting a bulge angle between two links, the currently selected link and its child link in the hierarchy. The bulge angle is the angle of the joint where the bulge has its full effect. When the joint has a different angle, Physique interpolates so the bulge can grow as the joint flexes toward that angle. See *Setting Bulge Angles (page 2–1095)* for more information.

  Note: The resulting bulge for any given frame in an animation is determined by the interpolated effects of all bulge angles for the link, based on the relationship of each bulge angle to the current joint angle. Bulge angles are not directly associated with keyframe parameters, but are relative to the skeleton's behavior.



Bulge shape is interpolated as joint movement approaches a bulge angle.

- Creating and shaping the cross sections associated with the bulge angle. A link's cross sections and its profile are spline controls of the shape of the skin. To create and shape cross sections, see *Shaping the Bulge (page 2–1096)*.

  Each bulge angle affects both neighboring links. Therefore, each link contains a set of cross sections for each bulge at both its parent and child joint angles. For example, the forearm link can be deformed by bulge angles associated with both the elbow and wrist joints.

- Adjusting bulge parameters, including the joint intersection parameters. Bulge parameters control the smoothness and the strength of the bulge and are found at the Bulge sub-object level.

  Joint intersection parameters control how the skin behaves when bulges would overlap each other if there were no collision detection for skin vertices. They are in a Physique rollout for Link sub-objects. See *Reinitializing Physique Settings (page 2–1098)*.

See also

Bulge Sub-Object (page 2–1141)

Setting Bulge Angles

A bulge angle associates an angle value and a name. By default, each link has one bulge angle whose default name is the name of the link followed by "Bulge 0". The default bulge angle's initial angle value is the angle between the link and its child when you first attach Physique to the skeleton.

Bulges are effective because they grow and shrink as the joint moves. The initial bulge angle defines one shape that the skin can deform to; this would normally be like a default musculature. With no other bulge angles defined, the skin would always look like the first bulge angle, regardless of pose. Additional bulge angles provide the other shapes the mesh can deform to.

In the simple case of a flexing biceps muscle, one bulge angle defines the relaxed position and another defines the muscle in its flexed pose. Both are exposed so you can add definition to the default character. As in a biceps muscle, a heavily-muscled character might have some shape even in a relaxed pose. This lets you change the baseline without disturbing the original mesh. You can actually use Physique bulge cross sections to model your character.

The cross section deformation for a given link is determined by interpolating between the contributions of all bulge angles that affect the link. This includes bulge angles for both the link's parent joint angle and its own (child) joint angle. Each bulge angle's contribution is determined by:

- Influence (how far away the bulge angle is from the current joint angle rotation)
- Power (an ease based on the influence of the angle)
- Weight (the relative strength of the bulge angle)

Keep in mind that bulge angles can be set for arbitrary rotations, and are not limited to single-axis hinge joints.

See also

Bulge Sub-Object (page 2–1141)

Fine-Tuning Bulges

The *Bulge Editor (page 2–1114)* duplicates many of the controls available at the Bulge sub-object level. It gives you a focused, two-dimensional view of the current bulge settings.

Shaping the Bulge

Once you've inserted the bulge angle, and recorded the angle setting, you then shape the bulge by inserting and adjusting cross sections. You can either move and scale the cross sections, or adjust each one's shape using control points.

Adding More Poses to Your Character

For more complex body movements, you may need to add bulges for different positions. Overall, you might want to build a generic character that works in a number of situations, regardless of the type of motion.

Some examples are:

- Moving the character into several extreme poses.
- Putting the arms high over the head and bending them into different positions.
- Making the character squat and lifting the legs into various positions.
- Adding bulge angles to the links at places where real muscles would bulge.

Once the skin reacts well to the full range of motions, you can then place the character into a less demanding pose and it should work perfectly.

Tendons

Tendons tie links together, extending the effect of moving one link to another link where the tendon is based. Their effect is similar to that of tendons in a body.

For example, raising an arm usually stretches the skin along the same side of the body. To get this effect using Physique, you could base a tendon on a spine link, then attach it to the upper arm or collarbone, so when the arm lifts, the skin around the torso stretches as well.

see *Creating Tendons (page 2–1096)* for more information about creating tendons.

See also

Tendons Sub-Object (page 2–1147)

Creating Tendons

Like tendons in an actual body, *tendons in Physique (page 2–1147)* link one bone to another. They spread the effect of moving one link to the skin around a different link. Tendons can improve the realism of skin movement when it is animated by Physique.

Basically, a tendon consists of base points that live on a cross section near the skin. A base point is attached to another link on the skeleton that pulls it. As the tendon base point is pulled, it deforms the skin around it. Several parameters control how much the point is pulled in each direction.

You insert the tendon cross section at the location where you want the skin to stretch. You then attach it to the link that influences the movement.

Each link can contain several tendon cross sections, and each control point may be attached to a separate link. In a practical application, however, such as in the area below the armpits, you

might have two control points attached to each of the left and right clavicles.

Procedures

To create and attach a tendon:

1.  Go to the Tendons sub-object level. In the Selection Level group of the Tendons rollout, turn on Link.
2. In a viewport, select a link.
3. In the Insert Settings group, click to turn on Insert.
4. In a viewport, position the cursor over the link.

   The cursor changes to a small star.
5. Click to create cross sections on the link.
6. Click to turn off Insert.

   Tip: You can also right-click in a viewport to turn off tendon creation.
7.  In the Selection Level group, turn on Cross Section.
8.  Use the Rotate tool to rotate the cross section so the control points are positioned in useful places.
9. In the Tendon Parameters group, use the Radius spinner to scale the cross section radially so the control points fall close to the surface of the skin.
10.  In the Selection Level group, turn on Control Point, then in a viewport, select one or more control points.
11. In the Edit Commands group, turn on Attach, then in a viewport, click a different link.

    Tendons span the area between the links.
12. Adjust the Pinch, Pull, and Stretch settings to adjust the skin behavior.
13. Adjust the Upper and Lower Boundary settings to control the extent of the effect on neighboring links.

To attach a tendon to another link:

1.  In the Selection Level group of the Tendons rollout, turn on Control Point to make it the active selection level.
2. In a viewport, select one or more of the tendon's attach points.

   Click a point to select it; use Ctrl+click to add points to the selection; or drag a rectangular region to select multiple points.
3. In the Edit Commands group, turn on Attach.
4. In a viewport, click the link to which you want to attach the tendon.

   In the viewports, red lines appear showing the tendon's connection from its base to the other link. The skin mesh might also deform as a result of attaching the tendon.

   Tip: The tendon can have fixed attach points that are not connected to another link. These are useful for giving some rigidity to the skin, as when (in an actual body) a bone lies close to the skin's surface. For example, you might leave two fixed attach points on either side of a character's chest area, to simulate the effect of the sternum. When all tendons are attached to other links, the skin over the base link can have a "squishy" appearance when it is animated. This is appropriate for some animated characters, but not for others.

To delete a tendon:

1.  At the Link selection level in the Tendons sub-object, select the link that has the tendon you want to remove.
2. In the Insert Settings group, click Delete.

Saving and Loading Physique Data

You can save Physique data to a Physique (*.phy*) file to save data common to all objects that share a given Physique modifier

Later, you can reload the data file, either to restore the data that belongs to a particular skin or portion of skin, or to transfer the Physique of one skin (or portion) to a different one.

See also

Physique Rollout (page 2–1106)

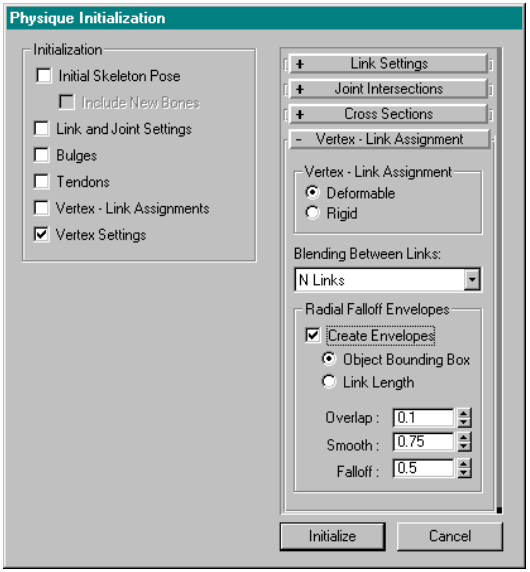
Physique Load Specification Dialog (page 2–1123)

Reinitializing Physique Settings

When you need to reset vertex, envelope, and other skin parameters on a Physique mesh, click Reinitialize to display the Physique Initialization dialog. Using controls in this dialog, select the category to update, and apply the new global settings.

For example, if you've added a new bone to the hierarchy and want it included and influenced by the Physique modifier, use the reinitialization mechanism to effect its inclusion. Or maybe you've repositioned the biped structure relative to the mesh, you'd need to reinitialize Physique settings to recognize those changes.

Reinitializing with changed settings applies them as new defaults in the areas you choose. Reinitializing without changing settings on the Physique Initialization dialog is a method of *erasing* unsatisfactory changes you've made to the Physique settings, again in those areas you specify.



Reinitialization mode of Initialization dialog

Initialization Group

The Initialization group in the Physique Initialization dialog is where you name the category of settings to reset. (When you applied Physique to the mesh initially, these settings appeared unavailable; in fact, all settings are set at Initialization.)

Warning: Select only the check boxes for those settings you need to reset. If you select everything, your model is returned to the state it was in immediately after applying Physique. That's a quick way to discard changes that aren't working; just make sure that's what you want to do. If you change settings in rollouts on this dialog and fail to select the corresponding check box, those rollout settings are ignored and initial defaults are reestablished.

See also

Reinitialize Physique (page 2–1113)

Working with an Initial Pose

At times, you'll need to alter the fit of the hierarchy in the Physique mesh, or change its structure. To do so, you change the default, initial pose, which Physique uses as a reference for various operations, including reinitializing.

- For a biped, you change the structure in Figure mode, then reinitialize the Initial Skeleton Pose.
- For a Bones system with the HD IK solver, there is a Show Initial State check box in the Motion panel, on the IK Controller Parameters rollout.
- For any other hierarchy, use frame 0 for the initial pose. Be sure your animation does not affect the position of the skeleton in frame 0.

Procedure

To load a file created with different system units:

Note: If the System Unit Scale has a value that is different from that of the file you are loading, 3ds Max prompts you to rescale the scene. If you do so, any objects that have a Physique modifier exhibit a double scaling. Do the following:

1. Select the object.
2. Reinitialize with both Initial Skeletal Pose and Vertex Settings (the last check box) turned on.

See also

Figure Mode (page 2–982)

Improving Interactive Performance

Physique is multithreaded and optimized for modifier stack changes below the Physique modifier.

Level of Detail Controls

The controls in the *Physique Level of Detail rollout (page 2–1108)* help you optimize performance while working with Physique, by letting you specify which skin deformations are refreshed automatically in viewports. The more complex the skin object, the more effective these controls can be at speeding up your work.

Note: The 3ds Max renderer also heeds Level of Detail settings. Don't forget to reset these controls before the final render.

Turning Blending On and Off

When Physique is first initialized, the Blending Between Links parameter in the *Physique Initialization dialog (page 2–1111)* is set to *N Links (page 3–978)*. This means that every envelope must be considered when determining the influence on any *vertex (page 2–1123)*. You can reinitialize with No Blending (or only 2, 3, or 4 links) to reduce this calculation permanently, although at the cost of losing some or all blending at the joints.

Note: Setting to No Blending is advised only for those developing characters to be used by real-time game engines.

You can also set No Blending at the Vertex sub-object level. You can temporarily disable Link Blending on the Physique Level of Detail rollout, on the Modify panel.

Using the Optimize Modifier with Physique

The standard 3ds Max Optimize modifier allows you to reduce the level of mesh detail in order to increase display performance and speed up screen refreshes. Optimize can be helpful with Physique when you work with complicated skin objects such as detailed biped figures.

There are stack update options that determine how the stack is updated. See the *Physique Level of*

Detail rollout (page 2–1108) for a reference on the stack update options.

Warning: If the current stack update option is Reassign Globally, any modifier in the stack below Physique causes Physique to reevaluate vertex assignments each time the stack is changed. While this technique will improve performance in one way, it will hurt performance if you need repeatedly to make changes down in the stack.

Procedure

To optimize skin objects:

1. Select the skin.

   If the skin consists of multiple objects, select all of its component objects.

2. Apply an Optimize modifier to the skin.

   The Optimize modifier's Parameters rollout appears on the Modify panel.

3. In the Optimize group of the Parameters rollout, increase the value of Face Threshold.

   As you increase Face Threshold, Optimize reduces the number of faces in the mesh.

The optimized mesh displays more quickly and speeds up your work with Physique. However, it doesn't show full detail, and sometimes an optimized skin does not animate correctly, because of the reduced number of vertices.

Tip: Always turn off Optimize before you render the animated skin.

Scaling a Character

Using standard 3ds Max scale transform tools, you can adjust a biped's posture by scaling the size of its links. You must be in Figure mode to scale the biped links. If you try to scale a biped without going into Figure mode, nothing happens.

Procedures

To scale a biped that has a mesh attached to it by using Physique:

1. Select the biped.
2.  On the Motion panel, on the Biped rollout, turn on Figure mode.
3. On ths Structure rollout, change the biped's height.

   The biped and mesh scale together.

To reinitialize a scaled mesh:

Reinitializing a scaled mesh might be necessary if you use Initial Skeleton Pose in Physique. In this case, the mesh appears at its size before scaling. To correct for this, perform the following steps:

1. On the Modify panel, either make Physique inactive, or highlight the skin so the active level is below the Physique modifier in the stack.
2. Scale and move the mesh until it matches the biped or bones.
3. Re-enable Physique, or highlight its name once more in the stack display.

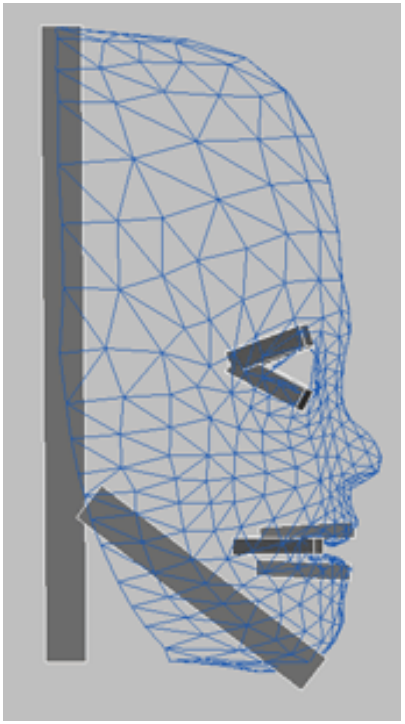
   The mesh will get large because it is doubly scaled.

4. Reinitialize with Initial Skeleton Pose turned on.

   The mesh will shrink to the size you set during scaling.

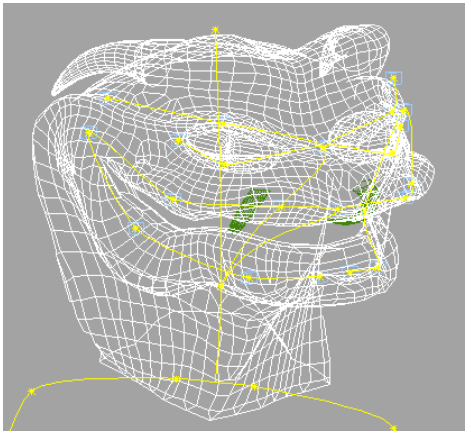
Facial Animation

As an alternative to morphing, Physique can be used effectively for detailed facial animation when a character's face is set up with extra links. By setting up 3ds Max bones or other objects at appropriate locations on the mesh head, you can define a skeletal structure for moving the facial features.



Boxes form the facial skeletal structure.

You can also use splines linked to the biped head to deform the mesh. You can add a spline to a character that already has Physique applied by using Add in the Physique Bones rollout and clicking the spline in a viewport.



Splines used for detailed facial animation.

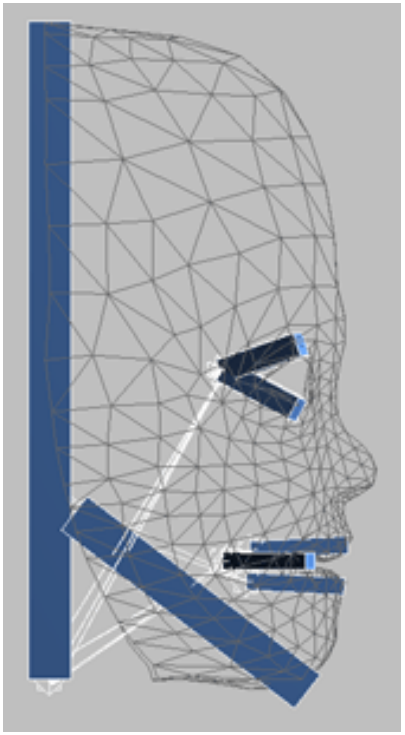
Facial Bone Structure

A common skeletal structure for the face includes:

- Bones to make the upper and lower lips open and close.
- Bones to push the outer corners of the mouth, as for a smile.
- Bones to make the eyes open and close.
- A bone to control the jaw opening and closing.
- The main head bone going straight up from the neck to the top of the head. This bone controls the head overall, and any vertices not affected by bones on the eyes, mouth and jaw. Use this bone to bend the head forward or backward.

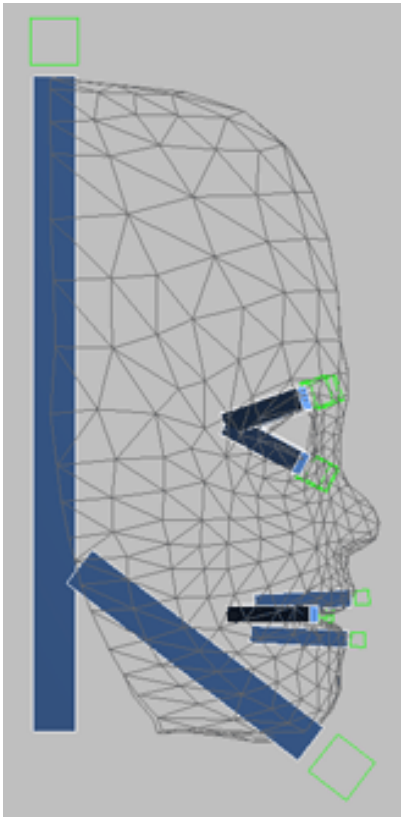
Tip: If you've never set up a facial bone structure before, refer to a text on character animation for information on the facial bones needed to create various expressions.

In general, the lower lip bones should be linked to the jaw, and upper lip, jaw and eye bones should be linked to the main head bone.



Facial bone linkage

If you use objects other than 3ds Max bones, you must create dummy objects at the end of each object used as a bone, and link each one to its corresponding bone before applying Physique. This is necessary because Physique extends links to the end of the second-to-last object in the hierarchy, not the last. In other words, if you don't create and link the dummy objects to the bones, Physique will not extend links to the ends of the bones.



Dummy objects at ends of bones

Mesh objects that are to be animated along with the face, such as the eyes and hair, should not be linked to the bone structure until after Physique is applied. If you link them beforehand, Physique will treat them like bones. You can link these objects to the main head bone after you apply Physique.

Separate or Connected Head

There a two ways a setup like this could be used with a character body: The head could be attached to the body mesh, or it could be separate.

If possible, use a head that is separate from the body mesh. For human characters, this often works fine; most character animation doesn't provide a view of the underside of the neck, so

your viewers will most likely never notice that they're separated. If you use a separate head, it will be easier to control the facial animation.

If you use a separate head, apply Physique to the head mesh alone. Click Attach to Node and pick the main head bone. Check the links to make sure they extend to the ends of the bones. Then, link the main head bone to the biped head.

If the character's head is connected to the mesh, link the main head bone to the biped head before applying Physique. Alternately, you can use the biped's head as the main head bone, linking the facial bones directly to the biped's head. When you apply Physique and click Attach to Node, pick the biped's COM as usual.

After applying Physique to the head or the entire structure, you can adjust envelopes and vertex assignments for the facial bones as you would for any biped bone.

Animating the Facial Structure

After you apply Physique to the structure, you can animate it by rotating or moving bones. 3ds Max *manipulators (page 2–27)* are useful for setting up a custom user interface for bone animation.





Facial expressions animated with Physique and facial bones.

Tip: Hierarchies other than bipeds can't use Figure mode, so special consideration is needed to establish the initial skeletal pose. Position the facial bones in the "at rest" position at frame 0, and start keyframing the face at frame 1 or later.

You can perform lip synch animation by loading a sound track in Track View. By scrubbing the time slider, you can locate a sound and keyframe the dummies to appropriate positions.

Procedure

To isolate lip vertices from influence by inappropriate links:

When working with a complicated facial bone structure, envelopes for the lower lip are bound to affect vertices in the upper lip, and vice versa. In general, this can be corrected in at the Vertex sub-object level:

1. Select the head and on the Modify panel go to the Vertex sub-object level.
2. Turn on Initial Skeletal Pose
3. Turn on Select, region-select vertices of the lower lip, then click Remove From Link.

4. Select the links of the upper lip, and click Lock Assignments.
5. Region-select vertices of the upper lip, and click Remove From Link.
6. Select links of the lower lip, and click Lock Assignments.

Physique and Modifiers

Combining Physique with Other Modifiers

Although you can apply Physique to an object with modifiers, this can affect performance. If you use other animated modifiers in combination with Physique, this reduction is unavoidable. If the other modifiers are not animated, you can collapse the stack to remove modifiers prior to Physique. This can significantly improve performance.

How Physique reevaluates vertices coming up the stack depends on the stack update options.

Stack Update Options

In Physique, there are three stack update options available from the *Physique Level of Detail rollout (page 2–1108)* to determine how Physique handles animated vertices coming up the modifier stack.

Add Change

Add Change adds in the changes from the stack and then applies mesh deformation. No vertex remapping or reassigning is done. This is the default option, and usually will give you the deformation that you want. There is no performance degradation from Physique.

Remap Locally

Remap Locally resets vertex position on the spline used for bending and the link position used to interpolate twist. Use this option when vertices are sliding along the length of the spline, and you want them to bend and twist based on the spline position but don't want the weights to change.

Reassign Globally

Reassign Globally does a complete vertex reassign for each frame. Use this option when vertices are moving to different envelopes and you want them reassigned to the new envelopes. This option should rarely be used. The only reason would be if vertices are sliding along a link and you want the twist to be interpolated based on the new position.

Using Physique with Changing Geometry

Physique can effectively handle dynamically changing geometry. That is, if the geometry coming through the modifier stack to Physique changes, Physique dynamically adjusts to accommodate the changes to the geometry, while maintaining its own parameter settings, including manual vertex assignments. This is why, when using Optimize, you can effectively change the resolution of the geometry deformed by Physique.

Only in situations where the changes dramatically increase the density of data is there ever a need to manually adjust vertex assignments. Physique remembers your manual vertex assignments, and uses these assignments to reassign vertices when the geometry changes.

Physique and Free Form Deformations (FFDs)

Physique can be applied to a Free-Form Deformation (FFD) space warp, which in turn can animate a mesh that is bound to the FFD. For

example, you could use this technique to animate a credit card or a box of cereal.

Procedures

To apply Physique to an FFD to animate the entire mesh:

1. Place an FFD (Box) space warp around the mesh to deform. The box should be large enough to encompass the mesh. The number of control points you use for the FFD can be fine-tuned later by going back down in the stack to the FFD to adjust the number of control points.
2. Select the FFD and apply the Physique modifier.
3.  Click Attach To Node, and then click the biped pelvis or the root node of your bone structure.
4. If necessary, select the FFD space warp and adjust the vertex assignments of the control points so they are assigned to the proper links.

   This can be done in the same way as done using Physique with a mesh, only there are fewer assignments to deal with. Although fewer assignments provide smoother surface deformation with the FFD, control points and their link assignments must be thoughtfully placed.
5.  Use Bind To Space Warp on the main toolbar to bind the mesh to the FFD space warp.
6. Link the mesh to the biped pelvis, or root node of the bone structure, so it follows the skeleton and FFD as they move around the scene.

To use an FFD to complement the effects of Physique on a portion of a character mesh:

Tip: You can use this procedure to animate clothes or amorphous shapes.

1. Place an FFD (Box) space warp around the mesh to deform. The box should be large enough to encompass the mesh. The number of control points you use for the FFD can be fine-tuned later by going back down in the stack to the FFD to adjust the number of control points.
2. Select the FFD and the mesh and apply a Physique modifier to both.
3.  Click Attach To Node and in a viewport, click the biped pelvis.
4. Select the mesh only, and add a Mesh Select modifier to the mesh *above* Physique in the modifier stack.

   Note: If you are using a NURBS model, use the NSurf Sel modifier (rather than Mesh Select) to select sub-object control vertices that lay inside the FFD lattice.
5. Go to the Vertex sub-object level (or Control Vertex), and select the set of vertices (or CVs) that lie inside the FFD space warp you are using.
6. While in Vertex sub-object level, bind the mesh to the FFD space warp. Now only the selected vertices or CVs will be affected by the FFD.
7. With the mesh selected, highlight the Physique level in the stack, and assign all vertices or CVs affected by the space warp as blue (rigid). These should be the same vertices or CVs selected in step 5, above.
8. Assign the rest of the vertices or CVs that fall outside the FFD as deformable, in the same way you normally *assign vertices to links (page 2–1150)* with Physique.
9. Select the FFD space warp. If necessary, adjust the vertex assignments of the control points so they are assigned to the proper links.

   You do this in the same way as you do when using Physique with a mesh, except there are fewer assignments to deal with. Although

fewer assignments provide smoother surface deformation with the FFD, control points and their link assignments must be thoughtfully placed.

Physique User Interface

Physique User Interface

Select a mesh. > Modify panel > Modifier List > Physique.

Select a mesh that has the Physique modifier applied to it. > Modify panel > Floating Bones, Physique, and Physique Level of Detail rollouts

Physique is a modifier. Controls for Physique appear on the Modify panel. The Physique and Floating Bones rollouts are for attaching the mesh to the biped, splines, or bones. Controls in the Physique Level Of Detail rollout are for troubleshooting envelopes, bulges, and tendons. The sub-object controls are for fine-tuning envelopes, creating and adjusting tendons and bulges, and for vertex editing.

See also

Physique (page 3–994)

Physique (page 2–1076)

Physique Rollout

Select an object that has the Physique modifier applied to it. > Modify panel > Physique rollout

You use the buttons on the Physique rollout to link a mesh to a biped, a bones hierarchy, or a spline; to reinitialize the Physique parameters on a mesh; to open the *Bulge Editor (page 2–1114)*; and to load or save Physique (PHY) files.

Procedures

To attach a mesh to a biped using Physique:

1. Select or create a biped.
2.  Go to the Motion panel. On the Biped rollout, click to turn on Figure mode.
3.  Position the biped inside your character mesh.
4.  Select the mesh.

   Important: If the mesh is made of multiple objects, select all of them.
5.  Go to the Modify panel. Use the Modifier List to apply Physique to the mesh.
6.  On the Physique rollout, click to turn on Attach To Node. In a viewport, select the biped's pelvis object.

   The *Physique Initialization dialog (page 2–1111)* is displayed.
7. Click Initialize, accepting the default settings.

   By default, the Object Bounding Box option is selected: Physique sizes envelopes to approximate the biped limbs.
8. Experiment with animating the biped, and adjust envelopes around problem areas. Use the Envelope sub-object level to edit the newly created envelopes.

To attach a mesh to a bones hierarchy using Physique:

1. Select or create a bones hierarchy.
2.  Position the bones hierarchy inside the mesh.

   Tip: When you use a bones hierarchy with Physique, use frame 0 as your "figure mode."

Don't include frame 0 in your animation. Use it as the a place where you position the bones and fit them to the mesh.

3.  Select the mesh.
4.  Go to the Modify panel. Use the Modifier List to apply Physique to the mesh.
5.  Click to turn on Attach To Node. In a viewport, select the root bone in the bones hierarchy.

   The Physique Initialization dialog displays.
6. Click Initialize, accepting the default settings.
7. Experiment with animating the bones, and adjust envelopes around problem areas. Use the Envelope sub-object level to edit the newly created envelopes.

To add a bone after you've already used Attach To Node:

1. Link the bone to the biped.then click Reinitialize.
2.  Select the mesh, then go to the Modify panel. On the Physique rollout, click Reinitialize.

   The Physique Initialization dialog is displayed.
3. On the dialog, in the Initialize group at the upper left, click to turn on Initial Skeleton Pose. The toggle Include New Bone becomes available. Click to turn on Include New Bone, and then click Initialize.

To save Physique data:

1.  With Physique active on the Modify panel, click Save Physique File on the Physique rollout.

   A file save dialog is displayed.
2. Enter a name for the new Physique file, and then click OK.

To load Physique data:

1.  With Physique active on the Modify panel, click Open Physique File on the Physique rollout.

   A file open dialog is displayed.
2. Choose the Physique (PHY) file to open, and then click OK.

   A Physique Load Specification dialog is displayed.
3. Use the dialog to choose the kind of data to load, and then use the lists to match links in the PHY file with links in the scene.

   Important: The number of links selected from the file (left column) *must* match the number of links in the active Physique modifier (right column).
4. Click OK.

   Physique updates the links you chose with the data that was saved.

Interface



 **Attach to Node**—Attaches the mesh objects to the biped or to a bones hierarchy.

Important: Put the biped in Figure mode and fit it to the mesh character before you click Attach To Node.

Turn on Attach To Node, then in a viewport, click the biped's pelvis or the root node of a bones hierarchy.

Note: If you click a biped's center of mass instead of its vertex, Physique corrects this by attaching to the pelvis instead.

After you select a node in the viewports, the *Physique Initialization dialog (page 2–1111)* is displayed. Accept the default values in this dialog, then click Initialize. Physique traverses the biped or bone hierarchy, and creates an *envelope (page 2–1130)* for each link it finds. (It can include additional bones that are linked to the biped.) The envelopes control Physique's influence on the vertices of the mesh. When the hierarchy is animated, the mesh vertices move along with it.


Reinitialize—Displays the Physique Initialization dialog and resets any or all of the Physique attributes to the default values. For example, *reinitializing (page 2–1113)* with the Vertex Settings option selected reestablishes the relationship of a vertex and its original position relative to the Physique deformation spline. Settings for vertex-link assignments, bulges, and tendons can be reset from this dialog.


Bulge Editor—Displays the *Bulge Editor (page 2–1114)*, which is a graphical alternative to the Bulge sub-object level for creating and editing bulge angles.


Open Physique File—Loads a saved Physique (*.phy*) file, which stores envelope, bulge angle, link, tendon, and vertex settings.

When you click Open Physique File to load a *.phy* file, the *Physique Load Specification dialog (page 2–1123)* is displayed. Select the links you want to import in the list on the left of the dialog, and apply them to the links you select in the list on the right of the dialog.

Any aspect of your Physique work can be loaded into any other character. You can choose from Link Settings, Bulges, Tendons, and Envelopes. You can also apply links from the file to differently named links in the current scene. For example, save and reload a *.phy* file to copy a bulge angle created on one character to a link on another character.

Tip: A bulge created on a bones hierarchy can be loaded and applied to Physique links created using a biped.

Note: Loading a *.phy* file will not overwrite locked vertices. If you want to overwrite them, you must either go to the Vertex sub-object level and unlock them, or reinitialize and turn on the option Vertex - Link Assignments.


Save Physique File—Saves a Physique (*.phy*) file, which contains envelope, bulge angle, link, and tendon settings.


About Physique—Displays a dialog that shows version and copyright information about Physique.

Physique Level of Detail Rollout

Select an object that has the Physique modifier applied to it. > Modify panel > Physique Level of Detail rollout

Controls in the Physique Level Of Detail rollout not only optimize the viewports, but also affect the rendered result. The primary purpose of this rollout is for troubleshooting. After creating bulge angles and tendons, you can turn off their influence to see exactly what they add to the deformation of the skin. This rollout also has controls for how changes in the modifier stack below Physique are handled.

Tip: Choose Deformable and turn on all its parameters during editing; this allows you to spot problem areas. Choose Rigid for the fastest viewport redraw.

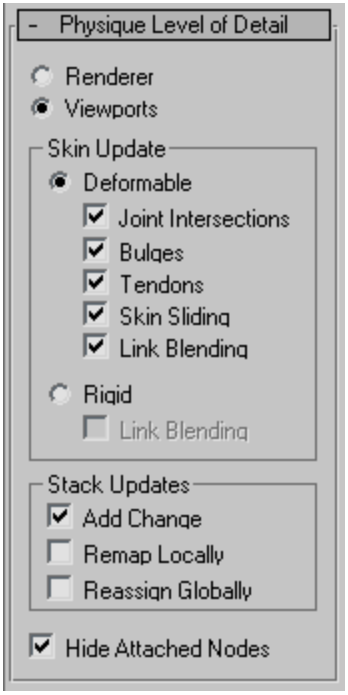
Procedure

To troubleshoot bulges and tendons:

1. Select the mesh of a character that has bulges or tendons.
2.  Go to the Modify panel.
3. On the Physique Level Of Detaul rollout, in the Skin Update group, make sure Deformable is chosen.
4. Turn Bulges off, then on. Do the same with Tendons.

   Examine the mesh with and without the influence of bulges and tendons.
5. Use the Bulge Editor or the sub-object levels to adjust bulge and tendon settings.

Interface



Renderer—When chosen, settings in the Skin Update group affect rendered images.

Viewports—When chosen, settings in the Skin Update gropu affect viewports.

Skin Update group

Deformable—When chosen, Physique deformation is active. Deformable yields the highest-quality rendering. The Deformable toggles are unavailable unless Deformable is chosen.

- **Joint Intersections—**Turn off to remove joint intersection influence. Joint intersection influence allows the mesh to overlap itself; for example, at the elbow and knee joints. Default=on.
- **Bulges—**Turn off to remove any bulge angle influence. Default=on.
- **Tendons—**Turn off to remove any tendon influence. Default=on.
- **Skin Sliding—**Turn off to remove skin sliding influence. Default=on.
- **Link Blending—**Turn off to remove the influence of link blending. Default=on.

Rigid—When chosen, forces all vertices to use Rigid assignments rather than Deformable. This is an easy way to isolate deformation problems. It also provides the quickest viewport redraw speed. You might choose this option while adjusting the animation of your skeleton.

- **Link Blending—**Turn off to remove the influence of blending on rigid links. This toggle is unavailable unless Rigid is chosen. Default=on.

Stack Updates group

The controls in this group handle changes to vertex count that arise from (nonanimated) changes to the modifier stack.

If the vertex count changes, vertices are reassigned globally. Add Change adds in changes based on the vertices' initial position. The other options reset the initial position at each frame to do the

remapping and reassigning. For this reason, turning on Add Change or making nonanimated stack changes should always be done at the initial position (Figure mode or frame 0).

Add Change—Adds in changes from the stack and then applies Physique deformation. No vertex remapping or reassigning is performed. Default=on.

This option will generally give you the deformation you want. There is no performance penalty (from Physique) when this option is used.

Remap Locally—For deformable vertices, this resets the vertex position on the Physique deformation spline used for bending, as well as the link position used to interpolate twist. For rigid vertices, this option resets the link position used to interpolate twist. Default=off.

When vertices are sliding along the length of the spline and you want them to bend and twist based on the spline position, but don't want vertex weights to change, turn this option on.

Reassign Globally—Re-weights, and resets the position on the spline used for bending for moved vertices globally. The vertex link assignment, weighting, and spline position are reset for all moving points on every frame. (This is equivalent to Physique 2.2.) This option is like reinitializing on every frame. Default=off.

When vertices are moving to different envelopes and you want them reassigned to the new envelopes, use this option.

Hide Attached Nodes—Toggles the display of the underlying skeletal system. This allows you to hide and unhide the biped, for example.

Floating Bones Rollout

Select an object that has the Physique modifier applied to it. > Modify panel > Floating Bones rollout

"Floating" bones provide a way to use Physique without using a biped character. The Floating Bones rollout specifies the splines, bones hierarchy, or unattached bones you are using to deform a mesh. For example, Physique lets you animate a mesh by animating spline vertices or a bones system. When you attach the spline or parent bone to Physique, it creates envelopes for the selected spline or bones.

Procedure

To have a spline influence a mesh:

1. Create a spline and place it within a mesh that has Physique applied.
2.  Select the mesh, and go to the Modify panel.
3. On the Floating Bones rollout, click Add.

   A Select Bones dialog is displayed.
4. Use the Select Bones dialog to select the spline.
5. Animate the spline vertices to animate the mesh.

Tip: If you want to control the mesh by animating bones, follow the previous steps using bones instead of a spline.

Interface



Add—Displays a Select Bones dialog so you can select splines or bones to use with Physique.

Reset—Resets splines or bones in their initial position: Reset reassigns vertices, but leaves envelopes as is.

Delete—To delete a spline or bone, click to highlight its name in the list, and then click Delete.

Physique Shortcuts

 Turn on the *Keyboard Shortcut Override toggle (page 3–872)* to enable the **character studio** keyboard shortcuts.

All **character studio** keyboard shortcuts activate when the Motion panel is active and the Keyboard Shortcut Override button is active.

See Also

Keyboard Shortcuts (page 3–871)

Keyboard Panel (page 3–793)

Customize User Interface Dialog (page 3–792)

Action	Shortcut	Description
Copy Envelope	Ctrl+C	
Delete	Ctrl+D	
Next	PageDown	
Next Selection Level	I	
Paste Envelope	Ctrl+V	
Previous	Page Up	
Previous Selection Level	Shift+I	
Reset Envelopes	Ctrl+E	Resets envelopes for selected links to their default values.

* Not available for customization in the Customize User Interface dialog.

Physique Dialogs

Physique Initialization Dialog

Select an object that has the Physique modifier applied to it. > Modify panel > Physique rollout > Click Attach to Node. > Select a root object in a viewport. > Physique Initialization dialog

Select an object that has the Physique modifier applied to it. > Modify panel > Physique rollout > Click Reinitialize. > Physique Initialization dialog.

Use the Physique Initialization dialog to specify link parameters and the type and size of envelopes to create for Physique links. This dialog is displayed under two circumstances:

-  When you apply Physique initially, using Attach To Node on the *Physique rollout (page 2–1106)* on the Modify panel.
-  When you click Reinitialize on the Physique rollout.

Creating Physique Links and Envelopes for the First Time

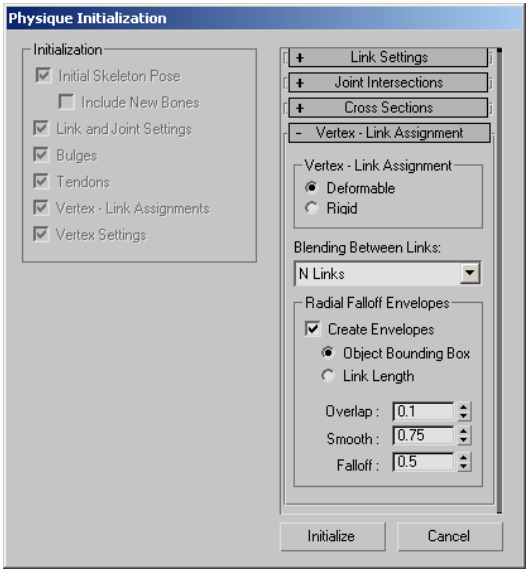
When you use Attach To Node to attach the mesh to a skeleton for the first time, you can use parameters on the *Vertex-Link Assignment rollout (page 2–1123)* to specify envelope parameters. In most cases, the default parameters on the Vertex-Link Assignment rollout work well as a starting point for envelope type, envelope sizing, and blending between links.

Normally, you want deformable envelopes. Blending Between Links is set to N Links (a vertex can be influenced by all envelopes that overlap it). The Object Bounding Box option is chosen by default, and bases the size of the envelopes on the size of biped limbs or bone diameters.

Link Settings, Joint Intersections, and Cross Section Rollouts

Controls on these rollouts match those found in *bulge sub-object (page 2–1141)* and *link sub-object (page 2–1135)* and are used here to set default values. Unless you want to change the default settings for *all* link parameters and bulges, you do not need to change the settings in these rollouts.

Interface



Initialization group

The Initialization check boxes are normally unavailable and not changeable during initialization. Although these options are unavailable, they are all turned on by default (except Include New Bones) for initialization, since *all* these settings are new when you use Attach To Node for the first time.

These check boxes become available when you click Reinitialize on the Physique rollout.

See *Reinitialize Physique (page 2–1113)* for a description of the controls in this group.

Link Settings rollout

The Link Settings rollout contains the default link values that will be assigned to all links. See *Link Sub-Object (page 2–1135)* for a command reference.

Joint Intersections rollout

The Joint Intersections rollout contains default joint intersection values that will be assigned to all links. See *Link Sub-Object (page 2–1135)* for a command reference.

Cross Sections rollout

The *Cross Sections rollout (page 2–1125)* sets the default Bulge Angle parameter settings for new bulges.

Vertex-Link Assignment Rollout

On the *Vertex-Link Assignment rollout (page 2–1123)* you can choose to create new envelopes and pick the type of blending between links.

Reinitialize Physique

Select an object that has the Physique modifier applied to it. > Modify panel > Physique rollout > Click Reinitialize.

The Reinitialize button on the Physique rollout is used mainly for resetting a particular attribute of your character's skin behavior. Among common uses are resetting bulges, deleting tendons, reassigning vertex-link assignments, or repositioning the biped relative to the attached mesh. Essentially, you can reinitialize to reset any single item in the Initialization group, or all items.

The toggles in the Initialization group of the Physique Initialization dialog are available only when you click Reinitialize.

Reinitializing Old Files

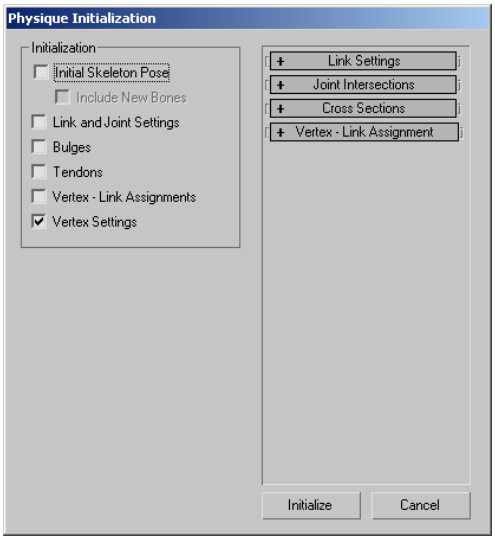
Files from versions of **character studio** prior to **character studio** go through a semi-automatic Physique reinitialization process. When you load an old file that uses Physique, you're prompted to put the biped or bones system into Figure mode or in the initial position, and then reinitialize.

To reinitialize, select the mesh, open the Modify panel, and click Reinitialize. A dialog warns you to put the skeleton into the initial position. Click OK to reinitialize, and save the file after you have done so.

See also

Physique Initialization Dialog (page 2–1111)

Interface



Initialization group

Initial Skeleton Pose—Uses the current hierarchy position (for example, the biped pose) as the default pose for defining how the skeleton fits inside the skin. The current pose becomes the new initial pose, replacing the pose used when you first attached the mesh to the biped. Default=off.

Note: This toggle does not reassign vertices that you reassigned manually unless you also turn on Vertex-Link Assignments.

When using Physique with a biped, turn on Initial Skeleton Pose if you want to use Figure mode to reposition the biped relative to the mesh; for example, to reposition the biped shoulder joints.

It is helpful first to turn off the Physique modifier, using the lightbulb icon in the modifier stack display. This lets you use Figure mode to scale and rotate the biped limbs independent of the mesh skin. Use this option after you scale a character (see *Scaling a Character (page 2–1100)*).

- **Include New Bones—**Creates new links for any new object (bone) linked to the biped. Use this to add a bone to a mesh that has the Physique modifier applied. This toggle is unavailable unless Initial Skeleton Pose is on. Default=off.

Link and Joint Settings—Resets link parameters and joint intersection parameters to their default values. Link parameters include Bend, Twist, Sliding, and Radial Scale values. For Radial Scale, this includes Scale, Amplitude, Stretch, and Breath. Default=off.

When this option is on, you can set new default values in the *Link Settings rollout (page 2–1136)* and the *Joint Intersections rollout (page 2–1140)*.

Bulges—Resets bulge angles to the default values. By default, one bulge angle is created per link. Turning on Bulges deletes any new bulge angles you might have created. Default=off.

When this option is on, you can set new default values for cross sections in the *Cross Sections rollout (page 2–1125)*.

Tendons—Resets tendons to the default values. Turning on Tendons deletes any tendons you might have created. Default=off.

Vertex-Link Assignments—Re-evaluates which vertices fall within each envelope. (This can vary from the original default settings, depending on whether you have made changes to the envelopes.) Turning on Vertex-Link assignments also removes any custom vertex assignments. Turn on Vertex-Link Assignments if you want to replace all manual vertex reassignments with the default vertex assignments. Default=off.

When this option is on, you can set new default values in the *Vertex-Link Assignment rollout (page 2–1123)*.

Vertex Settings—Recalculates vertex parameters at the link to which each vertex is assigned. Turning on Vertex Settings does *not* reassign manually reassigned vertices. Use this option when you have changed the link parameters and want to recompute the vertex parameters based on these changes. This check box is automatically turned on if you turn on either Initial Skeleton Pose or Vertex-Link Assignments. Default=on.

Bulge Editor 

Select an object that has the Physique modifier applied to it. > Modify panel > Physique rollout > Click Bulge Editor.

Creating and editing cross sections allows you to "bulge" the mesh. The Bulge Editor is an alternative to using the Bulge sub-object level to create and edit bulge angles. The difference is that the Bulge Editor allows you to create, select, and edit cross sections in schematic Cross Section and Profile views. Using the Bulge sub-object level, creating and editing bulge angles takes place in the viewports. All of the parameter changes you make in the Bulge Editor are also reflected on the mesh in the viewports.

On the Modify panel, you access the Bulge Editor from the *Physique rollout (page 2–1106)*, from the Bulge rollout at the *Bulge sub-object level (page 2–1141)*, or from the Link Settings rollout at the Link sub-object level.

Note: The Bulge Editor works exclusively with bulge angles. To create and edit Tendons, use the *Tendons sub-object level (page 2–1147)*.

Using an Animation to Preview Limb Orientations

For easier creation of bulge angles, you should create a simple animation that moves the limb into extreme orientations. In the case of the human arm, you might set keyframes with the arms down against the body; extended to the sides; bent at the elbows; and finally, more relaxed with the hands touching the shoulders. By scrubbing the time slider, you can easily choose one of many intermediate or extreme orientations. This saves time when you create Bulge Angles, because you won't need to exit Physique to manipulate the skeleton to change the orientation of the mesh. You can create and set multiple Bulge Angles without ever leaving the Bulge Editor or Bulge sub-object level.

Procedures

To create a new bulge angle using the Bulge Editor:

1.  Click the Bulge Editor button at the top level of the Physique modifier, or at the Link or Bulge sub-object level.

   The Bulge Editor is displayed.

   Tip: Accessing the Bulge Editor at the Bulge sub-object level provides the added benefit of letting you work in either the Bulge Editor or the Bulge rollout interchangeably.

2. In a viewport, click to select a link.

3.  On the Bulge Editor toolbar, click Insert Bulge Angle.

   Physique creates a new bulge angle. The number of the bulge angle name in the Current Bulge Angle field increments.

4. Type a descriptive name in the Current Bulge Angle field.

   Tip: It is a good idea to change the color of the newly created bulge angle. Do this by clicking Bulge Angle Color at the Bulge sub-object level, and selecting a color with the Color Selector.

5.  If no appropriate cross sections exist, then on the Bulge Editor toolbar, click Insert CS Slice. Click the Profile view to create and position a cross section.

   The cross section is created at the location you clicked.

6. In the viewports, rotate the joint to the desired angle.

   This is most easily done by creating a preview animation, as described in *Using an Animation to Preview Limb Orientations (page 2–1115)*: just move the time slider to a frame that has the angle you want. Otherwise, you will have to exit the Bulge sub-object level, select the appropriate limb, and rotate it. For example, to bulge the biceps on a biped, you might rotate the biped forearm to ninety degrees.

7. "Bulge" the mesh by editing cross section control points using the Bulge Editor's Cross Section view.

   As you scale or move control points in the Cross Section view, the mesh "bulges" in the viewports.

8.  On the Bulge Editor toolbar, click Set Bulge Angle on the Bulge Editor toolbar.

   Physique saves the current angle of the joint.

   When the joint angle is reached, the mesh bulges. By default, Physique creates one bulge angle when it is first attached to the mesh. So to make an arm or leg that bulges when biped joints are rotated, you need to create and set only one additional bulge angle.

To change a bulge angle value:

1. On the Bulge Editor, make sure the bulge angle you want to reset is displayed in the Current Bulge Angle field.
2. In a viewport, change the angle between the active link and its child link.

   Rotate a biped limb or use the time slider to move to a frame that has the limb rotated correctly.
3.  On the Bulge Editor toolbar, click Set Bulge Angle.

   The angle of the joint is set for the displayed bulge angle. When the biped limb rotates to this angle, the mesh bulge effect is at full strength.

To choose a specific bulge angle for editing:

1. Expand the Current Bulge Angle drop-down list (click the downward pointing arrow at its right).

   The full list of bulge angle names appears.
2. Click the name of the bulge angle you want to edit.

   The Bulge Editor's Cross Section and Profile views update to show the cross sections associated with the newly selected bulge angle.

To delete a bulge angle:

1. Expand the Current Bulge Angle drop-down list (click the downward pointing arrow at its right).

   The full list of bulge angle names appears.
2. Click the name of the bulge angle you want to delete.
3.  On the Bulge Editor toolbar, click Delete Bulge Angle.

   Note: You can't delete the default bulge angle: a link must always have at least one bulge angle defined.

To add a cross section:

1. On the Bulge Editor, make sure that the bulge angle you want to edit is displayed in the Current Bulge Angle field.
2.  On the Bulge Editor toolbar, turn on Insert CS Slice.
3. In the Profile view, click the link at the point where you want the cross section to be added.

   The cross section is created at the location you clicked.

   Tip: You can also add cross sections to the child link.

To delete a cross section:

1. In the Profile view, click to select a cross section.

   The selected cross section turns red.
2.  On the Bulge Editor toolbar, click Delete CS Slice.

   Note: You cannot delete the default cross section at the joint between the link and its child.

To make a cross section the active cross section:

- In Profile view, click a cross section to select it.

  The selected cross section displays in Cross Section view. You can change its shape by selecting and editing control points in Cross Section view. This will bulge the mesh at the location of the cross section.

To select multiple cross sections:

You can change parameters for multiple cross sections by selecting them and then adjusting parameters in the *Cross Sections panel (page 2–1121)*.

1.  On the Bulge Editor toolbar, turn on Select And Translate CS.
2. In Profile view, click to select a cross section.
3. Select additional cross sections by doing one of the following:
   - Use Ctrl+click to add cross sections to the selection.
   - Drag a rectangular region in the Profile view. All cross sections that the region surrounds or crosses are selected.

   Tip: Use Alt+click to remove a cross section from the selection.

To move cross sections along the link:

Moving a cross section repositions where on the mesh a bulge occurs.

1. Select the cross sections to move.
2.  On the Bulge Editor toolbar, click Select And Translate CS.

   Tip: If you are moving a single cross section, you can skip step 1 and click the cross section after you click Select And Translate CS Slice.
3. Drag left or right in the Profile view to move the cross section (or cross sections) left or right.

   Note: You cannot move cross sections through each other.

To copy and paste cross sections:

1. Select the cross section to copy.
2.  On the Bulge Editor toolbar, click Copy Selected CS.
3. Select a cross section to paste over.
4.  On the Bulge Editor toolbar, click Paste Selected CS.

   The cross section you pasted now has the shape of the cross section you copied.

   Note: You can copy a cross section and paste it to a similar cross section on another link. Create a bulge angle and an appropriate cross section on the target link before pasting. If necessary, use Mirror Selected CS after pasting.

To change the shape of a cross section:

1. In the Profile view, click to select a cross section.
2.  In the Cross Section view, use the transform tools to move, scale, or rotate the cross section's control points.

    You can also add or delete cross section control points.

    You can also use freehand drawing to add control points.

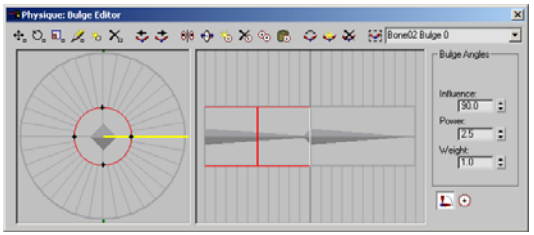
To change the Profile view orientation:

1. In the Cross Section view, move the cursor to the outer end of the orientation bar.

   The cursor changes to a shape like the Cross Section view.
2. Drag to rotate the orientation bar.

   When you release the mouse, the Profile view updates to show the profile at the angle of the orientation bar.

Interface



Toolbar


Select, Scale, and Rotate Control Points—Lets you select, scale, and rotate control points of a cross section in the Cross Section view.

This is comparable to using the Move transform on geometry in viewports. Control points are simultaneously scaled and rotated around the link, which in effect "moves" them.

Select a cross section in the Profile view first. Then select and move control points in the Cross Section view to change the bulge.


Select and Rotate Control Points—Lets you select and rotate control points in and around the center of the link in Cross Section view.

This does *not* twist the mesh. If the cross section is not a circle, a bulge migrates around the mesh.


Select and Scale Control Points—Lets you select and scale control points about the link's center in the Cross Section view.

If all cross section control points are scaled, the mesh bulges uniformly.


Draw Control Points—Lets you add control points by drawing freehand in the Cross Section view. The shape of the bulge changes based on the shape you draw.

In the Cross Section view, you can draw the cross section interactively, placing control points at each resolution step. You can freeze existing control points by selecting them in the Cross Section view and holding down the Ctrl key as you draw.

Note: The number of control points created depends on the cross section Resolution (see *Cross Sections panel (page 2–1121)*).

Tip: You can also use this too In Profile view. When you draw in Profile view, you change the profile of the bulge while placing a control point on each cross section of the bulge angle. This is useful mainly when the bulge angle uses multiple cross sections. You can rotate the orientation bar to draw whichever profile you desire. (Imagine the Cross Section view being like drawing around the arm, and the Profile view like drawing along the length of the arm.)


Insert Control Points—Lets you insert control points by clicking the cross section in the Cross Section view. A single control point is created each time you click.

Adding control points and changing their position lets you sculpt the cross section and control exactly where a bulge occurs on the mesh.


Delete Control Points—Deletes control points. First select the control points in the Cross Section view, and then click this button.


Previous Link and Next Link—Move to the next or previous link in the hierarchy. The Cross Section and Profile views update to display the appropriate bulge geometry.

Important: Profile view displays the parent link on the left and the child link on the right.


Mirror Selected CS—Mirror the cross section across the vertical plane running between the green profile reference points at the top and bottom of Cross Section view. In other words, clicking this button causes the bulge to appear on the opposite side of the link.


Select and Translate CS—Lets you select and move a cross section along its link. The bulge on the mesh migrates up and down the link as the cross section is moved.

 **Insert Cross Section Slice—**Lets you insert a cross section by clicking Profile view in the location you want the cross section to appear. A new cross section is created where you clicked, on either the parent or the child link.

Extra cross sections give you more control of how and where the mesh bulges.

 **Delete Cross Section Slice—**Deletes the currently selected cross section. First, select the cross section in Profile view.

 **Copy Selected CS—**Copies the selected cross section.

You can copy a cross section from one link to another link. Create a new bulge angle with an appropriate cross section in the target link first, then copy the cross section parameters form the source cross section, select the target cross section, and click Paste Selected CS.

 **Paste Selected CS—**Pastes copied cross section parameters onto another cross section.

First, select a cross section in Profile view, and click Copy. Then select another cross section in the Profile view, and click Paste.

 **Set Bulge Angle—**Associates the effect of the current bulge angle with the skeleton's current joint angle.

First, set the joint to the orientation at which you want the bulge to appear, then click this button. During animation, whenever the joint rotates near this angle, the mesh will bulge.

 **Insert Bulge Angle—**Adds a new bulge angle for the selected link.

The Current Bulge Angle field displays a bulge angle name with the number incremented. You can enter a descriptive name for any new bulge angle; for example, *Arm at 90*.

By default, one initial bulge angle is created when Physique creates links. Only one additional bulge angle will allow you to bulge the mesh. You can create more bulge angles for further control, if you like.

 **Delete Bulge Angle—**Deletes the current bulge angle. The current bulge angle displays in the Current Bulge Angle field.

First select a bulge angle in the Current Bulge Angle dropdown, and then click Delete Bulge Angle.

 **Select Nearest Bulge Angle—**When on, selects the bulge angle nearest to the current joint angle.

If a joint bends over time, you can use this button in conjunction with the time slider to select a bulge angle. If you play the animation, you can see the Current Bulge Angle field change to reflect the bulge angle nearest to the current angle of the selected link and its child (if two or more bulge angles for the limb exist).



Current Bulge Angle field and drop-down list—Displays the current bulge angle. You can choose a bulge angle from the drop-down list to edit the bulge or change its name. To change the name, simply enter a new one in the edit field.

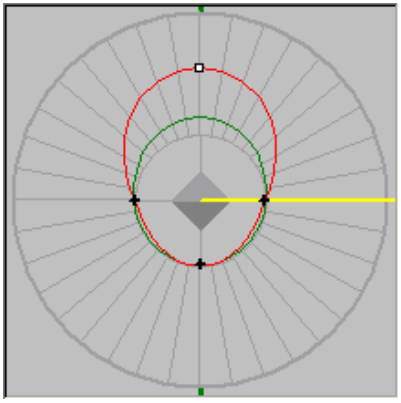
Tip: When you create a bulge angle, give it a descriptive name.

When you create a new bulge angle by clicking Insert Bulge Angle, this field displays the default bulge angle name, which is the previous name with its sequence number incremented.

Cross Section view

Cross Section view displays an outline of the active cross section. In *Profile view (page 2–1120)*, the active cross section is shown in a bright red color. In Cross Section view, you can edit cross sections to bulge the mesh.

Note: It is possible to select multiple cross sections in Profile view using either a rectangular selection region or Ctrl+click. This allows you to enter settings on the *Cross Sections panel (page 2–1121)* for all selected cross sections at once. Only one active cross section, however, can be viewed and edited in Cross Section view. (Selected but inactive cross sections are shown in a dark red color in the Cross Section and Profile views.)



These are the elements of the cross section display:

Gray square at center—Represents the link at the center of the cross section.

Control points—When unselected, are shown as small black crosses on the control spline. When selected, they are shown as black squares with a white center.

Resolution lines—Display as gray lines that surround the link radially. Control points snap to these lines as they are positioned.

You can increase or decrease the resolution by changing the Resolution setting. This control is on the Bulge rollout in Bulge Sub-Object. The Cross Section view, however, always displays 36 resolution lines.

Red line—The control spline for the active cross section.

The shape of the spline determines where mesh deformation occurs.

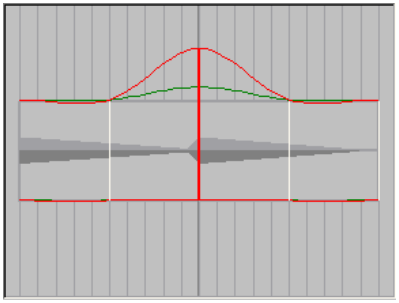
Green line—Represents actual mesh deformation. (Typically this is less exaggerated than the control spline.)

When you drag the time slider or play an animation, this updates according to the angle of the limb and the bulge angle parameters.

Orientation bar—The bright yellow radius indicates the orientation of the "slice" shown in profile view. Profile view always shows a vertical profile of Cross Section view, indicated by the green dots at the top and bottom of the view, just outside the circle that contains the resolution lines. By default, the orientation bar is level at 3 o'clock. You can change the Profile view orientation by dragging the end of the orientation bar (the cursor changes to a miniature cross section). This lets you view different angles of the bulge in Profile view.

Profile view

Use Profile view to select, move, and copy cross sections on the selected link and its child. Highlight a cross section in Profile view to display and edit it in *Cross Section view (page 2–1120)*.



Profile view is a schematic profile of two links. The currently selected link is on the left, and its child link is on the right. If the selected link is an end link, the outline of the right half of the Profile view turns gray.

Cross sections are shown as vertical bars across the profile. The active cross section is red. Unselected cross sections are white. Cross sections that are selected but not active are dark red.

You can use a rectangular selection region or Ctrl+click to select multiple cross sections. Although only one cross section at a time is active in Cross Section view, you can use the *Cross Sections panel (page 2–1121)* to change parameters for multiple cross sections.

Profile view shows a profile of the bulges you create. As in Cross Section view, the control spline is red and the deformation spline is green.

The profile is always a vertical profile of the Cross Section view. You can drag the orientation bar in Cross Section view to change the angle of the profile.

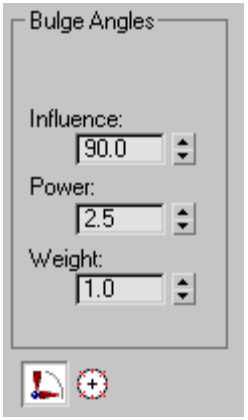

To insert a new cross section, turn on Create CS Slice and then click Profile view to set the cross section's location.


You can use the Draw tool in Profile view to change the control spline by freehand drawing. Drawing updates the cross section shapes.

Bulge Angles panel

At the right of the Bulge Editor are two panels for setting parameters. Use the Bulge Angles panel to change bulge settings for the currently active bulge angle.

Note: These parameters are the same as those in the Bulge Angle Parameters group on the Bulge rollout at the Bulge sub-object level. Changing a value in one location changes it in the other.




Bulge Angles—Click to display the Bulge Angles panel.

Influence—The range of angles through which the bulge influences the skin. Range=0 to 180. Default=90 degrees.

For example, if you've set a bulge angle for the joint at 90 degrees, an Influence value of 40 means that the bulge effect begins to appear when the rotating joint reaches 50 degrees (90 minus 40) or 130 degrees (90 plus 40).

Power—Controls how smoothly or abruptly the bulge takes effect.

At 0, the bulge takes effect immediately, without interpolated easing. As values increase, the bulge eases in gradually. A value of 10 will bulge the mesh abruptly when the set angle is reached. Range=0 to 10. Default=2.5.

Weight—Increases the effect of the current bulge angle relative to the effect of any other bulges. Range=0.0 to 100.0. Default=1.0.

Cross Sections panel

At the right of the Bulge Editor are two panels for setting parameters. Use the Cross Sections panel to change the cross section settings for a particular link.

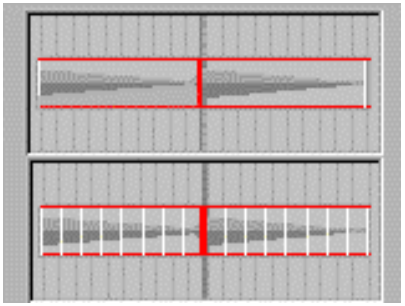
Note: These parameters are the same as those in the Cross Section Parameters group on the Bulge rollout at the Bulge sub-object level. Changing a value in one location changes it in the other.



 **Cross Sections—**Click to display the Cross Sections panel.

Sections—Sets the number of cross sections for the selected link.

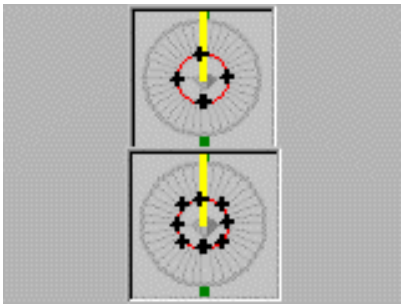
Tip: Rather than manually inserting cross sections in Profile view, use Sections to create cross sections. For a thigh or biceps bulge, you might need only one additional cross section in the middle of the link.



Above: Sections=1

Below: Sections=8

Divisions—Sets the number of control points on the selected cross sections. By default, control points are evenly distributed around the cross section control spline.



Above: Divisions=4

Below: Divisions=8

Resolution—Sets the number of radial divisions around the cross section. Control points snap to the nearest resolution line.

Entire Link—When on, selects all cross sections for all bulge angles in the selected link. Use this to globally change parameters for all cross sections on a link.

Entire Link also affects Copy Selected CS and Paste Selected CS. While Entire Link is on, Copy Selected CS copies all cross sections for all bulges, and Paste Selected CS pastes all cross sections for all bulges. Entire Link is useful for copying all the bulges from one arm or leg to another.

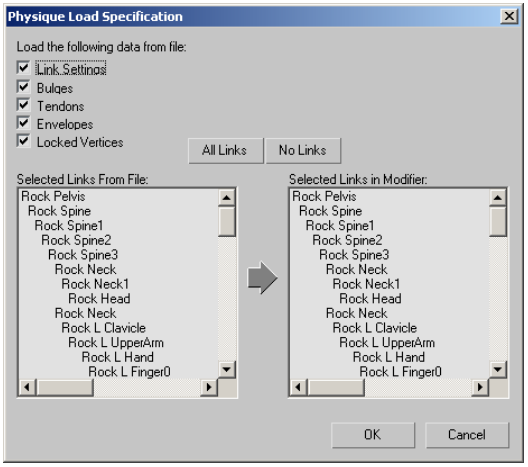
Note: When you turn on Entire Link, a number of toolbar controls, including the Current Bulge Angle name field, become unavailable. On the Bulge rollout at the Bulge sub-object level, the Current Bulge Angle changes to read "Entire Link."

Physique Load Specification Dialog

Select an object that has the Physique modifier applied to it. > Modify panel > Physique rollout > Click Open Physique (.phy) File. > Choose a file. > Physique Load Specification dialog

This dialog appears after you have chosen a Physique (*.phy)* file to load. It lets you specify which information you want to obtain from the file.

Interface



At the upper left of the dialog are check boxes that let you specify which kind of data to load. Turn off a check box to avoid loading that kind of data.

Link Settings—Loads link parameters. Default=on.

Bulges—Loads cross sections and bulge angles. Default=on.

Tendons—Loads tendons. Default=on.

Envelopes—Loads envelopes. Default=on.

Locked Vertices—Loads locked vertices.

All Links—Selects all links to load.

This button is enabled only if the number of saved links equals the number of open links. Click to select all links in both lists. If All Links is disabled, you must choose links by hand. Click the name of a link in a list to select it. Before you click OK, the number of links to load must equal the number of links to update.

No Links—Deselects all links in both lists.

After deselecting all links, you must select the links to load and to update, before you click OK.

Selected Links From File and Selected Links In Modifier—At the left of the dialog is a list of the links that were saved in the file, and at the right is a list of links in the currently open Physique skin and skeleton. You can choose which links to update with the saved data.

Physique Initialization Rollouts

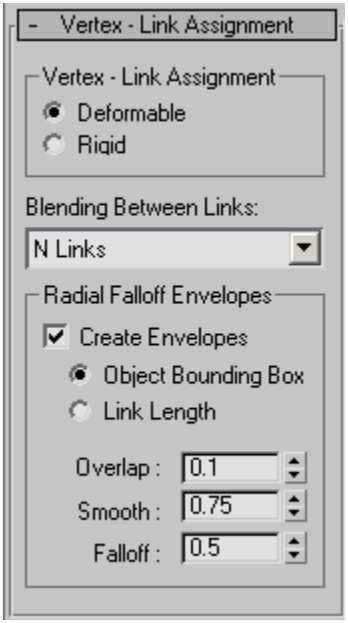
Vertex-Link Assignment Rollout

Select an object that has the Physique modifier applied to it. > Modify panel > Physique rollout > Click Attach to Node. > Select a root object in a viewport. > Physique Initialization dialog > Vertex-Link Assignment rollout

Select an object that has the Physique modifier applied to it. > Modify panel > Physique rollout > Click Reinitialize. > Physique Initialization dialog > Vertex-Link Assignment rollout

On the Vertex-Link Assignment rollout, you can choose to create new envelopes and pick the type of blending between links.

Interface



Deformable—Creates deformable envelopes for the links. Default=chosen.

Deformable envelopes determine vertex-link assignment based on the deformation spline that Physique creates through the links.

Rigid—Creates rigid envelopes for the links. Default=not chosen.

Rigid envelopes determine vertex-link assignment based upon the linear links in the hierarchy.

Blending Between Links—This drop-down list lets you choose how links are initially blended. Default=N Links.

Tip: Typically, you should leave this set to the default of N Links.

- **N Links—**Vertices are influenced by all overlapping envelopes.
- **No Blending—**Vertices are influenced by only one (1) link. This allows a mesh with the Physique modifier applied in **character studio 1** to work with **character studio 3**. Choose No Blending if you are developing characters for a game engine that doesn't support blending, or if you intend to use strictly **character studio 1**-style vertex-link assignments.
- **2, 3, or 4 Links—**Vertices are influenced by more than one link, but the number of links is limited to 2, 3, or 4, depending on your choice. Choose one of these options if you are developing characters for a game engine whose support of blending is limited.

Radial Falloff Envelopes group

Create Envelopes—When on, creates envelopes for the links. When off, envelopes are not created. Default=on.

Note: With no envelopes, you have to use hand-weighting for vertices. If you want to add envelopes when a model has none, reinitialize Physique.

- **Object Bounding Box—**Envelope size is based on the size of objects in the hierarchy, such as biped limbs, or bones.
- **Link Length—**Each envelope has a Radial Scale that is one-third of the length of the link.

Overlap—Sets how far the envelopes overlap. Default=0.1.

Smooth—Sets the distance between the inner and outer bounds of an envelope by scaling the outer boundary. Range=0.0 to 10.0. Default=0.75.

Falloff—Sets the rate of falloff between the inner and outer boundary of an envelope. The falloff rate is a Bezier function. Range=0 to 10. Default=0.5.

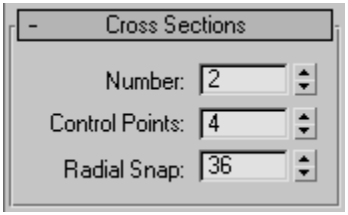
Cross Sections Rollout

Select an object that has the Physique modifier applied to it. > Modify panel > Physique rollout > Click Attach to Node. > Select a root object in a viewport. > Physique Initialization dialog > Cross Sections rollout

Select an object that has the Physique modifier applied to it. > Modify panel > Physique rollout > Click Reinitialize. > Physique Initialization dialog > Cross Sections rollout

The Cross Sections rollout lets you globally set how cross sections are initialized.

Interface



Number—Sets the initial number of cross sections per link. Default=2.

Control Points—Sets the initial number of control points per cross section. Default=4.

Radial Snap—Sets the initial number of snap positions per cross section (this is the value labeled "Resolution" in the *Bulge Editor (page 2–1114)* and at the *Bulge sub-object level (page 2–1141)*.

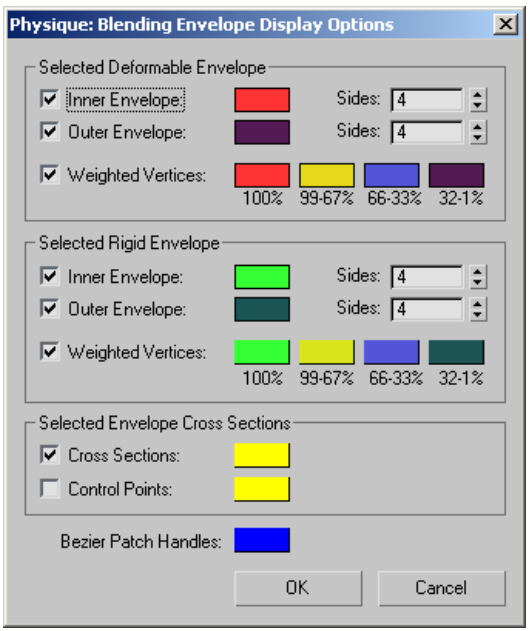
Sub-Object Options Dialogs

Blending Envelope Display Options Dialog

Select a mesh that has the Physique modifier applied to it. > Modify panel > Envelope sub-object level > Blending Envelopes rollout > Display group > Click Display Options.

This dialog lets you customize how Physique displays envelope blending at the Envelope sub-object level.

Interface



Note: All check box settings except Control Points are on by default.

Selected Deformable Envelope and Selected Rigid Envelope groups

These two groups have the same controls. By default, deformable envelopes are shown in red, and rigid envelopes are shown in green.

Inner Envelope—The check box toggles viewport display of inner envelopes. Click the color swatch to change the viewport color of the inner envelope boundary.

- **Sides—**Sets the number of inner envelope sides shown in the viewport display.

Outer Envelope—The check box toggles viewport display of outer envelopes. Click the color swatch to change the viewport colors of the outer envelope boundary.

- **Sides—**Sets the number of outer envelope sides shown in the viewport display.

Weighted Vertices—The check box toggles viewport display of weighted vertices. Click a color swatch to change viewport colors of weighted vertices in a particular percentage range.

Selected Envelope Cross Sections group

Cross Sections—Toggles viewport display of cross sections. Click the color swatch to change the viewport color of selected cross sections.

Control Points—Toggles viewport display of control points. Click the color swatch to change the viewport color of selected control points.

Bezier Patch Handles—Click to set the viewport color of Bezier patch handles.

Exclude Envelopes Dialog

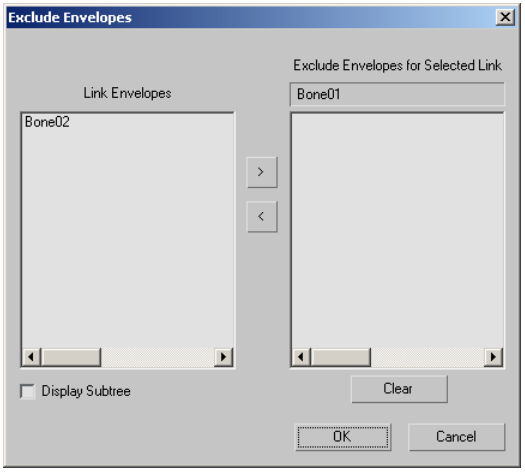
Select an object that has the Physique modifier applied to it. > Modify panel > Bulge sub-object level > Select an envelope. > Blending Envelopes rollout > Edit Commands group > Click Exclude.

At the *Envelope sub-object level (page 2–1130)*, this dialog lets you exclude links from influencing other links. For example, you can exclude the right thigh link from influencing the left thigh link. Or, rather than scale the index finger envelopes to avoid influencing middle finger vertices, you might exclude the middle finger links from the index finger links.

Tip: You can leave this modeless dialog open while you choose links in the viewports.

Note: If you have a Physique model created with **character studio 1** or **2**, excluding links will not work until you *reinitialize Physique (page 2–1113)* using the initial skeleton pose. In the Physique Initialization dialog, make sure Vertex-Link Assignments is turned on.

Interface



Link Envelopes list (left side)—Displays the links available for exclusion. Selected links are not in this list.

Exclude Envelopes for Selected Link list (right side)—Displays envelopes to exclude from the current link selection. The text field displays the name of the selected link, or "Multiple Links Selected."

Right arrow—Select links in the Link Envelopes List on the left, then click the right arrow to add them to the Exclude Envelopes For Selected Link list on the right.

Left arrow—Select links in the Exclude Envelopes For Selected Links list on the right, then click the left arrow to remove them from the exclusion list.

Display Subtree—When on, the Link Envelopes list increments links according to their level in the hierarchy. This visual aid can help you to find and select links. Default=off.

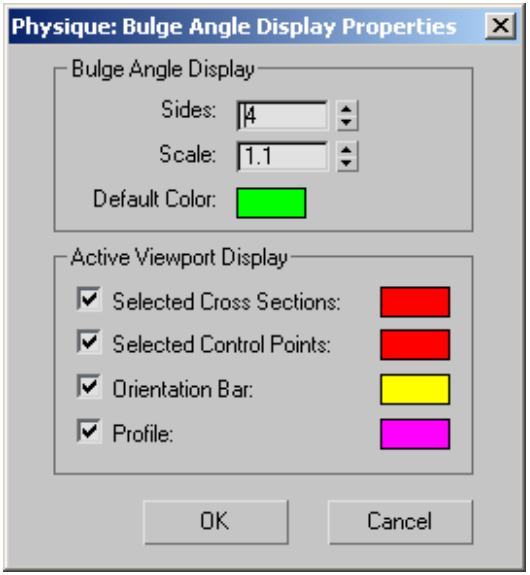
Clear—Clears all links from the Exclude Envelopes For Selected Link list.

Bulge Angle Display Properties Dialog

Select a mesh that has the Physique modifier applied to it. > Modify panel > Bulge sub-object level > Bulge rollout > Display group > Click Display Options.

This dialog lets you customize how Physique displays bulge angles at the Bulge sub-object level.

Interface



Bulge Angle Display group

Sides—Specifies the number of sides the bulge angle envelope displays in the viewports. Default=4.

Scale—Specifies the scale of envelope display in the viewports. Default=1.1.

Default Color—Click to change the bulge angle envelope color in viewports.

Active Viewport Display group

Selected Cross Sections—The check box toggles the display of selected cross sections. Default=on.

Click the color swatch to change the color of selected cross sections displayed in viewports.

Selected Control Points—The check box toggles the display of selected control points. Default=on.

Click the color swatch to change the color of selected control points displayed in viewports.

Orientation Bar—The check box toggles the display of the orientation bar. Default=on.

Click the color swatch to change the color of the orientation bar displayed in viewports.

Note: The orientation bar is displayed in the viewports when the Bulge Editor is open.

Profile—The check box toggles the display of the profile. Default=on.

Click the color swatch to change the color of the profile displayed in viewports.

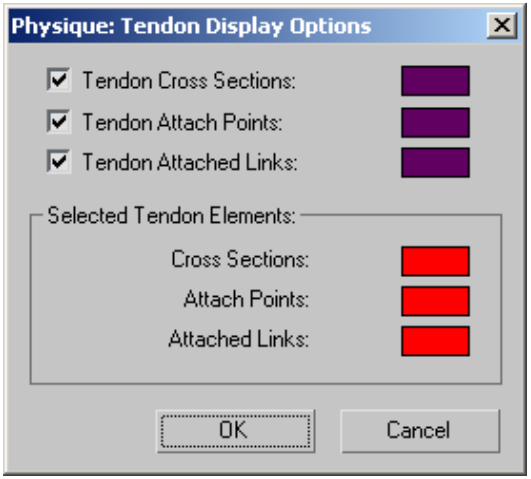
Note: The bulge profile displays in the viewports when the Bulge Editor is open.

Tendon Display Options Dialog

Select a mesh that has the Physique modifier applied to it. > Modify panel > Tendons sub-object level > Tendons rollout > Display group > Click Display Options.

This dialog lets you customize how Physique displays tendons at the Tendons sub-object level.

Interface



Tendon Cross Sections—The check box toggles the display of tendon cross sections. Default=on.

Click the color swatch to change the color of tendon cross sections displayed in viewports.

Tendon Attach Points—The check box toggles the display of tendon attach points. Default=on.

Click the color swatch to change the color of attach points displayed in viewports.

Tendon Attached Links—The check box toggles the display of attached links. Default=on.

Click the color swatch to change the color of attached links displayed in viewports.

Selected Tendon Elements group

Cross Sections—Click to change the viewport color of selected cross section.

Attach Points—Click to change the viewport color of selected attach points.

Attached Links—Click to change the viewport color of selected attached links.

Type-In Weights Dialog

Select a mesh that has the Physique modifier applied to it. > Modify panel > Vertex sub-object level > Vertex-Link Assignment rollout > Vertex Operations group > Click Select by Link. > Click a link in a viewport. > Lock Assignments > Click Type-In Weights.

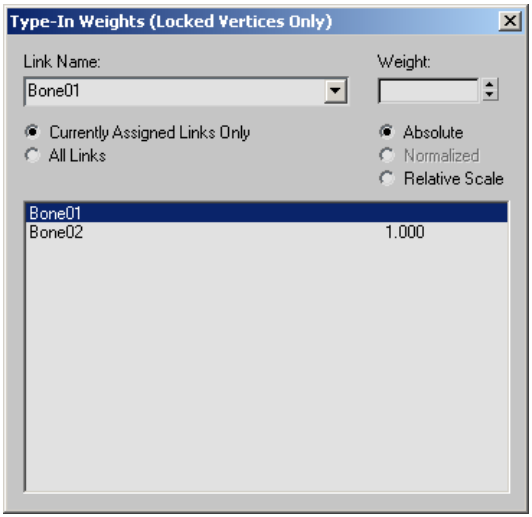
Lets you enter a weight for selected locked vertices. To use this dialog, select one or more vertices, then lock them, click Type-In Weights to display the Type-In Weights dialog, use the lists in the dialog to select a link, and then change the weight to that link in order to position the vertices (vertex positions update in viewports as the weight changes).

Note: You'll primarily use Type-In Weights to correct flaws on low- to medium-resolution meshes. On a high-resolution mesh, adjusting envelopes should be used to correct deformation.

Tip: Adjust weights for stubborn vertices that fall in blended regions between two or more envelopes. Subtle adjustments can be made to the vertex

weights of specific links that would be difficult to achieve by adjusting envelopes alone.

Interface



Link Name—Use the drop-down list to select a link and display the vertex weight to that link.

Currently Assigned Links Only—When chosen, displays only those links that presently influence the selected vertices.

All Links—When chosen, displays all available links.

Weight—Displays the vertex weight for the active link (the link currently displayed by the Link Name list). Use the spinner to change the weight value, if necessary.

Absolute—Uses an absolute value for vertex weight.

Normalized—This option is unavailable in **character studio**.

Relative Scale—Uses relative values for vertex weights.

This is the mode of choice when you work on multiple vertices. If a particular link is not having a strong enough effect on the selected vertices, you should lock them, open the Type-In Weights dialog, select the link, and choose Relative Scale. Then you can enter relative values in the Weight field. For example, entering **1.0** leaves weights the same, while entering **2.0** would double the effect of the selected link on the selected vertices.

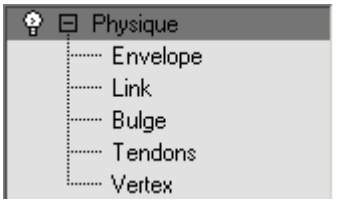
List of links—Displays a list of links, with weight values at the right. The weight field is blank when it displays absolute weights. (When more than one vertex is selected, each is likely to have a slightly different weight.) Click a link in this list to select it.

Physique Modifier Sub-Object Levels

Physique Sub-Objects

Select an object that has the Physique modifier applied to it. > Modify panel > In the modifier stack display, expand the Physique sub-object hierarchy.

The Physique modifier has five different sub-object levels. Each kind of sub-object has various controls.



- *Envelope (page 2–1130)*

  Envelopes define a link's area of influence on mesh vertices.

- *Link (page 2–1135)*

  Links are the segments of the *Physique deformation spline (page 2–1135)*. Link

sub-object controls adjust the way link movement affects the mesh.

- *Bulge (page 2–1141)*

  Bulge controls are for creating flexing muscles.

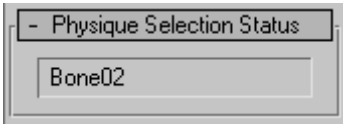
- *Tendons (page 2–1147)*

  Tendons provide a way for multiple links to affect an area of the mesh.

- *Vertex (page 2–1150)*

  For Physique, Vertex controls let you fine-tune the influence of envelopes.

Physique Selection Status Rollout



This rollout appears for all of the Physique sub-object levels. For all levels except Vertex, it shows the name of the hierarchy object associated with the active link. For the Vertex sub-object level, it prompts you to "Select Vertices" unless Select By Link or Assign To Link is active, in which case it displays the name of the hierarchy object associated with the active link, as it does at other sub-object levels.

Envelope Sub-Object

Select an object that has the Physique modifier applied to it. > Modify panel > Bulge sub-object level

Use Envelope sub-objects to adjust the way skin behaves. Each link has an envelope, and the shape of the envelope determines which vertices are affected by the link's movement. An envelope has an inner and an outer bound: outside these bounds, vertices are not effected. Inside these bounds, the envelope's influence is strongest at the inner bound. A Falloff setting lets you control how quickly influence falls as it approaches the outer bound.

Envelopes also control blending: they provide smooth deformation of the skin across joints. Vertices that fall between overlapping envelopes receive influence from each, thus creating a smoothly blended transition. The outer bounds of the envelopes of adjacent links should overlap enough to provide a smooth blend at the joints. You can adjust the relative strength of envelopes, giving one link's envelope more influence than another.

At the Envelope sub-object level, in addition to the settings already mentioned, you can adjust an envelope's Radial Scale. This is useful for making sure the envelope encloses vertices along the length of the link. The settings for Strength, Falloff, Radial Scale, and Overlap are grouped under *Envelope Parameters (page 2–1133)*.

Envelopes have cross sections. You can add cross sections to refine the shape of an envelope, and you can move or scale a cross section's control points to change the cross section's shape. (Cross sections are also used to create *bulges (page 2–1141)*.)

Tip: You can also use tools on the Main toolbar to adjust envelopes. For example, you can use Non-Uniform Scale on envelopes and cross sections, or use Select And Move to move envelopes, cross sections, or control points on cross sections.

Workflow

The goal is to modify the envelopes so each vertex in the mesh is encompassed by at least one link's envelope.

Usually the first adjustment is to assign a rigid envelope to the character's head. A rigid envelope blends with other envelopes but retains the shape of the mesh. The head should not deform with the deformation spline.

With Initial Skeletal Pose turned on, check that the envelopes enclose all areas of the mesh. If you turn off Initial Skeletal Pose, the character adopts its animated position at the current frame. Use an animation that stretches the character around, such as a run or dance motion. Find a frame where the envelopes need adjusting, and edit the envelope parameters. Changing the envelopes with the character in an animated position will always reference the Initial Skeletal Pose.

These are some frequently encountered problems with the default vertex assignment:

- Vertices that don't follow the skeleton.

  To correct this, increase the Radial Scale of an envelope's outer boundary.

- Vertices at joints get pushed or dented in.

  Too many envelopes might be affecting these vertices. Try reducing the Radial Scale of the inner bound, and reduce the Overlap values of the links on both sides of the joint.

- Bending appears too linear or broken at the joint.

  Increase the envelope's outer Overlap value toward the link on the other side of the joint.

Some further adjustments are described in the "Procedures" section that follows.

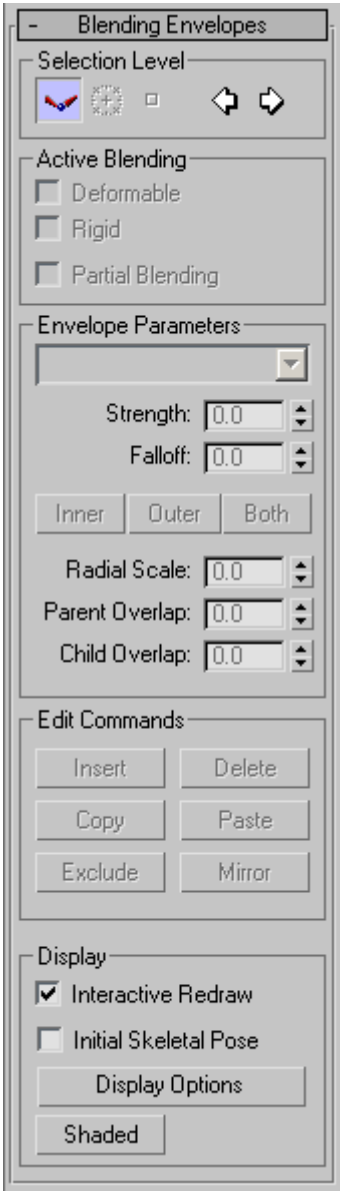
Note: An envelope does not need to enclose Patch tangent handles. It needs to enclose only the patch vertices. This relieves you from having to scale the envelopes to a large size to encompass tangent handles.

Procedures

See *Adjusting Default Envelope Shape (page 2–1086)*

Interface

Controls for Envelope sub-objects are on the Blending Envelopes rollout.



Selection Level group




Link—Turn on to select links in the viewports and edit the selected link's envelope parameters.

For example, you might turn on Link, select a biceps link, Ctrl+click to add the opposite bicep to the selection, then edit envelope parameters for both links at the same time.


Cross Section—Turn on to edit envelope cross sections, changing the envelope's shape and thus its area of influence.

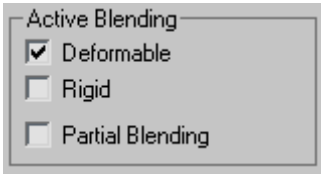
For example, you might turn on Cross Section, select a cross section on the inner or outer boundary of an envelope, and then move or scale it. You might Non-Uniform Scale the cross section of a collar envelope so it avoids vertices in the chest area.


Control Point—Turn on to edit the control points on a cross section.

For example, you might turn on Control Point, select a point on a cross section of an envelope, and move the point to change the envelope's shape and area of influence.


Previous and Next—Click to move to the next or previous link, cross section, or control point, depending on which selection level is active.

Active Blending group



Deformable—When on, enables a deformable envelope for the selected links. Default=on.

By default, deformable envelopes are shown in red.

Rigid—When on, enables a rigid envelope for the selected links. Default=off.

By default, rigid envelopes are shown in green.

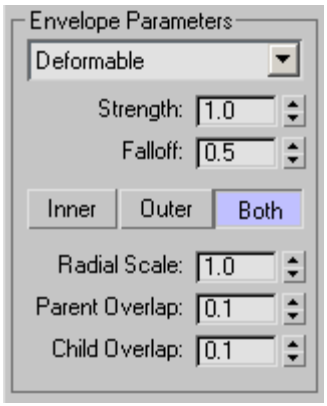
Important: You can have both a deformable and a rigid envelope turned on for the same link. Normally, you use one or the other. By non-uniform scaling the rigid and deformable envelopes for one link, you can position one envelope on top of the other. For example, you could control the shoulder with a rigid envelope and the armpit with a deformable envelope. Both envelopes can also be turned off for a link.

Partial Blending—Turns on partial blending for the selected links. Leave either a deformable or rigid envelope on, and then turn on Partial Blending. Physique calculates the weights of each link on a given vertex. If Partial Blending is *off*, and the total weight is less than 1, Physique normalizes the combined weight to 1. If Partial Blending is *on*, the weights are allowed to remain at a value of less than 1. The remainder is filled in by the root node, the equivalent of no influence. See *Partial Blending and Weight Assignments (page 2–1092)* for a detailed explanation of how the program calculates these values.



Partial Blending on all the deformable envelopes for the bones in this character's jaw allows smooth mesh deformation when the bones are moved.

Envelope Parameters group



Envelope type drop-down list—Shows the type of the selected envelope. If the link has both a rigid and a deformable envelope, you can use this list to choose which envelope's parameters you are adjusting.

Strength—Changes the strength of the envelopes. Range=0.0 to 100.0. Default=1.0.

Strength applies to both the inner and outer envelope bounds.

Used primarily for areas where envelopes overlap, and you want one to be more influential than the other.

Falloff—Changes the rate of falloff between the inner and outer bounds of an envelope. This is a Bezier curve function. Range=0.0 to 1.0. Default=0.1.

Vertices within the inner bound are fully influenced (weight=1.0) and those beyond the outer bound get no influence from the link (weight=0.0). Falloff determines the rate at which the influence falls off from 1.0 to 0.0.

The Inner, Outer, and Both buttons determine whether the controls that follow, Radial Scale, Parent Overlap, and Child Overlap, apply to the envelope's inner bound, its outer bound, or both at once. First, use the buttons to choose which bounds to adjust, then change values in the spinners.

Inner—Turn on to change the values of the inner bound.

Outer—Turn on to change the values of the outer bound.

Both—Turn on to change the values for both inner and outer bounds at the same time.

When Both is selected, the values displayed for Radial Scale and Parent and Child Overlap reflect the values for the *inner* bound.

Radial Scale—Radially scales the envelope bounds. Range=0.0 to 100.0. Default=1.0.

Parent Overlap—Changes the envelope's overlap with the parent link in the hierarchy. Range= -1.0 to 10.0. Default=0.1.

A value of 0.0 causes the end of the envelope to fall on the joint. Values less than 0.0 bring the envelope inside the link, and values above 0.0 will overlap onto the adjoining link.

Child Overlap—Changes the envelope's overlap with the child link in the hierarchy. Range= -1.0 to 10.0. Default=0.1.

A value of 0.0 causes the end of the envelope to fall on the joint. Values less than 0.0 bring the envelope inside the link, and values above 0.0 will overlap onto the adjoining link.

Edit Commands group



Note: Which buttons are available in this group will vary, depending on whether Links, Cross Sections, or Control Points is the active selection level.

Insert—Inserts a cross section or control point on a cross section.

Delete—Deletes a cross section or control point.

Copy—Copies an envelope or cross section.

Paste—Pastes an envelope or cross section.

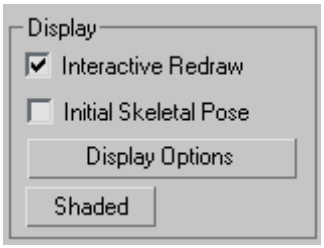
Exclude—Clicking this button displays the *Exclude Envelopes dialog (page 2–1126)*. You can exclude a link from influencing another link. For example, you can exclude the right thigh link from influencing the left thigh link. Or, rather than scale the index finger envelopes to avoid influencing middle finger vertices, exclude the middle finger links from the index finger links.

Mirror—Mirrors the envelopes on a selected link, or mirrors selected cross sections in an envelope.

After a Mirror operation, you can adjust the orientation by clicking Rotate on the Main toolbar, choosing Local coordinates, and then clicking and dragging the link or cross section.

Tip: To copy an envelope to its mirror, the sequence should be Copy, select the opposite link, Paste, then Mirror.

Display group



Interactive Redraw—When on, Physique dynamically updates the mesh while you adjust envelopes. When off, the mesh updates only when you enter a final value (you press Enter or Tab, or release the mouse). Default=on.

Initial Skeletal Pose—When on, puts the mesh character in the pose it was in just before Physique was applied. Default=off.

Display Options—Clicking this button displays the *Blending Envelope Display Options dialog (page 2–1125)*, which lets you customize envelope display.

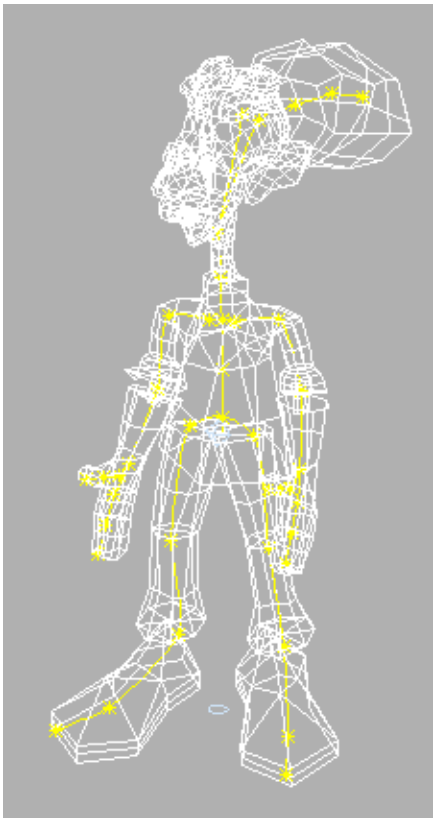
Shaded—Toggles shaded display of vertex weights in the viewports. Default=off.

Link Sub-Object

Select an object that has the Physique modifier applied to it. > Modify panel > Link sub-object level

Use parameters at the Link sub-object level to change how deformation around joints occurs. When a joint in the skeleton bends or rotates, Physique, by default, deforms vertices uniformly on either side of a joint. You can change these defaults by using the tools at the Link sub-object level. For example, you can adjust the amount of skin sliding that occurs along a limb as the limb bends, or change the angle of the crease between the upper arm and chest.

The Physique Deformation Spline



The deformation spline displays as a yellow curve that runs through the mesh.

Like a spline object, the *deformation spline* created by Physique is a continuous curve through several points. While a spline object runs continuously from vertex to vertex, the deformation spline is a smooth curve running from joint to joint. The Bend, Bias, and Tension spinners can change the shape of the curve, much as you can rotate or scale the handles at a spline vertex.

The deformation spline also takes into consideration twisting and scaling of the skeleton's links. At the Link sub-object level, you take control of the behavior of the deformation spline, and subsequently gain full control of the skin's behavior relative to the skeleton's movement.

Interface

Link Settings Rollout (page 2–1136)

Joint Intersections Rollout (page 2–1140)

Link Settings Rollout

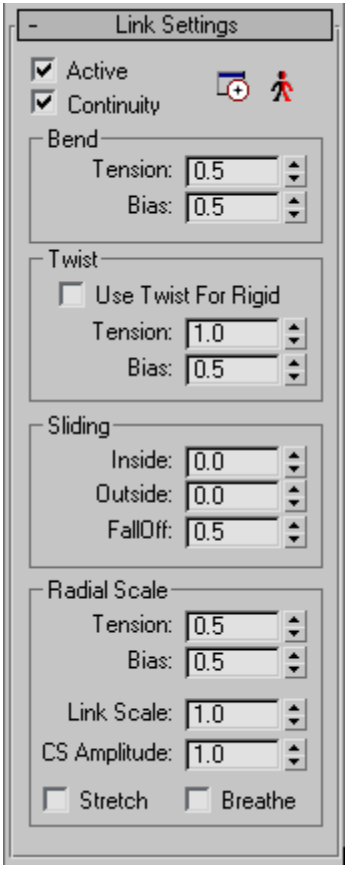
Select a mesh that has the Physique modifier applied to it. > Modify panel > Link sub-object level > Link Settings rollout

Controls on the Link Settings rollout let you adjust the behavior of the selected link.

See also

Joint Intersections Rollout (page 2–1140)

Interface



Active—When on, activates the selected link. Default=on.

Turning off Active makes the link unavailable for vertex assignment, meaning that the link has no influence on *any* vertices. Vertices within range of this link are not influenced by it, and can be "picked up" by other nearby envelopes, or can be manually assigned (without blending) to any other link. Turning off Active makes Physique ignore the link as if it were never part of the skeleton.

Continuity—Maintains a smooth transition across the joint from the parent link to the current link. Default=on.

When Continuity is on, the effect of the link parameters passes smoothly across the joint to the connected link. When Continuity is off, Bend, Twist, and Radial Scale settings are limited to the current link, which produces an abrupt transition across the joint. This is analogous to breaking the spline handles on a spline shape.



Bulge Editor—Displays the *Bulge Editor (page 2–1114)*. The Bulge Editor displays bulge cross sections schematically, letting you create and edit bulge controls.

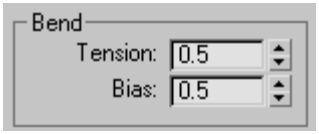


Reinitialize Selected Links—Reinitializes the link and its vertex parameters based on the current link parameter settings. Clicking Reinitialize Selected Links does not change vertex assignments or manual reassignments.

For example, increasing the tension of a link can cause the spacing to change for the link's cross sections. Reinitialize Selected Links smoothes the spline, and makes cross section spacing even again.

Think of the skin being snugly stuck to the link. As the link parameters change, the vertices can bunch up or stretch out. Clicking Reinitialize lets the skin "slip" to its original shape, establishing a new relationship to the deformation spline.

Bend group



Tension—Sets the smoothness of a joint. Tension affects the curvature of the deformation spline through the joint. A low value, near 0.0, makes the spline linear and creates a sharp angle at the joint, like bending a hinge. High values, near 2.0, make the spline smooth through the joint, creating a rounded joint, like bending a firm hose. Range=0.0 to 2.0. Default=0.5.

Bias—Displaces the pivot point about which vertices are bent. Bias pushes the effect along the spline toward one side of the joint or the other. The default value of 0.5 centers the bend at the joint. Values lower than 0.5 move the pivot onto the child link. At 0.0, the bend effect is limited to the selected link. Range=0.0 to 1.0. Default=0.5.

Twist group



Twist parameters control the way the skin deforms when a joint rotates along its length, as in turning a doorknob.

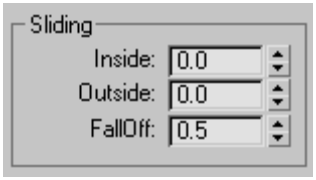
Use Twist for Rigid—When on, twist is used for rigid as well as deformable envelopes. Default=off.

Turn off for the rigid envelope for a character's head: the head should not twist along the length of the link. You might want to turn this on for the forearm or upper-arm link if you're using rigid envelopes.

Tension—Values lower than 1.0 concentrate the effect closer to the joint. Values higher than 1.0 move the effect away from the joint. Range=0.0 to 2.0. Default=1.0.

Bias—Shifts the distribution of the twist from one side of the joint to the other. The default value of 0.5 twists the selected link and the child link equally. Values greater than 0.5 shift the twist to the child link. Range= 0.0 to 1.0. Default=0.5

Sliding group

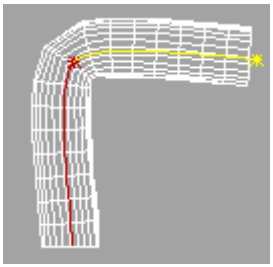


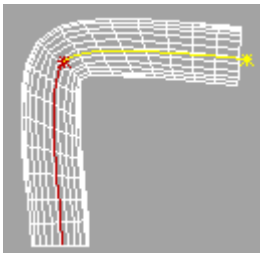
Skin sliding parameters control the amount of skin sliding that occurs when a joint rotates. Without skin sliding, vertices closest to the joint tend to compress on the inside and stretch apart on the outside, generally revealing the segments of the mesh. Outside sliding causes the vertices around the joint to move toward the joint, preventing localized stretching on the side that is greater than 180 degrees. Inside sliding causes the vertices to relax and slip away from the joint, preventing bunching of vertices on the side having an angle less than 180 degrees.

Tip: Use sliding for knees and elbows.

Inside—As values increase, skin moves away from the joint. Range=0.0 to 1.0. Default=0.0 (no inside sliding).

Outside—As values increase, skin moves toward the joint. Range=0.0 to 1.0. Default=0.0 (no outside sliding).



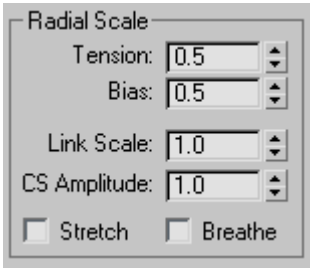


Above: Inside and Outside set to 0.0 (no sliding)

Below: Inside and Outside set to 0.25 (sliding creates a smoother bend)

FallOff—As values increase, the effect is localized to the joint. Range=0.0 to 1.0. Default=0.5.

Radial Scale group



Radial Scale parameters expand or contract the skin by scaling the radial distance perpendicular to the link. They apply to any combination of user-definable scale, bulge settings, or link length.

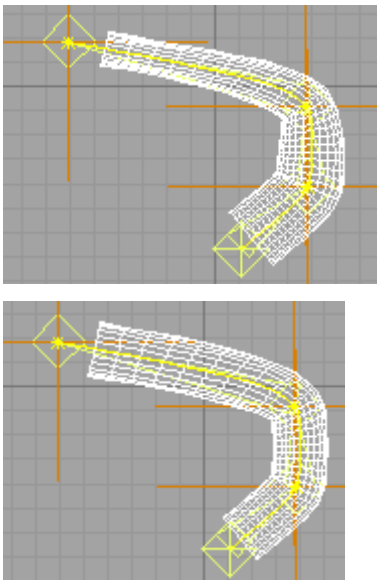
Tension—Values between 0.0 and 1.0 concentrate the effect closer to the joint. The effect is closest to the joint when Tension is 0.0. Values greater than 1.0 move the effect away from the joint. Range=0.0 to 2.0. Default=0.5.

Bias—Shifts the effect of radial scaling. At the default value of 0.5, scaling affects both the selected link and the child link. Values lower than 0.5 shift the scaling effect onto the selected link, and values greater than 0.5 shift scaling onto the child link. At 0.0, expansion and contraction are limited to the selected link. At 1.0, expansion and contraction are limited to the child link. Range=0.0 to 1.0. Default=0.5.

Link Scale—Scales the entire link radially, independent of the effect of any cross sections. At 1.0, the link is "actual size" and this setting has no effect. Other values increase or decrease the radial scale of the link. Range=0.0 to 10.0. Default=1.0.

CS Amplitude—Cross section amplitude has no effect unless the link has bulge angle cross sections. At 0.0, cross section deformation is turned off. Values greater than 1.0, up to the maximum of 10.0, exaggerate the effect of cross sections. Default=1.0.

Stretch—When on, preserves the volume of the link's skin when the length of the link changes. The effect is similar to stretching or squeezing a material that is only partly elastic, such as a tough rubber hose. Default=off.

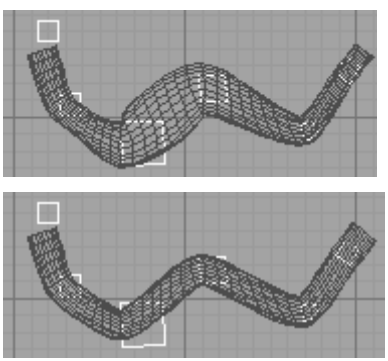


Above: Stretch is turned on.

Below: Stretch is turned off.

As the bones change in length, the mesh expands and contracts when Stretch is on.

Breathe—When on, scaling a skeleton node changes the radial scale of the link's skin. When off, scaling a node has no effect on the scale of the skin. Default=off.



Above: Breathe is turned on: the mesh is scaled as the object is scaled.

Below Breathe is turned off: object scaling doesn't effect the mesh.

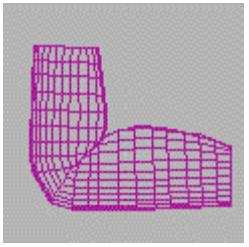
The linked objects are shown in white. The lowest, largest object is the one that has been scaled. These linked objects were used as the hierarchy for Physique.

Joint Intersections Rollout

Select a mesh that has the Physique modifier applied to it. > Modify panel > Link sub-object level > Joint Intersections rollout

When a joint bends, the skin can "collide." Without collision detection, it can overlap unrealistically. This is especially likely when one or both of the links have bulges. The joint intersection controls can detect skin collisions and correct overlap by creasing the skin.

By default, Physique draws a crease plane that bisects the joint and prevents vertices from each link from crossing the crease plane. The result can be a somewhat unnatural flattening of the bulges along the plane. Physique therefore provides controls that let you modify the effect of the crease plane.



Properly adjusted crease plane

You can modify the effect of the joint creases at both ends of the selected link. These joints are the *parent's joint* and the *link's joint*. Blend From and Blend To points define the distance along the link that is affected by the crease plane.

- The parent's joint is between the selected link and its parent. The parent's joint is numbered 0.

  For the parent joint, the joint is at 0.0 and the crease plane has an effect from the joint to the Blend From point. Between Blend From and Blend To, the crease plane has a partial effect, adjusted by the Bias setting. Vertices at locations greater than the Blend To point are not affected by the crease plane.

- The link's joint is between the selected link and its child. The link's joint is numbered 1.

  For the link's joint, the joint is at 1.0 and the crease plane has an effect from the joint to the Blend To point. Between Blend To and Blend From, the crease plane has a partial effect, adjusted by the Bias setting. Vertices at locations less than the Blend From point are not affected by the crease plane.

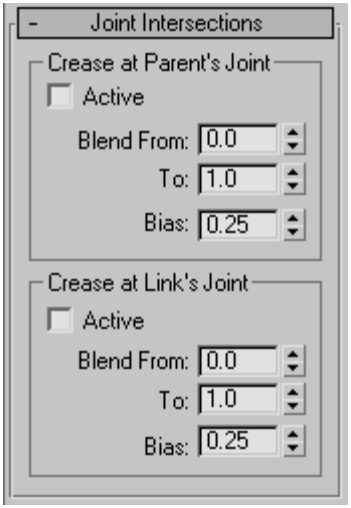
Procedure

To adjust joint intersection parameters:

1. With Physique active on the Modify panel, go to the Link sub-object level.
2. In a viewport, click the link you want to work with.
3. Select Active for the crease you want to manipulate: choose Crease At Parent's Joint or Crease At Link's Joint.
4. Change settings on the Physique Joint Intersections rollout.
5. Select the opposite link and adjust the opposite set of joint intersection parameters.

For example, if you are working with the joint between the selected link and its child (such as the biceps link), turn on Active for Crease At Link's Joint and change settings. Then select the child link (the forearm) and activate its Crease At Parent's Joint. Adjustments to the Parent's Joint for this link affect the other side of the same crease.

Interface



Both the parent's and the link's joint have the same controls:

Active—Turns off the effects of the joint intersection controls. When this box is cleared, Physique makes no compensation for overlapping bulges. Default=off.

Blend From—The area between Blend From and Blend To contains vertices that are partially affected by the crease plane. These vertices will be shifted, but not as much as those between the crease plane and the point where this joint begins to blend (Blend From in the parent's case, Blend To in the link's case). The distance of the vertices from the crease plane determines the exact amount of shifting. Vertices within the blend region that are closer to the intersection plane are shifted more than the vertices further away from it. Default=0.0.

Blend To—Indicates the distance beyond which the crease plane has no effect. Default=1.0.

Bias—Sets the strength of the crease plane effect within the blended region. A value of 0.0 means that the intersection plane will have no effect within the blended region; a value of 1.0 means that it will have a full effect within the blended region. Default=0.25.

Bulge Sub-Object

Select an object that has the Physique modifier applied to it. > Modify panel > Bulge sub-object level

After you have edited *envelope sub-objects (page 2–1130)* for good overall mesh deformation, you can create bulges to simulate muscle contraction and expansion when a character's joint rotates.

A bulge angle requires two links: the selected link, and the child link of that selection (for example, an upper arm and forearm). The joint that separates the two is referred to as the *bulge joint*. Any rotation applied to this joint becomes the angle used to center the effect of a bulge angle.

Physique creates a single default bulge angle for each bulge joint in the attached skeleton. The angle of this joint is that of the initial skeleton pose. This is the pose the skeleton had when the Physique modifier was applied to the skeleton, or when you use *Reinitialize (page 2–1113)* and turn on Initial Skeleton Pose in the Physique Initialization dialog. To create a bulge, you need to add only a single additional bulge angle. You can add further bulge angles to refine the bulge effect.

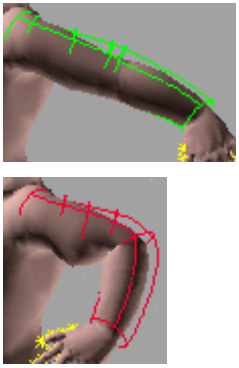
Workflow to Create a Biceps Bulge

The overall steps for creating a biceps bulge are:

- Insert a new bulge angle.

  Give the new bulge angle a custom name and color.
- Set the bulge angle at a point where the joint is flexed; for example, at 90 degrees.
- Insert a cross section on the upper arm.
- Use control points (or the *Bulge editor (page 2–1114)*) to shape the bicep bulge.

The following paragraphs elaborate these steps.



Above: Default bulge angle at the initial pose.
Below: Bicep bulges at the custom bulge angle of 90 degrees. The new bulge angle has been assigned a red color.

To create a biceps bulge angle, go to the Bulge sub-object level. In the Selection Level group, turn on Link. Select the upper arm link, then click Insert Bulge Angle; the bulge angle name increments from 0 to 1 in the Current Bulge Angle field, indicating that a new bulge angle has been created. Type a descriptive name for the new bulge angle in the Current Bulge Angle field; for example, *Arm at 90*.

Click Bulge Angle Color and select a color for the newly created bulge angle. When Select Nearest Bulge Angle is active, the bulge angle color changes as the joint bends; this is a good visual indicator of which angle has the most influence at any given frame.

Scrub the time slider to a frame where the arm is bent to 90 degrees, then click Set Bulge Angle (you can also rotate the joint).

Note: The Set Bulge Angle command records only the angle between the biceps and forearm, not the frame number where that angle occurs. The bulge will occur at *any* frame where the elbow is at the bulge angle.

In the Selection Level group, click Cross Section. Then in the Cross Section Parameters rollout, click Insert. In a viewport, position the cursor over the upper arm link, and click to place the new cross section for the biceps.

On the Selection Level rollout, click Control Point. To shape the cross section, move the cross section control points in the viewports, creating a bulge. Click Play. The biceps should grow as the angle of the character's arm approaches 90 degrees. If Select Nearest Bulge Angle is active, the bulge angle gizmos change colors, depending on which angle has the greatest influence.

Go to the Link selection level and use Copy, Paste, and Mirror to create an identical bulge angle for the opposite arm.

Tip: When you use Physique with a biped, load a *.bip* file that will rotate the character's limbs. Turn on In Place mode before you scrub the time slider to locate appropriate limb angles. This keeps the character (the biped) in the viewports.

Bulge Editor

The *Bulge Editor (page 2–1114)* duplicates many of the controls at the Bulge sub-object level. The benefit of using the Bulge Editor is that bulge angle data is represented graphically, and provides an alternative way of creating, selecting, and editing cross sections. Rather than editing cross section control points in the viewports to "bulge" the mesh, you might prefer to open the Bulge Editor and move control points in the Bulge Editor's Cross Section view.

Procedures

To create a new bulge angle on a selected link:

1.  With Bulge sub-object active, turn on Link in the Selection Level group.
2. Select a link in a viewport.
3.  Click Insert Bulge Angle.

This creates a new bulge angle. The name in the Current Bulge Angle field increments.

Note: By default, one bulge angle is created by Physique in the initial skeletal pose; the arm is usually straight in this pose.

4. In the Current Bulge Angle field, enter a name, such as **Arm at 90**.
5. Click the Bulge Angle Color swatch and choose a color in the color selector. This makes identification easier.

   Now you need to set a joint angle for the newly created bulge angle.
6. Move the time slider until the joint is at the desired angle.

   If the skeleton is not yet animated, you can exit the Bulge sub-object level, rotate the skeleton's joint until it is at the desired angle, then return to the Bulge sub-object level.
7.  Click Set Bulge Angle. The joint angle between the link and its child is recorded (for example, if a biped's upper arm link is selected, Set Bulge Angle records the angle created by the upper arm and the forearm.)
8.  In the Selection Level group, turn on Cross Section. Then in the Cross Section group, click Insert. In a viewport, click the link at the point where you want the new cross section to be.
9. Scale the cross section or move its control points to create the bulge.

    To move individual control points, use the Control Point selection level.
10. Enter **80.0** in the Influence field.

    If the bulge angle is for a biceps bulge when the forearm and upper arm create a 90-degree angle, the bulge begins to appear at 10 degrees.

To copy all Bulge angles from one link to its opposite:

1. Go to the Bulge sub-object level.
2.  In the Selection Level group, turn on Link. In a viewport, select a link to copy.
3. Choose Entire Link from the Current Bulge Angle drop-down list.
4. In the Cross Section Parameters rollout, click Copy.
5. In a viewport, select the opposite link.
6. In the Cross Section Parameters rollout, click Paste, and then click Mirror.

   All the bulge angles from the first link are pasted to the opposite link, then mirrored.

To choose a specific bulge angle for editing:

1. Click the arrow of the Current Bulge Angle drop-down list.

   The full list of bulge angle names appears for the current link.
2. Click the name of the bulge angle you want to edit.

   The viewports update to show the cross sections associated with the newly selected bulge angle.

To change a bulge angle value:

1. Make sure the bulge angle's name is displayed in the Current Bulge Angle drop-down list.
2. Move the time slider to a frame that displays the angle you want to record, or use transforms to change the angle between the active link and its child link.
3.  In the Bulge Angle Parameters group, click Set Bulge Angle.

To delete a bulge angle:

1. Click the arrow of the Current Bulge Angle drop-down list.

The full list of bulge angle names is displayed.

2. Click the name of the bulge angle you want to delete.
3.  Click Delete Bulge Angle.

   You can't delete the default bulge angle: a link must always have at least one bulge angle defined.

To use Select Nearest Bulge Angle:

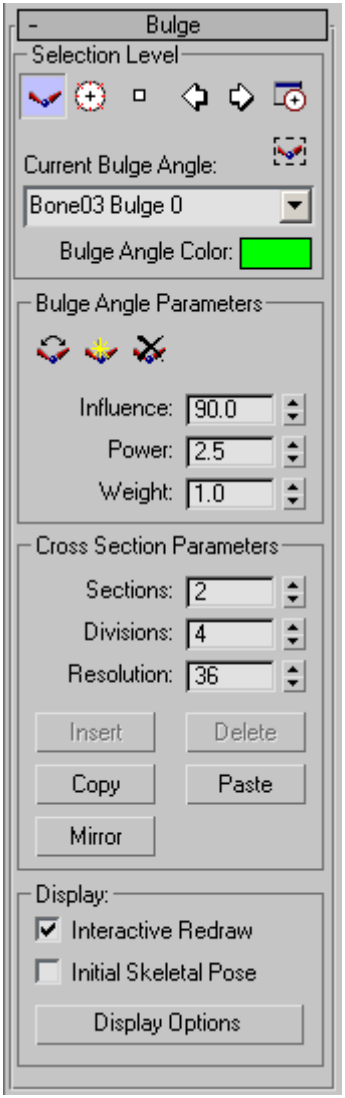
The Select Nearest Bulge Angle button can help you identify the bulge angle that has the greatest effect at the current pose.

1.  In the Selection Level group, click Select Nearest Bulge Angle.
2. Scrub the time slider or rotate the joint.

   The bulge angle named in the drop-down list changes to show the most influential bulge at each skeleton pose.

Interface

Controls for the Bulge sub-object level are on the Bulge rollout.



Selection Level group



 **Link—**Turn on to select links in the viewports. The Copy, Paste, and Mirror commands are enabled for links.

 **Cross Section—**Turn on to select and edit cross sections in order to "bulge" the mesh. The Insert, Delete, Copy, Paste, and Mirror commands are enabled for cross sections.

 **Control Point—**Turn on to edit cross section control points.

 **Previous and Next—**Click to select the next or previous link, cross section, or control point, depending on the selection level.

 **Bulge Editor—**Click to display the *Bulge Editor (page 2–1114)*, which lets you create and edit bulge angles using the schematic Cross Section and Profile views

 **Select Nearest Bulge Angle—**Turn on to select the bulge angle nearest to the current joint angle. If a joint bends over time, this lets you use the time slider to select a bulge angle. If you click Play, you can see the Current Bulge Angle field change to reflect the nearest bulge angle to the current angle of the selected link and its child (if two or more bulge angles for the limb exist).



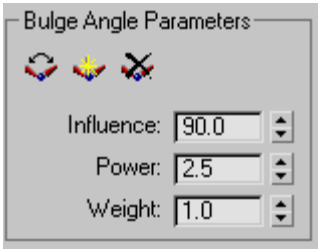
Current Bulge Angle (field and drop-down list)—Displays the current bulge angle. You can enter a descriptive name for a newly created bulge angle. When you click Insert Bulge Angle, this field displays the new bulge angle name. Use the drop-down list to select other bulge angles.

Entire Link—This choice in the Current Bulge Angle drop-down list selects all cross sections for all bulge angles in the active link. Use this to change cross section settings globally for a link. This also affects the Copy and Paste commands. While Entire Link is chosen, Copy Cross Section copies all cross sections for all bulges, and Paste Cross Section pastes all cross sections for all bulges.

Tip: Entire Link is useful for copying all the bulges from one arm or leg to another.

Bulge Angle Color—Click the color swatch to change the color of the current bulge angle. Giving each bulge angle a different color is a convenient way to distinguish between them.

Bulge Angle Parameters group



 **Set Bulge Angle—**Changes the angle value to the skeleton's current joint angle on the current bulge angle (visible in the Current Bulge Angle field in the Selection Level group). First use Move or Rotate in a viewport to change the angle between the selected link and its child, then click Set Bulge Angle.

 **Insert Bulge Angle—**Adds a new bulge angle for the selected link. The Current Bulge Angle field displays a bulge angle name with the number incremented. You can enter a descriptive name for a new bulge angle, such as **Arm at 90**.

By default, one Bulge Angle is created in the Figure mode pose when Physique creates links. Only one additional bulge angle is needed to bulge the mesh. You can create more bulge angles for further control, if you want.


Delete Bulge Angle—Removes the current bulge angle from the selected link.

You cannot delete the default bulge angle.

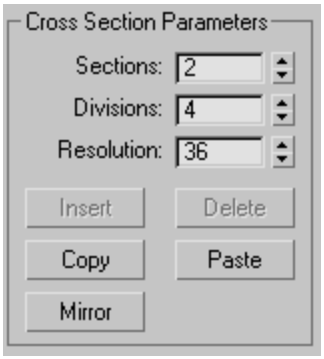
Influence—The range of angles through which the bulge influences the skin. Range =0 to 180. Default=90 degrees.

For example, if you've set a bulge angle for the joint at 90 degrees, an Influence value of 40 means that the bulge effect begins to appear when the rotating joint reaches 50 degrees (90 minus 40) or 130 degrees (90 plus 40).

Power—Controls how smoothly or abruptly the bulge takes effect. At 0.0, the bulge takes effect immediately, without interpolated easing. As values increase, the bulge eases in gradually. A value of 10.0 will bulge the mesh abruptly when the set angle is reached. Range=0 to 10.0. Default=2.5.

Weight—Increases the effect of the current bulge angle relative to the effect of any other bulges. Range=0.0 to 100.0. Default=1.0.

Cross Section Parameters group



Sections—Sets the number of cross sections for the selected link.

Divisions—Sets the number of control points around the cross section. By default, control points are evenly distributed around the circumference of the cross section.

Resolution—Sets the number of radial divisions around the cross section. Control points snap to the nearest resolution line.

Note: The buttons that follow, Insert, Delete, Copy, Paste, and Mirror, work according to which level is active in the Selection Level group.

Insert—Creates a new cross section or control point.

Delete—Deletes the active cross section or control point.

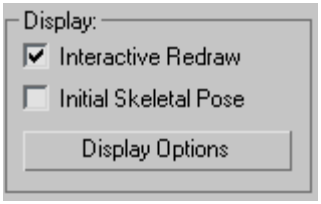
Copy—Copies the active link (with its bulge angles) or cross section.

Paste—Pastes a link (with its bulge angles) or cross section.

To copy and paste all the bulge angles from one link, choose Entire Link from the Current Bulge Angle drop-down list, and then click Copy. Select the opposite link and click Paste. Click Mirror to mirror the pasted link parameters.

Mirror—Mirrors bulge angles and cross sections on bulge angles. If necessary, you can then use Select And Rotate to orient the bulge angles and cross sections.

Display group



Interactive Redraw—Turn on to deform the mesh in real time while you edit cross sections in the viewports. Turn off to update the mesh only when you release the mouse. Default=on.

Tip: Turn off Interactive Redraw when working on a slow system, or with a dense mesh that takes a long time to compute.

Initial Skeletal Pose—When on, puts the mesh into its original pose; that is, the pose it was in when Physique was first applied. Default=off.

Display Options—Click to display the *Bulge Angle Display Properties dialog (page 2–1127)*, which lets you customize the viewport display of bulges.

Tendons Sub-Object

Select an object that has the Physique modifier applied to it. > Modify panel > Tendons sub-object level

While envelopes provide smooth skin deformations, tendons provide additional stretching in much the same way that actual human tendons do. For example, you might create pulling in the wrist (several joints away) when the fingers are moved. After adjusting the *envelope sub-objects (page 2–1130)* for good mesh deformation, you can use tendons to control the amount of skin stretching across multiple links.

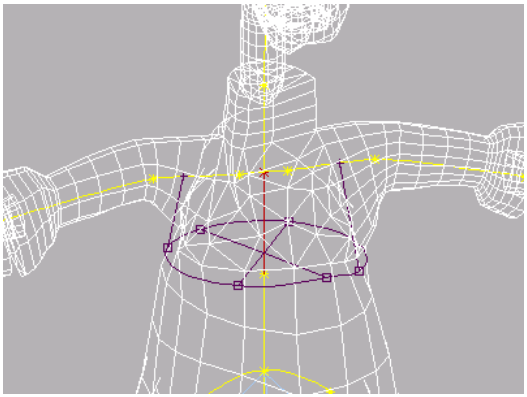
Workflow Example

To create a tendon from the spine to the upper arm, go to the Tendon sub-object level. Select a spine link. In the Insert Settings group, click to turn on Insert, and then click the selected link to position the tendons. A tendon cross section is added to the link. Turn off Insert.

In the Selection Level group, turn on Cross Section. Using the Rotate tool (on the Main toolbar), rotate the cross section so its control points are in useful locations. In the Tendon Parameters group, use the Radius spinner to scale the cross section radius so the control points fall close to the surface of the skin.

In the Selection Level group, turn on Control Point. Select a control point on the spine tendon's cross section, click Attach, and then click the arm link. A tendon stretches from the spine to the arm.

Finally, adjust the values of Pull, Pinch, and Stretch to refine the movement of the skin.



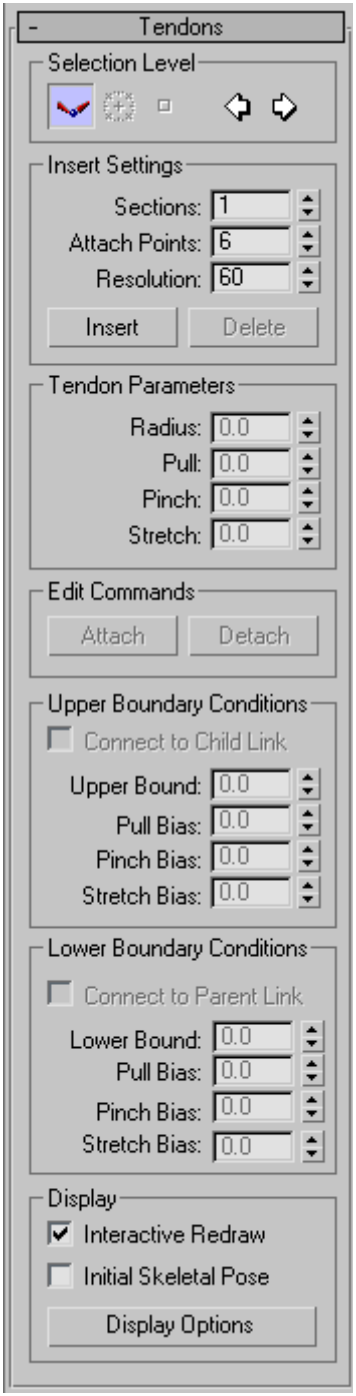
Tendons connecting upper chest to arms

Procedures

See *Creating Tendons (page 2–1096)*

Interface

Controls for Tendon sub-objects are on the Tendons rollout.



Selection Level group



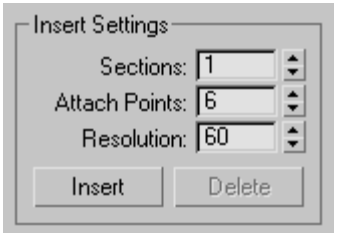
Link—Click to select a link.

Cross Section—Click to select and edit tendon cross sections.

Control Point—Click to edit control points on tendon cross sections.

Previous and Next—Click to go to the previous or next link, tendon cross section, or control point (depending on the active selection level).

Insert Settings group



Sections—Sets the number of base cross sections for the tendon. Default=1.

Attach Points—Sets the number of attach points around the tendon's base cross section. Default=6.

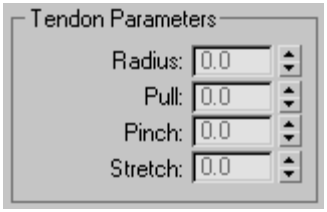
Resolution—Sets the radial resolution around the base cross section. Attach points are constrained to one of the radial resolution lines. Default=60.

Insert—Turn on to insert a new cross section on the selected link, or to insert a control point on a cross section. The action depends on the active selection level: Link, Cross Section, or Control Point.

Delete—Click to delete the selected cross section or control point.

The action depends on the active selection level: Link, Cross Section, or Control Point.

Tendon Parameters group



Radius—Scales attach points relative to the center of the cross section. Default=The average circumference of the skin where the tendon base is located.

Pull—Defines the strength of pull *along the length of the link*. Default=1.0.

Pinch—Defines the amount of pinch *around the link circumference of the tendon base*. Default=1.0.

Stretch—Defines the amount of stretch *toward the attached link*. Default=1.0.

The Pull, Pinch, and Stretch values work together to control, along a particular dimension, the strength of the one link (the attached link; call it link B) on vertices assigned to another link (the link where the tendon cross sections are inserted; call it link A). Vertices on link A behave as if they were under the influence of link B.

Pull, Pinch, and Stretch can all range from -2.0 to 2.0. At the default value of 1.0, the tendon's effect is full strength. Reducing the value of these parameters to a value between 0.0 and 1.0 reduces the strength of the tendon's effect. At values below 0.0, the tendon deforms the skin in the opposite direction. Pull, Pinch, and Stretch values between 1.0 and 2.0 cause an exaggerated effect, which is not usually needed for realistic effects.

Tip: Setting Stretch less than 0.0 can be useful in regions where movement typically pulls the skin inwards, as in the collarbone area or the buttocks.

Tip: As a general rule, leave the tendon values at 1.0 when the tendon is attached to a nearby link. For example, leave them at 1.0 between the upper spine (at the chest) and the arms. Reduce the Pull, Pinch, and Stretch values to reduce the tendon's effect when the tendon is attached to a more distant link. For example, reduce them between a lower part of the spine and the arms.

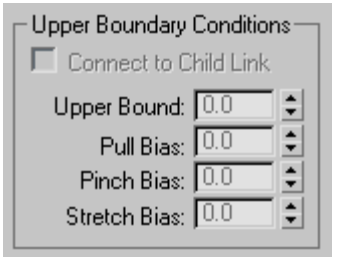
Edit Commands group



Attach—To attach a point, select the point, click to turn on Attach, then in a viewport, click a different link.

Detach—At the Control Point selection level, detaches the tendon of the selected control point. At the Cross Section selection level, detaches all of the cross section's attached tendons. At the Link selection level, detaches *all* attached tendons on *all* cross sections of the selected link.

Upper Boundary Conditions group



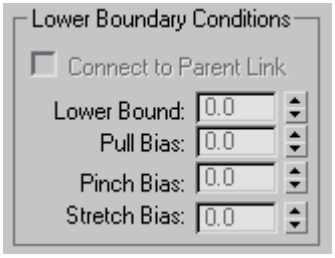
Connect to Child Link—When on, allows tendons to affect the child link. This lets you connect tendons across several links. Otherwise, there is a boundary between the links where no tendon effect occurs. Default=off.

Upper Bound—Sets the upper boundary overlap. Upper Boundary values greater than 1.0 affect the child link. Default=1.0.

Warning: Tendons that span several links can conflict with joint intersection parameters. In these cases, go to the *Link sub-object level (page 2–1135)*, and turn off Active to deactivate the joint intersection parameters for the joints that the tendon spans. (These controls are in the *Joint Intersections rollout (page 2–1140)*.)

Pull Bias, Pinch Bias, Stretch Bias—Set the upper boundary falloff effect for these tendon parameters. Values of 0.0 have no effect. Increasing the value shifts the Pull, Pinch, or Stretch effect toward the child. Default=0.5.

Lower Boundary Conditions group



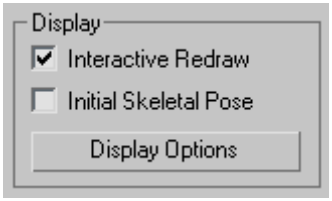
Connect to Parent Link—When on, allows tendons to affect the parent link. This lets you connect tendons across several links. Otherwise, there is a boundary between the links where no tendon effect occurs. Default=off.

Lower Bound—Sets the lower boundary overlap. Lower Boundary values less than 0.0 affect the parent link. Default=0.0.

Warning: Tendons that span several links can conflict with joint intersection parameters. In these cases, go to the *Link sub-object level (page 2–1135)*, and turn off Active to deactivate the joint intersection parameters for the joints that the tendon spans. (These controls are in the *Joint Intersections rollout (page 2–1140)*.)

Pull Bias, Pinch Bias, Stretch Bias—Set the lower boundary falloff effect for these tendon parameters. Values of 0.0 have no effect. Increasing the value shifts the Pull, Pinch, or Stretch effect toward the parent. Default=0.5.

Display group



Interactive Redraw—Turn on to deform the mesh in real time while you edit cross sections in the viewports. Turn off to update the mesh only when you release the mouse. Default=on.

Tip: Turn off Interactive Redraw when working on a slow system, or with a dense mesh that takes a long time to compute.

Initial Skeletal Pose—When on, puts the mesh into its original pose; that is, the pose it was in when Physique was first applied. Default=off.

Display Options—Click to display the *Tendon Display Options dialog (page 2–1128)*, which lets you customize the viewport display of tendons.

Vertex Sub-Object (Physique)

Select an object that has the Physique modifier applied to it. > Modify panel > Vertex sub-object level

Most often, you will want to use envelopes to correct the way skin behaves as the biped moves. However, you can override envelopes by manually assigning vertex properties. For example, you can remove the influence of inappropriate links from selected vertices. You can also change the weight distribution between links for a single vertex by using type-in weights.

Procedures

Tip: You don't have to work locally, one link at a time. You can use region selection to select groups of vertices or groups of links, working with areas of the body and its skeleton all at once.

To check vertex assignments:

1. Go to the Vertex sub-object level.
2. In the Vertex Operations group, click to turn on Select By Link.
3.  Make sure all three Vertex Type buttons are on.
4. In a viewport, click a link.

   Physique displays the vertices assigned to that link. If any vertices assigned to the link are out of place, reassign them to a different link, as described in the procedures that follow.
5. Repeat step 4 for other links in the skeleton.

To remove a link's influence on vertices:

For example, you might want to remove the influence of the index finger links from the vertices of the middle finger.

1. Go to the Vertex sub-object level.
2.  Turn on Select and select vertices in the viewports.
3. In the Vertex Operations group, turn on Remove From Link.
4. Click a link in a viewport.

   The vertices are removed from that link's influence.

   Note: You can click additional links to remove their influence as well.
5. In the Vertex Operations group, turn on Lock Assignments.

   You must use Lock Assignments to ensure that these envelopes will not influence these vertices if envelope parameters are adjusted later.

To remove deformable vertices from a link's influence:

1. Go to the Vertex sub-object level.
2.  Make sure all three Vertex Type buttons are on, and then in the Vertex Operations group, click to turn on Select By Link.

   Observe the assignments and determine which vertices are incorrectly influenced. Also note if the vertices are red (deformable) or green (rigid).

   Note: Blue vertices do not fall within the influence of any link and are thus assigned to the root.
3.  Turn on Select, and in a viewport, select only the out-of-place vertices.

   Note: Alternatively, you can use Alt+click to deselect the vertices that are in the correct place.
4.  In the Vertex Type group, click to turn off Rigid Vertices and Root Vertices. Leave Deformable Vertices on.
5. In the Vertex Operations group, click Remove From Link.
6. In the viewports, click each link that you want no longer to influence the selected deformable vertices.

To override vertex assignments manually:

This technique is used when envelope assignments are inappropriate, and you want to have specific vertices influenced by a specific link.

1. Go to the Vertex sub-object level.
2.  Make sure all three Vertex Type buttons are on, and then select the vertices you want to reassign.
3. In the Vertex Operations group, click to turn on Assign To Link.

4.  In the Vertex Type group, choose only the type of vertex assignment you want to use: deformable (red) or rigid (green.)
5. Choose No Blending from the Blending Between Links drop-down list.

   This setting will disregard the effects of blending envelopes and vertex weights, and let you assign the selection to any link manually.
6. In a viewport, click the link you want to influence the vertex.

   Note: When you assign deformable vertices, some vertices might turn blue; they are assigned to the root instead of the link you clicked. To assign these vertices as correctly deformable vertices, simply use Ctrl+click with the neighboring parent or child link. If the blue vertices now turn red, they are deformable.
7. In the Vertex Operations group, click Lock Assignments to preserve the manual assignment.

To make vertices rigid:

1. Go to the Vertex sub-object level.
2.  Turn on Select, and in a viewport, select the vertices you want to make rigid.
3. In the Vertex Operations group, click to turn on Assign To Link.
4.  Turn on Rigid Vertices (green) in the Vertex Type group, and turn off Deformable and Root Vertices.
5. Choose No Blending from the Blending Between Links drop-down list.

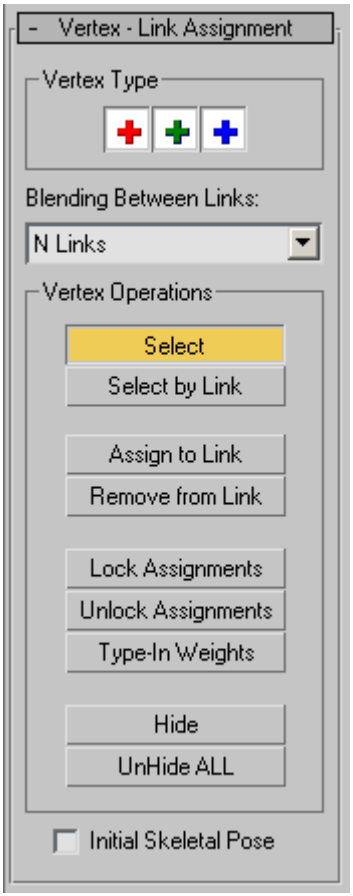
   This setting will disregard the effects of blending envelopes and vertex weights, and let you assign the selection to any link manually.
6. In a viewport, click the link these vertices are to be assigned to.
7. In the Vertex Operations group, click Lock Assignments.

   Small squares around the selected vertices indicate that the vertices are locked.

Interface

Controls for vertices in a mesh controlled by Physique are on the Vertex-Link Assignment rollout.

Note: This is *not* the same as the rollout used by the Physique Initialization dialog.



Vertex Type group



There are three vertex types, differentiated by color:

- **Red**—Deformable vertices that follow Physique's deformation spline.
- **Green**—Rigid vertices that do not deform but just follow the link they are assigned to.
- **Blue**—Vertices attached to the root node. Physique uses this color when it isn't sure which link to assign the vertices to. These vertices don't deform but follow the center of mass object. You should reassign blue vertices to be either rigid or deformable.

While a button of the appropriate color is turned on, you can select vertices of that color. Turn off a button to avoid selecting vertices of that color. You can't turn off all three buttons at once.

Blending Between Links drop-down list

This list provides several options:

- **N Links**—Vertices within all overlapping envelopes are influenced.

  While this option is active, vertices can be assigned *only* to the envelopes they fall within. Envelopes are used by Physique in this case to define the blended weight of each vertex.

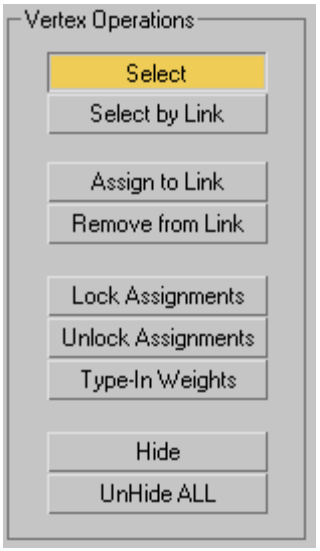
- **No Blending**—Vertices are influenced by only one link, as in character studio 1.

  This is useful as a general default if you are developing characters for a game engine that doesn't support blending. At the Vertex sub-object level it is useful to override blending weights defined by envelopes, and assign vertices to links as either Rigid or Deformable.

  For example, after selecting vertices of the skull and face, turn on No Blending, turn on Rigid vertices, turn on Assign To Link, and then click the head link. This will override the head envelope vertex assignments and assign these vertices as rigid to the head link. Use Lock Assignments to lock any changes (see below).

- **2, 3, 4 Links**—Vertices are influenced only by the specified number of closest links. This is useful when developing characters for a game engine with limited blending.

Vertex Operations group



Select—Selects vertices using the selection tools on the main toolbar.

Select by Link—Selects vertices by link.

Assign to Link—Assigns the selected vertices to a link.

To assign the currently selected vertices to a link, turn on Assign To Link and then click the link.

Assign To Link has a dependency on the envelopes defined for each link, and you must set Blending Between Links to "No Blending" to override

this. After vertices are assigned manually in this way, use Lock Assignments to keep them from being reassigned inadvertently during subsequent operations.

Note: You must select or assign a link by clicking the link; you can't use the Select Objects dialog to select or assign a link while you use Physique to work with vertices.

Remove from Link—Remove the selected vertices from a link.

Note: Select, Select By Link, Assign To Link, and Remove From Link will act only on vertices specified by the red, green, and blue Vertex Type ("+") buttons.

Lock Assignments—Locks vertex assignments.

Lock will prevent any changes from being made to the weights and blending presently assigned to the selected vertices. After *any* manual assignment of vertices, they should be locked. Before using type-in weights, vertices must be locked. When manipulating envelopes in difficult areas of the character, you might choose to lock vertices that are working well before changing the envelopes to affect other vertices.

Note: If additional data, such as more vertices, comes up the modifier stack, Physique notes the proximity of new vertices to the locked vertices. It then *locks* the new vertices based on proximity. Already locked vertices will not be unlocked.

Unlock Assignments—Unlocks vertex assignments.

Click Unlock Assignments on selected locked vertices prior to reassigning them to another link, or changing their blending settings.

Type-In Weights—Displays the *Type-In Weights dialog (page 2–1128)*, which you use to enter a weight for selected locked vertices; only the weight for locked vertices can be changed.

To use this feature: select one or more vertices, then lock them. Click this button to display the Type-In Weights dialog, select a link from the drop-down list, and change the weight of that link to position the vertices (vertex positions update in the viewports as the weight changes).

You can use type-in weights for stubborn vertices that fall in blended regions between two or more envelopes. Subtle adjustments can be made to the vertex weights of specific links that would be difficult to achieve by adjusting envelopes alone.

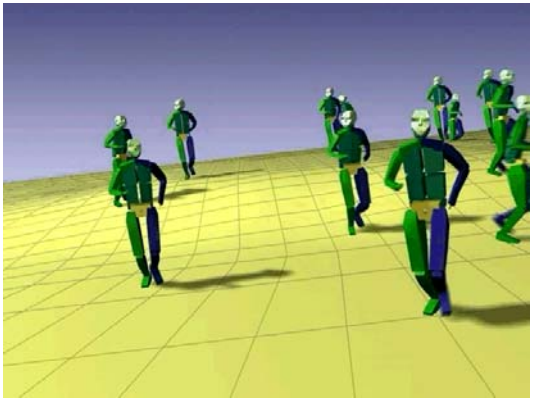
Tip: Type-In Weights is useful mainly for correcting flaws in low- to medium-resolution meshes. On a high-resolution mesh, adjusting envelopes should be used to correct deformation.

Hide—Hides the selected vertices

Unhide All—Unhides all vertices.

Initial Skeletal Pose—When on, puts the mesh into its original pose; that is, the pose it was in when Physique was first applied. Default=off.

Crowd Animation



The crowd-animation system in **character studio** is designed to simulate the behavior of real-life crowds. A crowd simulation emulates real-life situations by animating *delegates (page 3–928)* (helper objects that act as representatives). You

give the delegates overall guidelines on how to behave, and the crowd simulation calculates their motion.

You set these guidelines by assigning *behaviors (page 3–915)* to delegates. A behavior specifies a particular type of activity, such as moving toward an object in the scene, avoiding obstacles, following a path or surface, and so forth. You can combine multiple behaviors to create a rich and complex crowd simulation automatically. You can then link objects to delegates to complete the animation.

A crowd simulation can be used to animate bipeds, or to drive the use of mesh animation on objects linked to delegates. In addition, you can use *cognitive controllers (page 3–921)* to tell delegates how to behave in varying circumstances.

These topics provide further explanation of the principles of crowd animation:

Creating a Crowd System (page 2–1155)

Creating Crowd Helpers (page 2–1157)

Adjusting Delegate Parameters (page 2–1159)

Assigning Behaviors (page 2–1159)

Directing Delegates (page 2–1162)

Obstacle Avoidance (page 2–1164)

Changing Delegate Orientation and Speed (page 2–1167)

Solving the Simulation (page 2–1168)

Linking Objects to Delegates (page 2–1169)

Cognitive Controllers (page 2–1170)

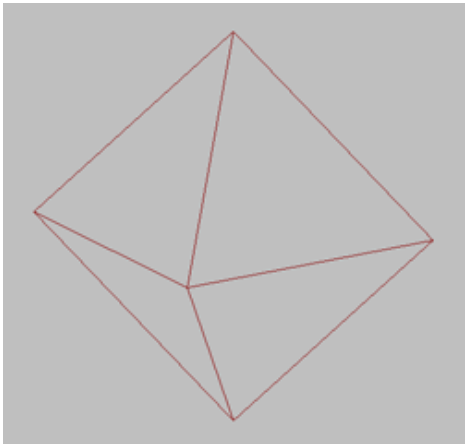
Biped Crowds (page 2–1172)

Non-Biped Crowds (page 2–1179)

Creating a Crowd System

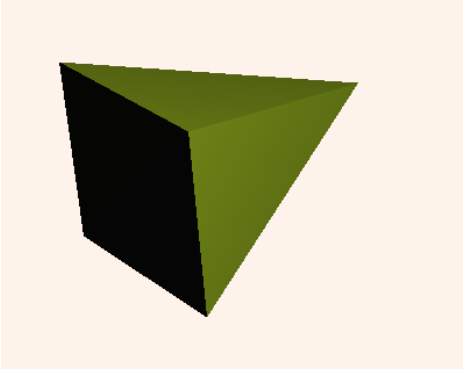
The following is a basic procedure for creating and using a crowd system:

1. Add a *Crowd helper (page 2–1157)* object.



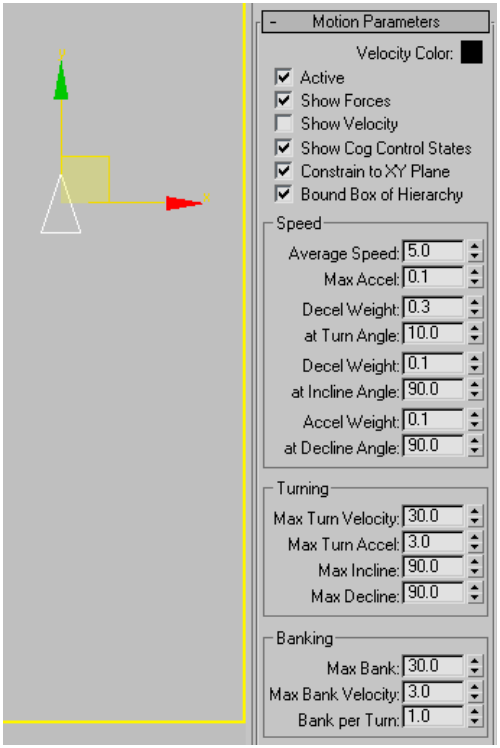
The crowd helper is the controlling object for the entire crowd simulation.

2. Create a *Delegate helper (page 2–1157)* object.



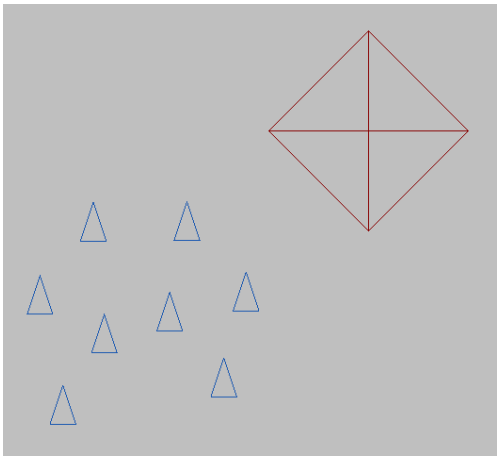
This is a prototype for your crowd; a representative member.

3. On the Modify panel > *Motion rollout (page 2–1183)*, set appropriate speed and turning limits for the delegate.

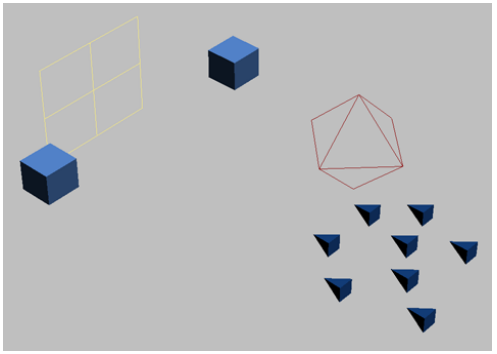


This is analogous to defining how a particular animal or object moves. Is it a bird, a fish, a slug, or an airplane?

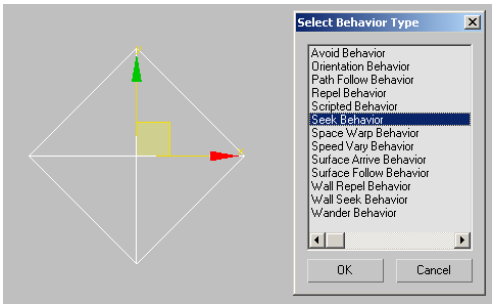
4. Make clones of the delegate, either with 3ds Max clone tools, or by *scattering delegates (page 2–1157)*.



5. Create 3ds Max objects to be used with behaviors, such as grids or objects to seek or avoid. These objects often correspond to objects in your scene. For example, you could create a grid in an open doorway to attract delegates, or a box at an obstacle to repel delegates.



6. Add one or more *behaviors (page 2–1159)*. Modify each behavior's settings to be appropriate for the members of your crowd.



7. Solve the simulation.
8. Watch the solved simulation. If it needs correction, make changes to behaviors or delegate parameters. Solve again, and continue until the simulation works as intended.
9. Create the objects that will follow the delegates, and *align and link them to delegates (page 2–1169)*.

Creating Complex Simulations

The process described above creates a simple crowd simulation. You can also use other tools to further control the simulation:

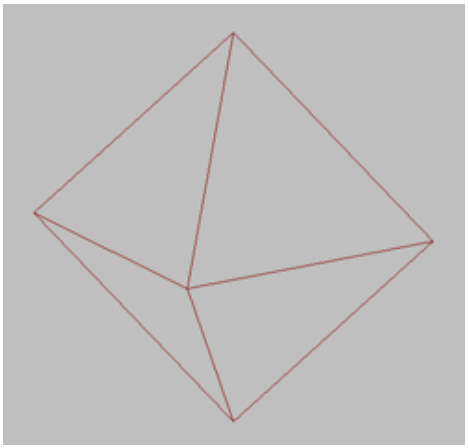
- Use *cognitive controllers (page 2–1170)* to cause delegates to switch behaviors based on their proximity to objects in the scene, time elapsed in the simulation, and other factors.
- Use *motion synthesis (page 2–1179)* to cause different portions of the linked object's animation to be used based on delegate speed, rotation, and other factors.
- Use bipeds and a motion flow network to animate the bipeds with delegates. See *Biped Crowds (page 2–1172)*.

Creating Crowd Helpers

The crowd system in **character studio** uses two helper objects: Crowd and Delegate. The first step in creating a crowd simulation is the creation of these helpers. The crowd helper serves as the command center for setting up and solving crowd simulations, while delegates provide stand-ins for animated objects. The crowd helper controls are used to animate the delegates, then later you link objects to delegates to create the finished scene.

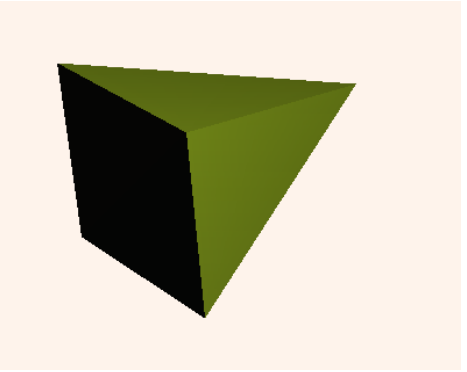
Crowd Helper

A crowd helper object appears in viewports as blue tetrahedron. The location and size of the crowd helper does not affect the simulation. For the easiest workflow, create the crowd helper with a relatively large size, and place it where you can easily select it. You'll rarely need more than one Crowd helper object per scene.



After you create delegates, you use the crowd helper to clone delegates and distribute the clones, add behaviors, apply behaviors to delegates, link delegates to animated objects, and much more.

Delegate Helpers



The Crowd system in **character studio** uses Delegate helper objects as intermediaries between the crowd simulation and the influenced objects. A delegate is a non-rendering pyramid-shaped object. Its apex points in the direction the crowd system uses as "forward" when it moves the delegate around.

Use one delegate per object to be animated by the crowd system. You can add delegates one by one, or use any of the standard 3ds Max methods for cloning objects, including Shift+clone and the Array function. However, the Crowd object offers a convenient *Scatter Objects (page 2–1189)* function that lets you clone delegates and distribute the clones over a surface or within a volume, with options for orientation and scaling.

Procedures

To create a Crowd helper object:

1. On the Create panel, click Helpers.
2. Click Crowd, and drag the helper object to a convenient size.

   Tip: Make the crowd helper object large enough so you can easily select it in two or more viewports. The size of the helper does not affect the crowd simulation in any way.

To create a Delegate helper object:

1. On the Create panel, click Helpers.
2. Click Delegate. Drag in a viewport to set the delegate position, and then drag vertically to set the size.

   Tip: The delegate always points "up" in the viewport you add it in. So if you want it to point "forward," that is, toward the positive Y direction in the World coordinate system, you should create it in the Top viewport.

To clone and scatter delegates:

1. Create a Crowd helper and a Delegate helper.
2. Decide how to distribute the delegate clones, and create an object to define how the clones are to be positioned: a grid object, a primitive box or sphere, a shape, or any object to serve as a surface. Alternatively, you can choose to distribute clones in a circular area without the use of a distribution object.
3. Select the Crowd object and access the Modify panel.
4. Click the Scatter Objects button to open the *Scatter Objects dialog (page 2–1189)*.
5. On the Clone tab, set the number of clones to create, and select the object to clone.
6. Click Generate Clones.

   This produces the specified number of clones in the same location. To vary positions, orientations, and/or sizes, proceed with the following steps.
7. On the Position tab > Placement Relative To Object group, choose a distribution option. Click the None button, and then select the distribution object you created earlier.

   Alternatively, to distribute delegates in a circular area, choose Placement In Area group > In Radial Area and set a center and radius.

8. Click Generate Locations.

   This distributes the clones.

9. On the Rotation tab, choose which of the cloned object's local axes are to look forward and up. Optionally, specify source and target objects for the clones' orientations, as well as limits for randomized deviations from the calculated orientation.

10. Click Generate Orientations.

11. On the Scale tab, for each axis, specify Average, Deviation, and, optionally, Same As settings and a random seed.

12. Click Generate Scales.

    At this point, you're basically finished. However, to create a series of randomized positions, orientations, and/or sizes for clones, follow the next two steps and repeat as necessary.

13. On the All Ops, turn on Positions, Rotations, and/or Scales. To vary the randomized positions, rotations, and/or scales, turn on the corresponding Inc Seed check boxes.

14. Click the Scatter button.

    You can combine the cloning, position, rotation, and scale functions in one operation by setting the options on their respective tabs without clicking the Generate buttons, and then using the All Ops tab to apply any or all scatter operations simultaneously. Also, you can use the All Ops tab > Select Objects To Transform to specify any objects to scatter; not just clones.

See also

Crowd Helper Object (page 2–1187)

Adjusting Delegate Parameters

Delegate parameters define the nature of their motion. You can change delegate parameters in one of two ways:

- To adjust parameters for an individual delegate, select the delegate, access the Modify panel, and adjust parameters. For descriptions of parameters, see .
-  To adjust parameters for several delegates at once, select the crowd helper object, access the Modify panel, and click Multiple Delegate Editing to access the *Edit Multiple Delegates dialog (page 2–1197)*.

The speed parameters let you describe a delegate's average speed and maximum acceleration. You can also define how much a delegate slows down as it turns or goes upward, and how much it speeds up when it goes down.

Turning parameters let you indicate how quickly a delegate can turn, and how much it can turn upward and downward. And banking parameters describe how much and how quickly a delegate banks and its banking limit.

All of these parameters work together to describe different types of creatures. A small fish, for instance, can turn more quickly than a large bird. A fish's speed would not be effected when the fish traveled up or down, but a bird's speed might be altered by its upward or downward direction.

See also

Crowd Helper Object (page 2–1187)

Assigning Behaviors

In the real world, different crowds exhibit diverse behaviors, and even members of the same crowd can conduct themselves in various ways. Included

with the **character studio** Crowd system is an assortment of behaviors that let you simulate a range of crowd activities.

Behaviors let you assign procedural activity types to *delegates*, which in turn affect objects linked to delegates. You can associate any number of behaviors with each *crowd object (page 2–1187)*, and then link delegates and teams of delegates to each behavior. A specific behavior assigned to a Crowd object belongs only to that crowd; it cannot be assigned to any other crowds.

Following is a list of available behaviors:

- *Avoid Behavior (page 2–1211)*—Prevents delegates from colliding with objects in the scene, or with each other. Avoidance can use any combination of turning, braking/stopping, repelling, and vector field.
- *Orientation Behavior (page 2–1214)*—Applies a fixed orientation or orientation range to delegates, so they face a specific direction instead of toward the destination. You can specify orientation in absolute terms, or relative to the direction the delegate currently faces.
- *Path Follow Behavior (page 2–1216)*—Restricts motion to a spline or NURBS curve; options include back-and-forth patrol-type movement.
- *Repel Behavior (page 2–1218)*—Forces delegates to move away from a target.
- *Scripted Behavior (page 2–1220)*—Uses MAXScript to specify behavior.
- *Seek Behavior (page 2–1220)*—Moves delegates toward a target or targets.
- *Space Warp Behavior (page 2–1221)*—Uses any dynamics-oriented space warp to control movement, including wind and gravity. Vector Field, a crowd-specific space warp that lets delegates avoid irregularly shaped objects while following their contours, is included with **character studio**.
- *Speed Vary Behavior (page 2–1222)*—Lets delegates change speed for more realistic movement.
- *Surface Arrive Behavior (page 2–1223)*—Lets delegates move toward and land on a surface, with custom speed and acceleration parameters.
- *Surface Follow Behavior (page 2–1226)*—Delegates move along a surface, which can be animated. Also, you can specify whether the delegates are to move straight ahead or skirt hills and depressions.
- *Wall Repel Behavior (page 2–1227)*—Uses a grid to repel delegates; ideal for keeping objects inside an enclosed, straight-sided room.
- *WallSeek Behavior (page 2–1229)*—Uses a grid to attract delegates. You can use this as a doorway for crowd-controlled bipeds to walk through.
- *Wander Behavior (page 2–1231)*—Induces a realistic semi-random movement for characters such as shoppers at a mall.

Tip: If you need custom behaviors, you can create your own with scripting.

The first time you add a behavior to the scene, a new rollout appears for the behavior below the Setup rollout. This rollout lets you change settings for the behavior. Certain behaviors, such as Seek and Avoid, let you specify "target" objects.

Tip: Behaviors are assigned generic names by default. It's a good idea to rename them with more meaningful descriptors; for example, "Seek Ball" or "Avoid Fire".

To display the rollout for a different behavior in the scene, choose it from the drop-down list in the Behaviors group. To see the controls available in the rollout for a behavior type, follow the link from its entry in the above list.

In addition to the controls available in behavior rollouts, you can use the *Behavior Assignments*

and Teams dialog (page 2–1200) to turn behaviors on and off (with the Active check box), and for all behaviors except Avoid, Orientation, and Surface Follow, you can set and animate Weight. The Active status is animatable for all behaviors.

Using Behaviors

To use a behavior, you apply it to a delegate or a team of delegates using the *Behavior Assignments and Teams dialog (page 2–1200)*. In this dialog, each assignment of a behavior to a delegate is given a weight. You can modify and/or animate these weights to influence the simulation.

Behavior assignment weights can profoundly effect a simulation. When applying two or more behaviors to the same delegate, the weights define the relationship between the behaviors, making one more or less powerful than the other. One way to visualize a behavior assignment weight is to examine the behavior's force vector during a crowd simulation. The vector's length indicates the behavior's weight upon the delegate.

Each behavior has its own parameters which appear in the Behavior rollout, available in the Crowd object's Modify panel. These parameters describe how the behavior works, and can sometimes contribute to the behavior's strength as well. For instance, Seek, Repel, Wall Seek, and Wall Repel, all have specific volumes of influence. Outside these volumes they have no effect and essentially have a weight of zero. This rollout lets you specify whether or not you wish to see behavior's force vector dynamically displayed during a Crowd simulation, and what color that vector should be.

When working with the Crowd system, it is critical to play with behavior assignment weights, as well as each behaviors' parameters. Typically, you run the simulation repeatedly, changing the weights and parameters to get the desired result.

A few behaviors cannot be weighted. These are Avoid, Surface Follow, and Orientation. Avoid and Surface Follow take over after all of the other behaviors have been applied to a delegate. They can take stringent measures to affect the delegate, possibly overpowering other behaviors in order to meet their constraints. Orientation simply sets the delegate's facing direction. It cannot be weighted and does not apply a force.

Behavior Tips

A few helpful things to know about behaviors in **character studio**:

- You can create conditional behavioral systems with Crowd's Cognitive Controller feature. This uses the MAXScript scripting language to determine when to effect a transition from one behavior to another; we've provided a number of sample scripts for you to learn from and adopt to your own simulations in *Cognitive Controller Editor (page 2–1206)* and *State Transition Dialog (page 2–1208)*.
- The Behavior rollout appears immediately after the Crowd object > Setup rollout in the Modify panel. However, it doesn't show up until you've added at least one behavior to the crowd object.
- The Crowd panel displays only one Behavior rollout at a time. To access a different one, choose its name from the drop-down list at the bottom of the Crowd object's Setup rollout.
- As with most scene entities in 3ds Max, it's a good idea to give behaviors custom names, such as "Seek Doorway" or "Follow Hilly Surface." You do this by clicking the behavior's name in the Setup rollout and entering a new one from the keyboard.
- The default behavior settings may not always give the ideal results. The optimal settings depend vary with the particulars of your simulation setup; in many cases, if not most, you'll need to experiment with the settings to

get the results you want. In some cases, you might need to animate settings as well.

One particularly useful feature of the delegate is its ability to display, using colored vectors, the strength and direction of the various forces acting upon it during solution of the crowd simulation. Each force can have a unique, identifying color. For example, the Seek behavior uses green by default, while the Wander behavior uses aqua. You can change these colors to any you like.

If a simulation isn't proceeding as expected, you can debug it by observing the vectors during the solution. And if the solution occurs too quickly, you can use the Step Solve feature to solve the simulation one frame at a time.

Procedure

To use assign behaviors to delegates:

This example shows how to create a basic crowd simulation with delegates and behaviors.

1. Run 3ds Max or reset the program.
2. Add a Crowd object and one or more *Delegate objects* to the scene. In general, add delegates in the Top viewport so that they point forward. The Crowd object's location is immaterial.
3. Select the Crowd object and open the Modify panel.
4. On the Setup rollout > Behaviors group, click the New button.
5. In the Select Behavior Type dialog, click a behavior and then click OK to close the dialog.

   A Behavior rollout appears for the behavior you chose.
6. If the behavior requires a target object or objects, such as Seek, click the None button and then select an object, or click Multiple Selection and select several objects.
7. Change other behavior settings as necessary.
8. Create and modify additional behaviors as necessary.
9.  On the Setup rollout, click the Behavior Assignment button.
10. In the Assignment Design group, the two upper lists should each contain a single entry: the delegate on the left, and the behavior on the right. Select both items.
11. Click the New Assignment button to the right of the Assignment Design group. It's a vertical button with five right-pointing arrows.

    This adds the new assignment to the list in the Behavior Assignments group.
12. Accept the changes and the OK button to close the Behavior Assignments and Teams dialog.
13. On the Modify panel, scroll down to the Solve rollout and click the Solve button.

    Keys are created as follows: The delegate turns to point toward the sphere, banking as it turns, and then moves directly toward the sphere. When it reaches its target, it moves slightly beyond the sphere, and then repeats the turn-and-move motion until the end of the simulation. To prevent this, try starting with the two objects farther apart, or animating the sphere's position.

Directing Delegates

One of the purposes of behaviors is to move delegates in a particular direction over the course of the simulation. These behaviors can be used to direct delegates:

- The *Seek behavior (page 2–1220)* is one of the most commonly used behaviors in the crowd system. With this behavior, you can cause delegates to move toward a specific object. The seek object can be animated, and the delegates will follow it.

- The *Wall Seek behavior (page 2–1229)*causes delegates to seek a rectangular area. Compare with the Seek behavior, which causes delegates to seek an object's pivot point.
- The *Path Follow behavior (page 2–1216)* is also useful for pushing delegates in a specified direction. With this behavior, delegates follow a path but can "stray" from the path by a specified distance, creating slightly random motion along the path.
- With the *Surface Arrive behavior (page 2–1223)*, delegates seek one or more objects as they do with the Seek behavior, but you can cause delegates to stop when they reach a target.
- You can use the *Surface Follow behavior (page 2–1226)* to cause delegates to move along a surface.
- With the *Space Warp behavior (page 2–1221)*, you can use a vector field to push delegates around the field to avoid obstacles. See *Obstacle Avoidance (page 2–1164)*.

After setting up a behavior, you must use *Behavior Assignments (page 2–1200)* to assign the behavior to a delegate or team.

Procedures

To use the Seek behavior:

1. Select the crowd helper and access the Modify panel.
2. In the Setup rollout > Behaviors group, click New. Choose Seek Behavior from the pop-up list.

   The Seek Behavior rollout appears below the Setup rollout.
3.  In the *Seek Behavior rollout (page 2–1220)*, click None to add one seek object, or click Multiple Selection to designate more than one object for the delegates to seek.
4. Change the default settings as desired.

To use the Wall Seek behavior:

1. On the Create panel, click Helpers, and click Grid. Create a grid in the scene.

   Tip: For best results, do not use Mirror to copy a grid to be used with behaviors. Use Shift+Clone instead.
2. Move and rotate the grid to position it.

   Tip: The Wall Seek behavior attracts delegates to the grid in the direction of grid's local Z axis. With the grid still selected, use the Local coordinate system to see the direction of the Z axis in viewports; the axis arrow points in the direction of the positive Z axis. Checking the Z-axis direction will speed the process of setting up the behavior.
3. Select the crowd object, and add a Wall Seek behavior.
4. In the Wall Seek Behavior rollout, click None, and pick the grid.
5. Choose Positive Axis to cause the delegates to be attracted to the Z-axis side of the grid. You can also choose Negative Axis to atrract delegates to the opposite side of the grid, or Both to attract them to both sides.
6. If you want delegates to be attracted to the grid only when they are within a specified area in front of the grid (rather than when they are anywhere in the scene), turn on Use Distance. Use the Inner Distance and Outer Distance parameters to set the area in which delegates will be attracted to the grid. Turn on Display Distance to see the distance in viewports.

To use the Path Follow behavior:

1. On the Create panel, click Shapes. Create a shape to be used as a path for delegates.
2. Add a Path Follow behavior to the Crowd object.

3. In the Path Follow Behavior rollout, click the None button and pick the shape for the delegates to follow.
4. Set the Radius parameter to the number of units by which delegates can stray from the path. Change other settings as desired.

To use the Surface Arrive behavior:

1. Add a Surface Arrive behavior to the Crowd object.
2. Add an object or objects to serve as the target surface to the scene.

   Note: If you use multiple objects, delegates will arrive at the surface of the closest object.

3.  In the Surface Arrive Behavior rollout, click None or Multiple Selection to designate the one or more target objects.
4. Use the settings in the Location group to determine where the delegate will stop when it reaches a target object.
5. Use the settings in the Approach group to determine how the delegate will behave when it approaches the target.

To use the Surface Follow behavior:

1. Add a Surface Follow behavior to the Crowd object.
2. Add an object or objects to serve as the follow surface to the scene.

   Note: If you use multiple objects, they must intersect to form a contiguous surface. Each delegate will move to the closest surface, follow it to the next closest that it encounters, and then start following that one, and so on.

3.  In the Surface Follow Behavior rollout, click None or Multiple Selection to designate the object or objects whose surface(s) the assignees are to follow.
4. If you like, use the settings in the Projection Vector group to force delegates to move in a specific direction.

Obstacle Avoidance

An important part of crowd behavior is avoidance of obstacles. Think of an obstacle as anything that impedes a crowd member's progress. Examples of obstacles include walls, telephone poles, and fences, as well as other crowd members. Encountering such objects can cause avoidance behavior, which consists of any combination of slowing down, turning, and stopping.

There are several ways to implement avoidance in **character studio**, including:

- The *Avoid behavior (page 2–1211)* is one of the most commonly used behaviors. Use this behavior to cause crowd members to avoid other crowd members, or spherical objects in the scene. It works by creating a spherical *volume of avoidance* around the avoided object, so it doesn't accommodate irregular objects.

  The Avoid behavior is unlike any other behavior in Crowd. After all the other behaviors exert their forces on the delegates, Avoid takes over and has the power to turn, slow down, and even stop a delegate in order to make it avoid an obstacle.
- Use the *Wall Repel behavior (page 2–1227)* to cause crowd members to avoid broad, flat objects such as walls and fences. You can set a maximum distance for the repel effect, and describe the rate at which the force away from the wall increases as a delegate approaches the wall.

  Unlike the Avoid behavior, which can stop or slow down a delegate, Wall Repel simply exerts a force on the delegate to turn it away from the wall. It does not guarantee wall avoidance. You must work with it's distance

and falloff parameters, as well as its weight in the Assignments and Teams dialog, to control its strength.

- Use the *Repel behavior (page 2–1218)* to cause crowd members to turn away from an object. It works exactly like Wall Repel except that it uses a spherical volume rather than a plane. You can set a maximum distance for the repel effect, and describe the rate at which the repel force increases as the delegate approaches the obstacle.

  Repel exerts a force on the delegate to turn it away from the obstacle. It does not guarantee avoidance. You must work with its Distance and Falloff parameters, as well as its weight in the *Behavior Assignments and Teams dialog (page 2–1200)*, to control its strength. Repel can be used instead of the Avoid behavior as a simple avoidance technique for non-terrestrial creatures such as fish or birds.

- Use a *Vector Field (page 2–1241)*. This is a special type of space warp that crowd members can use to move around irregular objects such as a curved, concave surface. You can use the Vector Field space warp in conjunction with the Avoid behavior to make delegates slow down when they approach a complex object, and then go around it. This guarantees that the delegate will not pass through the obstacle's surface.

  You can also use the Vector Field space warp in conjunction with the *Space Warp behavior (page 2–1221)*. This simply exerts a force on the delegate that mimics the contours of the object. It does not assure that the delegate will not pass through the surface of the obstacle. You can use a Vector Field with both the Space Warp and Avoid behaviors to combine their effects.

Procedures

To use the Avoid behavior:

1. Add an Avoid behavior to the Crowd object.
2. In the Avoid Behavior rollout, use the None button or the Multiple Selection button to designate the target object or objects to avoid.

   Tip: To make an entire team of delegates avoid each other, choose all delegates in the team. The crowd system will cause each delegate to avoid all others except itself.
3. Turn on Display Hard Radius to see the radial avoidance area in viewports. Adjust the Hard Radius to the appropriate size for your delegates.
4. Adjust the Brake Pressure to determine whether a delegate will slow down when it encounters an avoided object.
5. Adjust the Look Ahead parameter to determine how many frames ahead each delegate will look to determine the best way to avoid others.
6. Use *Behavior Assignments (page 2–1200)* to assign the behavior to a delegate or team.

To use the Wall Repel behavior:

1. On the Create panel, click Helpers, and click Grid. Create a grid in the scene.

   Tip: For best results, do not use Mirror to copy a grid to be used with behaviors. Use Shift+Clone instead.
2. Move and rotate the grid to position it.

   Tip: The Wall Repel behavior repels delegates from the grid in the direction of grid's local Z axis. With the grid still selected, use the Local coordinate system to see the direction of the Z axis in viewports; the axis arrow points in the direction of the positive Z axis. Checking the Z-axis direction will speed the process of setting up the behavior.
3. Select the crowd object, and add a Wall Repel behavior.
4. In the Wall Repel Behavior rollout, click None, and pick the grid.

5. Choose Positive Axis to cause the delegates to be repelled from the Z-axis side of the grid. You can also choose Negative Axis to repel delegates from the opposite side of the grid, or Both to repel them from both sides.
6. If you want delegates to be attracted to the grid only when they are within a specified area in front of the grid (rather than when they are anywhere in the scene), turn on Use Distance. Use the Inner Distance and Outer Distance parameters to set the area in which delegates will be attracted to the grid. Turn on Display Distance to see the distance in viewports.
7. Use *Behavior Assignments (page 2–1200)* to assign the behavior to a delegate or team.

To use the Repel behavior:

1. Add a Repel behavior to the Crowd object.
2. In the Repel Behavior rollout, use the None button or the Multiple Selection button to designate the object or objects that are to repel delegates.
3. Change the default settings as desired.
4. Use *Behavior Assignments (page 2–1200)* to assign the behavior to a delegate or team.

To add a Vector Field space warp:

Adding a Vector Field space warp object works the same as adding a Box geometry primitive.

1. On the Object Type rollout, click Vector Field.
2. Drag in a viewport to set the initial dimensions.
   - If using the Cube creation method, this sets all three dimensions simultaneously.
   - If using the Box creation method, release the mouse button, and then move the mouse vertically to set the height.
3. Click to create the space warp.

To use a Vector Field space warp with delegates:

This procedure presents general guidelines for using the Vector Field space warp with delegates in a crowd simulation.

1. Create an object to serve as an obstacle. This object must be an editable mesh or geometric primitive; it can have modifiers.
2. Add a Vector Field space warp.
3. Position and scale the space warp lattice so that it encloses the obstacle object.

   The lattice should be significantly larger than the object. The object should be located roughly at the lattice center.
4. In the Lattice Parameters rollout, increase the Length Segs/Width Segs/Height Segs settings so that the lattice segments intersect the object at reasonable intervals.

   To determine appropriate Segs settings for your obstacle objects, first examine the object complexity. If the obstacle has a lot of detail, and you want that detail reflected in the vector field, then you need a relatively high lattice resolution.
5. Click the Obstacle Parameters rollout > Compute Vectors group > Vector Field Object button (initially labeled "None"), and then in a viewport, click the object that will act as an obstacle in the crowd simulation.

   This specifies the obstacle object. The *range volume grid* appears on the object's surface as an olive-colored wireframe.
6. Increase the Obstacle Parameters rollout > Compute Vectors group > Range setting.

   As you increase this setting, you'll see the range volume grid expand. The range volume should enclose the space in which crowd members need to start turning in order to avoid the object.

7. Turn off Display group > Show Lattice and Show Range so that the vector field will be more easily visible when generated.
8. Turn on Display group > Show Vector Field.
9. In the Compute Vectors group, click the Compute button. This generates the vector field.

   Tip: To make the vector lines more evident, increase the Display group > VectorScale setting.

   The vectors appear as blue lines surrounding the obstacle object. One vector is computed for each lattice intersection within the range volume grid. Each vector matches the normal of the object at the point on the object's surface closest to the lattice point.

   The vector force falls off with distance from the object, as shown by the progressively shorter vector lines toward the grid perimeter.
10. Add *Crowd (page 2–1187)* and *Delegate* helper objects.
11. Select the Crowd object and open the Modify panel.
12. In the Setup rollout > Behaviors group, click New.
13. In the Select Behavior Type dialog that is displayed, choose Space Warp Behavior, and then click OK.
14. In the Space Warp Behavior rollout that appears, click the button (initially labeled "None"), and then in a viewport, click the Vector Field space warp.

    Tip: You might find it easier to use Select By Name to select the space warp.
15. In the Setup rollout, click the Behavior Assignments button, and use the *Behavior Assignments And Teams dialog (page 2–1200)* to assign your delegate or delegates to the space warp behavior.
16. Add any other objects or behaviors appropriate to the simulation.
17. Select the Crowd object, and then solve the simulation by clicking the Solve rollout > Solve button.
18. Fine tune the behavior associated with the Vector Field space warp by adjusting the *Lattice parameters (page 2–1242)* and *Obstacle parameters (page 2–1242)*.
19. Continue computing the vector field and then solving the simulation after each adjustment. In certain cases you might need to animate the vector field parameters to keep objects within the field.

Changing Delegate Orientation and Speed

You can use the *Orientation behavior (page 2–1214)* to change the direction in which a delegate moves, and the *Speed Vary behavior (page 2–1222)* to vary delegates' speed or make them stop moving altogether.

After setting up a behavior, you must use *Behavior Assignments (page 2–1200)* to assign the behavior to a delegate or team.

Procedures

To use the Orientation behavior:

1. Select the crowd helper and access the Modify panel.
2. In the Setup rollout > Behaviors group, click New. Choose Orientation Behavior from the pop-up list.

   The Orientation Behavior rollout appears below the Setup rollout.
3. To restrict the heading orientation with respect to the delegate's current heading, turn on Heading group > Relative. To restrict the

heading to a specific direction or range, leave Relative off.

4. To restrict the pitch orientation with respect to the delegate's current pitch, turn on Pitch group > Relative. To restrict the pitch to a specific direction or range, leave Relative off.
5. Change the other default settings as desired.

To use the Speed Vary behavior:

1. Add a Speed Vary behavior to the Crowd object.
2. Adjust the Center and Deviation as needed to control delegate speed. Change other settings as desired.

Solving the Simulation

After you have set up behaviors for delegates, you must solve the simulation to generate keyframes on delegates. You solve the simulation by selecting the crowd object and clicking Solve in the *Solve rollout (page 2–1232)*.

In order to generate the simulation as quickly as possible, delegate keys are saved *after* the solution is run, so there might be a pause between the end of the solution and return of control of the software to you. Also, any objects linked to delegates are hidden during the simulation.

You will most likely have to solve a simulation several times before it is correct. Watch the animation after the solution is calculated, correct behaviors or delegate parameters as necessary, then solve again.

You can speed up the solution calculation time by decreasing the frequency of keyframes or screen updates. You can also view the solution one frame at a time to help pinpoint trouble areas.

Procedure

To solve a simulation:

1. Set up a crowd simulation with a crowd helper, delegates and behaviors. See *Assigning Behaviors (page 2–1159)*.
2. Select the crowd helper.
3. On the Modify panel, in the Solve rollout, set the Start Solve and End Solve parameters to set the start and end frames for the solution.
4. Click Solve.

   Tip:  To use the keyboard shortcut for the Solve button, turn on the Keyboard Shortcut Override Toggle and then press the S key to run a solution.

   Wait a few moments while the solution is calculated. The progress bar at the bottom of the screen tells you the status of the solution process.

To speed up the solution time:

1. Under Save every Nth Key, set the Positions and Rotations parameters to a higher number, such as **5**.

   This causes the simulation to set delegate keyframes every five frames rather than at every frame.
2. Increase the Solve rollout > Display During Solve group > Frequency setting to a higher number, such as **100** .

   This will cause the display to update with the new delegate locations only every 100 frames.

To troubleshoot the simulation:

In a crowd simulation, it is not uncommon for a short portion of the animation to have a unique problem that doesn't appear on other frames. Use Step Solve to analyze short portions of your simulation when things don't go as expected.

1. Move the time slider to the frame at which you want to start solving one frame at a time.
2. Click Step Solve.

    To use the keyboard shortcut for the Step Solve button, turn on the Keyboard Shortcut Override Toggle and then press the T key to run a solution in step mode.

   You can start at any frame. Zoom in to examine the vectors of misbehaving delegates.
3. Press the spacebar to solve the simulation one frame at a time.
4. Press Esc to exit Step Solve mode.

Linking Objects to Delegates

You can link objects or bipeds to delegates to make them follow the simulated animation.

You can link objects to delegates with 3ds Max linking. Alternatively, the crowd system has a tool for automatically aligning and linking objects with delegates.

When you want to use bipeds with delegates, you associate bipeds with delegates rather than link them. This approach enables the bipeds to use a motion flow network to determine their motion during the simulation. For an explanation of how to set up a biped crowd simulation, see *Biped Crowds (page 2–1172)*.

Procedures

To link objects to delegates:

1. Set up a crowd simulation with a crowd helper and delegates. See *Creating Crowd Helpers (page 2–1157)*.
2. Create a series of objects to follow the delegates.

   Tip: In a later step, you will align each object's local Y axis to point in its corresponding delegate's forward direction. For this reason, you should take care to make sure your objects' local Y axes point in the direction you consider to be "forward". The easiest way to do this is to set up one object with its Y axis pointing in the forward direction, then clone the object.
3. Select the crowd helper, and access the Modify panel.
4.  On the *Setup rollout (page 2–1188)*, click Object/Delegate Associations.

   The *Object/Delegate Associations dialog (page 2–1196)* appears.
5. Click Add under the Objects list to add objects to link with delegates.
6. Click Add under the Delegates list to add delegates to link with the objects you added in the previous step.

   Each object in the Objects list will be associated with the delegate in the same position in the Delegates list. If necessary, reorder either list manually by highlighting entries and using the Shift Up/Shift Down buttons; these are the arrow buttons between the two lists.
7. Click Align Objects with Delegates to align each object with its corresponding delegate.

   The objects on the Objects list are moved and rotated to align with delegates. Each object's local Y axis points in the delegate's forward direction.
8. Click Link Objects to Delegates.
9. Click OK to exit.

To associate bipeds with delegates:

1. Set up a crowd simulation with a crowd helper and delegates.
2. Create several bipeds, one for each delegate.
3. Select the crowd helper, and access the Modify panel.

4.  On the *Setup rollout (page 2–1188)*, click Biped/Delegate Associations.

   The *Associate Bipeds with Delegates (page 2–1199)* dialog appears.

5. Click Add under the Bipeds list, and choose the bipeds to add to the list.

6. Click Add under the Delegates list, and choose the delegates to associate the bipeds with.

7. Turn on Set Delegates to Use Biped.

   This will turn on the Use Biped option for each delegate, which is a requirement for biped crowd simulations.

8. Choose Make Specified Associations to associate each biped with its corresponding delegate listed on the dialog. Alternatively, choose Associate Delegates With Closest Biped to associate each delegate with the biped nearest to it in the scene.

9. Click Associate to make the associations, and click Close.

   Unlike regular objects, bipeds are not aligned with delegates immediately. Each biped will align with its respective delegate when the simulation is solved.

   Note: In order for a biped crowd simulation to solve correctly, it must have a motion flow network for the bipeds to follow. See *Biped Crowds (page 2–1172)*.

Cognitive Controllers

You can use the cognitive controller feature to cause crowd members to change behaviors during a simulation depending on the circumstances. For example, a character could wander randomly until it comes within a certain distance of a target, at which point it could head straight for the target.

In technical terms, cognitive controllers let you influence crowd simulations with scripted conditionals, effectively implementing a form of artificial intelligence. You use the *Cognitive Controller editor (page 2–1206)*, a flowchart-style editor (much like the *Motion Flow Graph dialog (page 2–1045)*) to set up a network of behaviors and behavior combinations, known as states. Then you then apply MAXScript-based transitions that specify when delegates are to move from one state to another. Even with a relatively simple setup, you can create simulations that make your characters appear to be living, conscious beings, making decisions as they move through the scene.

You can find procedures describing various examples of cognitive-controller transitions in the *State Dialog (page 2–1207)* topic:

Testing a particle system parameter (page 2–1209)

Testing an object position

Testing an atmospheric property (page 2–1209)

Testing the distance between two objects

Testing a modifier parameter

Testing another delegate's behavior (page 2–1209)

Procedure

To set up and use a cognitive controller:

This procedure describes a typical setup routine for creating and using a cognitive controller. The procedure assumes basic knowledge of crowd simulation setup. For more information about crowd setup, see *Crowd Helper Object (page 2–1187)* and *Setup Rollout (page 2–1188)*.

1. Create a scene containing a crowd object and one or more delegates. See *Creating Crowd Helpers (page 2–1157)*.

2. Create at least two behaviors. See *Assigning Behaviors (page 2–1159)*.

3.  Open the *Cognitive Controller editor (page 2–1206)*.
4. Click the New button to create a cognitive controller.

   character studio gives the controller the default name of "Cognitive Controller." It's recommended that you give more descriptive names to cognitive controllers, such as "Seek/Wander". Do this by clicking on the name in the text box and editing it from the keyboard.

    Creating a new cognitive controller automatically places you in Create State mode.
5. Click in the editor window to create and place a state. Continue clicking in different places to add as many states as necessary.
6. Right-click a state to open the *State dialog (page 2–1207)*.
7. Again, it's recommended that you give more descriptive names to states, which you can do in the State dialog. Click the name (State or State#) in the text box and edit it from the keyboard.

   Next, define a behavior or behaviors for each state.
8. Click the Add button.
9. In the Select Behaviors dialog, choose one or more behaviors.

   If you choose multiple behaviors, you can specify different weights for each in the State editor. For example, you can combine a Seek behavior at full weight with a Wander behavior at half weight, so that the delegate will meander slightly as it seeks the target.
10. Close the Select Behaviors dialog, and then close the State editor.
11. Repeat steps 6–10 as necessary to define behaviors for the other states in the controller.

    Next, use Create Transition to define the sequence of states during the simulation.
12. Decide on the sequence in which the states are to occur.
13.  Click the Create Transition button.
14. Drag a line from one state to the next in the order that they are to execute. Click a state to create a transition from itself to itself.

    A transition arrow appears, pointing from the "source" state to the "destination" state.

    Each state can have any number of incoming and outgoing transitions. Specify different transition conditions for each to create as complex a state diagram as necessary.

    Next, use the State Transition dialog to define a conditional for each transition.
15. Right-click a transition line.
16. In the *State Transition dialog (page 2–1208)*, enter the name of the transition condition, and then click the Edit MAXScript button.
17. Use the MAXScript editor window to enter or load a script that defines the condition or conditions under which the transition is to occur.

    Typically, this is a function that tests a condition and returns 1 (if true) or 0 (if false).

    Note: See this topic in the online *User Reference* for sample MAXScript code for this function.
18. Use the State Transition dialog to set other parameters such as priority and duration.
19. Use the *Behavior Assignments and Teams dialog (page 2–1200)* to assign the cognitive controller to delegates or teams.

    Note: Crowd doesn't let you use multiple cognitive controllers with a delegate. You can assign them, but when you solve, the software notifies you that it will use only the first assigned cognitive controller.

See also

Cognitive Controller Editor (page 2–1206)

State Dialog (page 2–1207)

State Transition Dialog (page 2–1208)

Using Motion Synthesis

You can create advanced, complex crowd simulations in **character studio** with *motion synthesis*, which lets the software adjust the simulation results dynamically to account for differing conditions. Two different forms of motion synthesis are available: one for non-bipedal crowds, and a second for the more exacting requirements of biped crowds.

See also

Biped Crowds (page 2–1172)

Non-Biped Crowds (page 2–1179)

Biped Crowds

The biped crowd is a special case of crowd simulation necessitated by the complex nature of legged animal movement. Biped locomotion exhibits intricate dynamics and exacting IK foot constraints. As such, the smoothly curving trajectories computed from delegate motion parameters, while suitable for birds, fish, insects, and snakes, are not rich enough to animate the microstructure of bipedal motion. Therefore, several features in Crowd are focused on the special needs of bipeds.

In order to generate the required level of nuance, animated motion clips form the basis for the repertoire of biped movements. In other words, during a Biped/Crowd simulation, the delegates have no effect over the motion of the Biped, they only set goals to be achieved using clips available in the Motion Flow graph. With this approach, known as *motion synthesis*, the animator can precisely control details in the motion either by using hand animation or employing motion capture to produce a set of clips that describe how a member in the crowd behaves.

For example, if you wanted to animate a crowd of marathon runners making their way through the streets of a city, you would need motion clips for various kinds of walking, running, jogging, resting, drinking water, cheering, etc. In effect, each of the motions you might expect to see in a marathon race could be represented as a clip. But motion must be more than a fragmented collection of clips. You must also consider how motions might be sequenced. Which motion transitions are possible from a given motion clip?

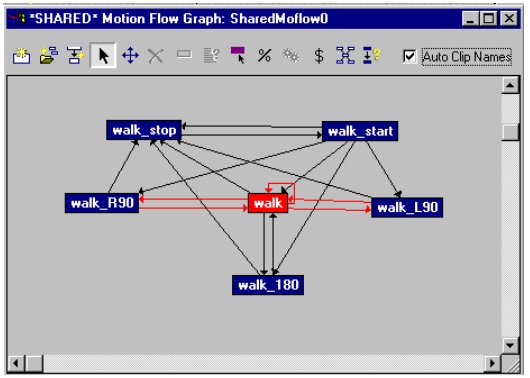
To best understand this process, study this topic and follow the procedure *Using bipeds in a crowd simulation (page 2–1175)*.

Motion Flow Network and Possible Scripts

Biped's *Motion Flow (page 2–1043)* functionality provides the mechanism for defining how separate motions fit together into a fluid animation. In effect, the motion flow network describes which motions can follow from other motions. Once the motion flow network is defined, a broad set of animated actions is possible by following different paths through the network. In Biped, a path through the network is called a motion flow script.

For example, shown below is a motion flow network used in the sample file *walkers.max*. You can find this file in the *cstudio\tutorials\biped_crowds* folder in your 3ds Max path. This folder also contains the BIP files used in the motion flow network. You can access these files only if you have installed tutorials

on your hard disk. For information on installing tutorials, see the *Installation Guide*.



This is a fairly simple network of possible motions, because the characters can only start, stop, turn at 90 degree angles left and right (*walk_L90* and *walk_R90*), and do an about-face (*walk_180*). However, for more natural crowd interaction, it's advisable to expand the motion flow network to include shorter, more finely tuned variations such as turning at 45 degree increments, moving in different directions while facing the same way, loitering motions, and moving at different speeds. The Biped Motion Library has a comprehensive list of clips for you to experiment with.

Tip: You can create motion clips that curve slightly to the left and right by applying Biped's footstep-bending operation to straight-line motion clips. If the clips are motion captured, you should employ footstep extraction during import in preparation for the bending operation. Adding clips that turn slightly will let the biped crowd simulation make minute corrections in heading in order to achieve goal locations more precisely.

Motion flow graphs that work best incorporate fine-tuned transitions. A good way to check your motion-flow transitions is by building test scripts as you build the graph: Add clips to the graph, add the necessary transitions, and optimize the transitions. Optimizing transitions works well as a starting point and, more often than not, produces the smoothest transitions. Next, make a new script that uses your transitions, and use the script to tweak the motion flow until the feet don't slide.

The Shared Motion Flow Networks feature lets many bipeds use a single motion-flow network. Therefore, it's practical to make motion-flow networks large without taxing your computer's memory.

Delegate-Directed Behavioral Goals

You can give a biped a behavioral goal by associating it with a delegate in the Crowd system, and then assigning behaviors to the delegate.

For example, in the *walkers.max* sample, the behavioral goals of each of the biped's delegates are to:

- Move toward the sphere using the Seek behavior.
- Avoid hitting each other using the Avoid behavior.

During a biped crowd simulation, the software attempts to compute the best motion flow script for each biped member of the crowd that satisfies the behavioral goals of its associated delegate. In other words:

- The biped's crowd delegate defines the behavioral goals of the biped.
- The possible motions to reach those goals are defined by the biped's motion flow network.
- The Crowd object's Solve operation computes the script, or path through the motion flow network, that best meets the goals of the biped's delegate.

So in the *walkers.max* sample, the simulation will always choose the best available walking clip in the network that directs the biped's delegate toward the sphere. Each biped's script evolves as the crowd "Solve" computes. This is somewhat like a real-time "game engine" in that the crowd solver's

choice of the next best clip for a given biped is restricted by that biped's currently active clip.

Biped Crowd Avoidance, Priority, and Backtracking

Because bipeds in crowds are always following motion flow scripts, the avoidance behavior for bipeds works differently. Unlike ordinary delegates, biped delegates can move only along motion flow-scripted paths, so if a collision takes place, the software will *backtrack* to the previous clip in the current script and find another path. This may take some time to compute when complex crowd interactions are present since a single backtrack may not be enough. The computation will explore all paths from a given backtracking clip, and if that fails, it will backtrack to the previous clip, and so on, until a solution is found.

In the example, if the current script of a biped is:

```
walk_start
walk
walk_L90
```

and a collision is encountered during the *walk_L90* clip, the biped will backtrack to the end of the *walk* clip and attempt to try a different clip in place of the failed left turn. If that fails, it will try the next best choice, and so on.

Tip: The inclusion of stopping and loitering motions in the motion flow network is sometimes helpful in preventing excessive backtracking since stopping is always an effective way to avoid collisions in a tight situation. In general, the more variation in speed and direction that is possible, the more quickly the backtracking feature will find a solution.

In order to make the backtracking computationally manageable, the biped crowd members are computed one at a time, in order of *priority (page 2–1235)*. Thus, the crowd interaction is accumulated with each successive biped added to the animation. In other words, each waits its turn to compute its complete animation, which entails avoiding the bipeds that have been computed before it. It follows that bipeds with the lowest priorities generally encounter the most collisions, since they must steer around all the bipeds that have higher priorities.

Preparing Characters for a Biped Crowd

To create a biped crowd simulation, you will need several bipeds. Because the crowd simulation factors in each biped's leg length when applying clips to the biped, your simulation will be more accurate if your bipeds are the correct size from the start.

One straightforward workflow would be:

- In a separate file, create or acquire a few character models for your scene.
- Skin the characters with Physique or with another method, such as the 3ds Max Skin modifier.
- Clone the characters and change something about each one, such as the color of clothing or hair, and the character's height. See *Scaling a Character (page 2–1100)*.
- Give each character a unique name that will allow you to identify it in the simulation to some degree. Example: BigMan02, LittleGirl03. See *Naming the Biped (page 2–847)*.
- For each character, disable the MeshSmooth modifier or other modifiers above Physique that add polygons.
- Hide the character meshes and leave the bipeds visible.

This setup makes it easy to merge the characters into the crowd scene. When characters are merged, the meshes themselves will be hidden, which will improve system performance.

Preparing Motions for a Biped Crowd

The clips you use for a biped crowd simulation can be loaded from the CD that comes with **character studio**, imported from motion capture files, or created from scratch.

The integrity of transitions between clips in the motion flow network is very important when creating a biped crowd simulation. When creating transitions for a single biped, it's a simple matter to correct individual transitions after the motion flow script has been created. With a crowd, having to correct transitions for each biped after the simulation is complete would be extremely time-consuming. It is much more efficient to ensure your transitions are correct before starting to solve the simulation.

One way to make this process easier is to set up all motions so each has an ideal transition of exactly the same number of frames, such as 10 or 15. Then you can optimize transitions for this length, and all or most transitions will work flawlessly.

To check transitions, create a simple script for one biped using several transitions in the graph. Check the motion and work with transitions individually until they look right. You might not catch all bad transitions with this method, but you'll catch most of them. A "bad" transition would be one where the feet skip or hop when they should be walking, or any other undesirable change that occurs during a transition. See *Customizing Transitions (page 2–1034)*.

After solving the biped crowd simulation the first few times, check again for bad transitions you might have missed the first time around. After a few simulations, most likely you will have found all or nearly all the bad transitions.

Restrictions

Some features of the crowd system are not designed to work with biped crowds. These behaviors and parameters have no effect on the crowd's movements when bipeds are used with the simulation.

- Vector fields
- The Look Ahead parameter in the Avoid behavior

Procedures: Using Bipeds in a Crowd Simulation

The following sequence of procedures gives the basic steps for creating and saving a motion flow graph, assigning several bipeds to delegates, and then sharing the motion flow graph among the bipeds so that the software automatically creates separate motion flow scripts for each biped, based on the behaviors assigned to its delegate.

Note: These procedures assume you know how to animate bipeds with footsteps and keyframe methods, and save the animations as BIP files. To prepare for using Shared Motion Flow, you should create and save a range of biped motions such as start, walk, turn right, turn left, stop, and wait. See *Loading and Saving BIP Animation (page 2–920)*.

The first step is to create an appropriate motion flow graph and save it to disk.

To create the shared motion flow:

1.  Add a biped, and then open the Motion panel.
2.  Click the Biped rollout > Motion Flow Mode button.
3.  Click the Motion Flow rollout > Show Graph button.
4.  In the Motion Flow Graph toolbar, click the Create Multiple Clips button.

   This displays an Open dialog that lets you select any number of BIP files from the same

directory to add simultaneously to the motion flow graph. Use Click+Shift+click (click, and then Shift+click) to choose several contiguous files, and Ctrl+click to choose non-contiguous files.

5. Choose the BIP files you want the software to use for motion synthesis, and then click the OK button.

   The files are added to the graph as clips. Each clip is automatically named after the file from which it's derived.

   Note: For best results, especially with simulations in which bipeds are to turn at different angles, use as many different turning clips as possible. A minimal setup would include separate left-turn and right-turn walks at angles of 45, 90, 135, and 180 degrees.

   The next step is to add transitions among the clips so **character studio** knows which actions can proceed to and from other actions. You can do this manually for greater control, but for initial testing, you can save time by letting the software add and optimize transitions automatically.

6.  Click the Synthesize Motion Flow Graph button. This uses the first 30 percent and the last 30 percent of each motion to create transitions.

   The graph now shows arrows to and from each clip, as well as from each clip to itself. If you like, delete transitions that obviously don't belong, such as the ones from the stop and start clips to themselves.

    Alternatively, you can use Create Transition to set up a custom graph.

7.  Optimize the transitions. See *Transition Optimization Dialog (page 2–1058)*.

   When you solve the simulation, Crowd automatically generates a motion flow script for the biped, based on this graph. When you have a graph with multiple clips, as in this case, it chooses the starting clip for the script from one or more clips you designate as random start clips.

   If you want to check the integrity of transitions before continuing, do so now. See *Customizing Transitions (page 2–1034)*.

8.  Click the Select Random Start Clips button, and then click a clip.

   This tells the software to start the script with this clip, and uses the default probability of 100 percent that the clip will be chosen.

   If you want the various bipeds to start with different clips, select multiple random start clips by pressing and holding the Ctrl key as you click. The default Random Start Probability setting of 100 for all clips means that the software will choose randomly among them for a starting clip for each biped's script.

   To change the likelihood of starting with specific clips, right-click a clip and modify its Random Start Probability setting. For example, say you want to start each biped's script with any of three clips: clips A, B, and C. You want clip A to be chosen twice as often as clip B or C. In that case, using the Random Start Clips tool, you'd first click clip A, and then Ctrl+click clips B and C. Then you'd right-click each in turn, assigning a Random Start Probability of 60 to clip A, and 30 to both clips B and C.

   Note: The Probability values are arbitrary; what counts is their ratios. For example, values of 80/40/40 or 20/10/10 would work the same.

   Note: You can also set and change random starting clips and start probabilities in the Motion Flow graph after loading the MFE file

into the Shared Motion Flow dialog, described later in this procedure.

Note: If a motion flow script already exists for a biped, for example after you've solved a crowd simulation, the software can use the first clip in the script for subsequent solutions.

9.  In the Motion Flow rollout, click the Save File button, and save the graph in the MFE format.

   The next step in using bipeds in a crowd simulation is to create the simulation.

To set up the crowd simulation:

1. Reset 3ds Max.
2. Set up a crowd simulation with any number of delegates, using behaviors appropriate to the crowd scene you want to create. Do not use bipeds at first; work with delegates only to rough out the motion of the crowd members.

   Tip: When first starting out with motion synthesis, use smaller crowds of eight or so delegates.
3. Solve the simulation and adjust the settings as necessary to obtain the desired motion.
4. Merge bipeds and characters into the scene, and associate each delegate with a different biped using the delegate's *Motion Parameters rollout (page 2–1183)* > Biped group settings.

   You must select the biped's center of mass (COM) object (typically named Bip0#), as indicated by the mouse cursor turning into a crosshairs icon when it's over the COM in the active viewport.
5. Turn on Use Biped for each delegate.

   Tip:  To link up any number of delegates with bipeds, all at the same time, select the Crowd helper object and click the Setup rollout > *Biped/Delegate Associations button (page 2–1199)*. Use the dialog to connect the pairs and turn on Use Biped for each delegate. See *Linking Objects to Delegates (page 2–1169)*.
6. Set the biped/delegates to use a random start clip as the first clip. You can set this simultaneously for multiple delegates with the *Edit Multiple Delegates dialog (page 2–1197)*.

To apply the shared motion flow:

Next, you use the Shared Motion Flow function to apply the saved motion flow graph to the bipeds.

1. Select any biped and open the Motion panel.
2.  Click the Biped rollout > Motion Flow Mode button.
3.  Click the Motion Flow rollout > Shared Motion Flow button.
4. In the Shared Motion Flow dialog, click the New button.

   This creates a new shared motion flow and assigns it a default name. You can change the name if you like.

   Next, load a motion flow file.
5. Click the Parameters group > Load .mfe button, and use the Open dialog to load a motion flow file. Typically, this would be the one you saved earlier in the procedure.

   Next, specify the bipeds that will share this motion flow.
6. In the Parameters group, click the Add button, and use the Select dialog to specify the bipeds that will share the motion flow.

   For your convenience, the Select dialog shows only center of mass objects for the bipeds in the scene.

   After you click the Select button, the bipeds appear in the dialog, in the list under "Bipeds Sharing this Motion Flow."

7. To correctly share a motion flow, bipeds' legs must be scaled the same. If any of the bipeds are scaled differently than the one you started with, an alert appears, and then, when you click the OK button in the alert box, the wrong-scale bipeds are noted as such in the list. At this point, you can select one of the bipeds in the list, and then click the Set Shared Motion Flow Scale button to match the others' scale to that biped. Or you can click one of the Reset Wrong Scales buttons to rescale the wrong-scale bipeds or just their legs. Be sure to take one of these measures before proceeding.

   One more step in the Shared Motion Flow dialog is necessary: You must activate Motion Flow mode for all the bipeds sharing the motion flow. A special button in the dialog lets you perform that action in one step.

8.  Click the Put Multiple Bipeds in Motion Flow button.

   This activates Motion Flow mode for all the bipeds sharing the motion flow.

9. Click the OK button to exit the dialog.

To offset the delegates and test the simulation:

Delegate-controlled bipeds can begin their animation with their motion flow scripts' first clip, if it exists, or with a random motion clip. But when you load a motion flow file into the Shared Motion Flow dialog, any scripts in the file are ignored. Thus, delegate-associated bipeds using motion flow in an unsolved crowd simulation have no existing scripts, and you must specify that they use the random start clip that you set in the motion flow. You do this via the delegates.

1. Use the Edit Multiple Delegates dialog > Biped group to specify Random Start Clip for all the delegates.

   Solve the simulation.

2. Select the Crowd object, and go to the Modify panel.

3. In the Solve rollout, set the desired End Solve frame.

4. Click Solve to run the simulation.

   Crowd solves the simulation.

5. Check the solved simulation by dragging the frame slider and/or playing back the animation.

To fine-tune the simulation:

Chances are good that adjustments will be required. You can resolve problems in a number of different ways: Change the behavior and/or motion flow setups, change delegate parameters, and so on.

If you find that bipeds are colliding and interpenetrating, you can take advantage of Crowd's special Priorities and Backtracking features. In fact, it is strongly recommended that you use both options for most Crowd/Biped simulations.

1. Use the *Priority rollout (page 2–1235)* controls for assigning different priorities to your delegates. Typically, delegates at the head of the crowd should have the highest priorities (that is, the lowest Priority settings).

2. In the Solve rollout > Bipeds group, turn on Biped/Delegates Only, then turn on Use Priorities, and then turn on Backtracking. If you've solved previously, it's a good idea to turn on Delete Keys Before Solve as well.

   Now, when it solves the simulation, Crowd solves for one biped/delegate at a time, starting with the lowest-priority biped/delegate. As it solves for each subsequent biped/delegate, it looks for collisions, and when the occur, it backs up the solution to the end of the previous clip, and if necessary, previous clips, and then tries different paths through the motion flow graph. This method can take longer, which is

why Backtracking is off by default, but it's often the best way to resolve problems with colliding bipeds.

Non-Biped Crowds

Motion synthesis in **character studio** lets the software derive character motion from a combination of crowd behaviors and either motion flow networks, when animating bipeds (see *Biped Crowds (page 2–1172)*), or clip controllers, when working with non-bipedal creatures. In the latter case, using the Global Motion Clip and Master Motion Clip controllers, you can animate groups of creatures such as birds, butterflies, schools of fish, and insects. You can create clip controllers either as block controllers in Track View, or, more directly, with the Crowd helper controls on the Global Clip Controllers rollout.

Two Approaches to Animation

You can animate your creature either in place with looping animation but no transformational motion (such as a bird flapping its wings), or you can incorporate transformational motion into the animation as well (the bird moves upward while flapping its wings). In-place cyclic motion lends itself to flying or swimming motions like birds and fish, while adding lateral motion lends itself to crawling type animation where feet should be planted on the ground and not sliding. Depending on which you use, you toggle options on the Motion Clips tab of the Synthesis dialog. In both cases, you use crowd delegates driven by behaviors to motivate the creatures, which are linked to those delegates.

Note: To animate a model for motion synthesis, apply modifiers to the model and animate their parameters. Modifiers such as Bend, Taper, Wave, and Xform produce animation you can use with motion synthesis. Do not use sub-object animation, such as animation of vertices on an Editable Mesh object.

Cyclic In-Place Animation

First you create a creature with a few short loop cycles, like the beating of wings, gliding, turning left and turning right. This creature is assigned as the *Global Object* or the master object from which the motion clips will be derived. Then clones of the original creature are created. The clones are positioned and linked to delegates. States are created to select which clips will play based on a state.

For example, if a bird (delegate) is pitching up or accelerating, the fast-beating clip is used; if the bird (delegate) slows to a stop, the wings-at-rest clip is used, and so on. During synthesis, the software determines which state should be active depending on the speed and direction of the delegates. An active state determines which clip should be applied to the clones of the original object. Clips are blended together to create the animation. Available states are speed, acceleration, pitch, pitch velocity, heading velocity, or script (MAXScript).

Animation with Lateral Motion

For multi-legged creatures that walk, you can animate lateral motion as well as the cyclic motion of the legs moving. This is done to ensure that the creatures' feet do not slide as they move. The software then uses the lateral motion information to create a state that perfectly matches the actual motion. The software then strips the actual motion out. When a delegate approaches the speed and heading recorded in that state, the appropriate motion clip is triggered. This technique minimizes sliding feet.

Use the **character studio** crowd tools to create the initial motion for the delegates. Use a seek or avoid behavior to steer birds, for example. Your

object with the loop animation is then copied and the copies are linked to the delegates to create the complete animation. The delegate handles the path and the clip controllers handle the looped animation.

You can create Master Motion Clip and Global Motion Clip in Track View by assigning a controller to the available controller under Block Control. It is, however, simpler to use the Crowd helper controls on the Global Clip Controller rollout to apply and use the clip controllers.

See also

Synthesis Dialog (page 2–1246)

ClipState Dialog (page 2–1253)

Global Motion Clip

Global Motion Clips store the clips to be shared among multiple Master Motion Clips, which are assigned to the cloned objects during synthesis. Global Motion Clips also contain the logic for performing motion synthesis on a collection of objects with trajectories and states associated with clips. Controls for motion synthesis are found in the Synthesis dialog.

The way the motion clip keys are scaled and ordered depends upon user-defined states. Each state contains one or more motion clips that will be played when the state is active.

Master Motion Clip

Master Motion Clips are controllers that contain *motion clips*, or individual clips of animation. These motion clips are sequenced to create animation, and overlap slightly with automatic blending to smooth the transitions between clips.

Procedure

To use Motion Synthesis with non-bipedal creatures:

All of the work involved in copying and synthesizing clips takes place using controls in the *Synthesis dialog (page 2–1246)*. This dialog has three tabs: Motion Clips, State, and Synthesis.

This procedure assumes that the global object is static and has animation that loops. For creatures with many legs, you can animate lateral motion on the global object and then strip out the lateral motion with the Synthesis dialog. This latter approach serves to minimize foot sliding in a multi-legged creature.

1. Animate an object.

   Create animation in one position, like a bird's beating wings. Create a variety of animation like a gliding motion (wings still), wings beating slowly, and so on. To animate the object, apply modifiers and animate their parameters.

   This will be the Global object, from which animation clips will be derived.

2. Using the Create panel > Helpers > Object Type rollout, add a Crowd object and a Delegate object.

   Create the objects in the Top or Perspective viewport.

3. Use *Scatter Objects (page 2–1189)* in Crowd function to clone the delegate and optionally distribute the clones (you can also distribute them manually).

   Make sure you end up with an equal number of delegates and animated object clones.

   Next, associate and link the objects to the delegates.

4.  On the Setup rollout, click the Object/Delegate Associations button.

   The *Object/Delegate Associations dialog (page 2–1196)* appears.

5. Add the objects and delegates into their respective columns.
6. Click Align Objects With Delegates, and then click Link Objects To Delegates, and then click OK to exit the dialog.

   The objects align themselves with the delegates and are linked to the delegates.

   Next, animate the delegates with behaviors. See *Assigning Behaviors (page 2–1159)* for information on using behaviors.

   When you solve the simulation, the cloned objects follow the delegates, which are guided by behaviors. You then generate motion synthesis based on the delegate movement.
7. Select the Crowd object. On the Modify panel > Global Clip Controllers rollout, click New, and use the Select dialog to select the Global object from step 1.

   The object appears in the Global Clip Controllers rollout list.
8. In the list, click the object, and then click the Edit button.

   The Synthesis dialog displays. Its name is that of the object.
9. On the Motion Clips panel, turn off all check boxes in the From Global Object and Remove Local groups.

   Use these options only if your original object has lateral motion to coordinate with footsteps.
10. On the Motion Clips panel, click New.

    This opens the MotionClip Parameters dialog, which lets you set the name, color, and frame range for a motion clip.
11. Choose a descriptive name and a frame range for the motion clip.

    For example, frames 0 through 10 might be your glide animation.
12. Continue to define clips using different frame ranges from the Global object's animation. Give the clips descriptive names.

    Next, you define states whose parameters determine when motion synthesis is to use each clip.
13. Click the State tab, and then click New State.

    A new state is added to the drop-down list at the top of the Synthesis States group. Give the state a descriptive name. In many cases, the state can use the same name as the motion clip that's to be associated with it.

    Next, you specify the state's parameters; that is, how the delegate should be moving when the associated object is to use its motion clip. For example, your glide motion should be active only if acceleration is less than 0.
14. Click Edit Properties and define how the software should activate the clip, based on any combination of speed, acceleration, and so on. Click each appropriate tab, turn on its Use ... check box, and set parameters. Click Exit to exit the dialog.

    When using a range, make sure the Min setting is lower than the Max setting. For example, when using a negative range such as -180 to -10, enter the number with the higher absolute value (-180) as the Min setting.

    You should already have several motion clips. Now you need to associate a clip with this state. For example, if you've defined a state whose acceleration is less than 0, you might associate the Glide clip with it.
15. In the MotionClips group, click Add Clip, and in the Select MotionClip dialog, highlight a clip and click OK.
16. Repeat steps 14 and 15 for each state to be used in motion synthesis.
17. Go to the Synthesis panel, click New Master Motion Clip, and add all of the cloned objects.

18. Click Auto Blend All and then click Synthesize All.

    The synthesis occurs as a progress bar displays. When the synthesis is completed, the Synthesis dialog reappears. You can now view the ClipState parameters' ranges and average values by clicking State panel > Edit properties. This is useful in fine-tuning state properties.

    Click OK to exit the Synthesis dialog.

19. Click Play. It's not necessary to re-solve the simulation.

    The objects follow the delegates and are animated using clips that are activated according to delegate motion and the states you created.

Crowd Animation User Interface



Crowd animation lets you simulate the behavior of crowds of people, animals, or other beings parametrically, using several different types of objects. The topics that follow describe the user interface for setting up a crowd simulation.

- The *Crowd helper object (page 2–1187)* includes facilities for replicating and grouping objects and *assigning behaviors (page 2–1200)* to objects and groups. The Crowd object works directly on *delegate helper objects*, and indirectly on bipeds and other objects via delegates.
- *Cognitive controllers (page 2–1206)* can change delegate behavior based on circumstances in the simulation.
- *Motion synthesis (page 2–1241)* uses animation from the linked object based on delegate orientation, speed and other factors.

Crowd Shortcuts

 To enable the **character studio** keyboard shortcuts, turn on the *Keyboard Shortcut Override toggle (page 3–872)*.

All **character studio** keyboard shortcuts activate when the Motion panel is active and the Keyboard Shortcut Override button is active.

See Also

Keyboard Shortcuts (page 3–871)

Keyboard Panel (page 3–793)

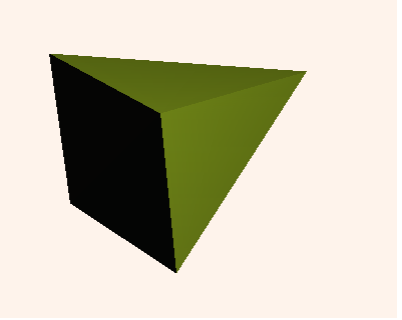
Customize User Interface Dialog (page 3–792)

Action	Shortcut	Description
Solve	S	Solves crowd simulation.
Step Solve		Step-solves crowd simulation.

Delegate Helper Object

Create panel > Helpers > Object Type rollout > Delegate

The Delegate is a special helper object used in crowd animation. It serves as an agent for motion created by a *Crowd object (page 2–1187)* and its behaviors. The Crowd object controls a delegate or delegates, whose motion can then be imparted to a biped or other object. Delegates cannot be rendered.



The delegate object is shaped like a pyramid. By default, the point of the pyramid indicates the forward direction, the delegate's local Y axis.

The delegate object uses the following rollouts:

Geometry Parameters Rollout (page 2–1183)

Motion Parameters Rollout (page 2–1183)

Also, you can set parameters for multiple delegates simultaneously, optionally with random variation, using the Crowd object's *Edit Multiple Delegates dialog (page 2–1197)*.

See also

Creating Crowd Helpers (page 2–1157)

Crowd Helper Object (page 2–1187)

Geometry Parameters Rollout

Create panel > Helpers > Object Type rollout > Delegate > Geometry Parameters rollout

Select a Delegate object. > Modify panel > Geometry Parameters rollout

Use these parameters to modify the delegate object's size.

Interface

Width, Depth, Height—Sets the width, depth, and height of the Delegate object. These fields also act as readouts when you create the delegate.

Note: The Delegate is a helper object, and thus cannot be rendered. Thus the size of the Delegate object is primarily for use in scene setup, and for determining bounding box extents.

Motion Parameters Rollout

Create panel > Helpers > Object Type rollout > Delegate > Motion Parameters rollout

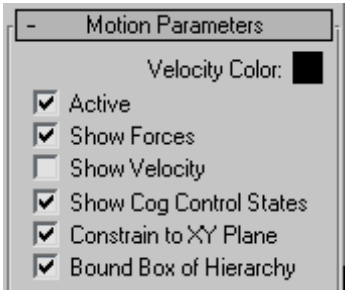
Select a Delegate object. > Modify panel > Motion Parameters rollout

The Motion Parameters rollout lets you specify a Delegate object's characteristics, including speed, acceleration, and other factors. It also lets you associate the delegate with a biped.

Important: When using delegates with bipeds, only the settings in the Biped group have any effect in the simulation, because Biped gets all its speed and turning information from motion flow clips and behaviors.

Note: You can set any or all motion parameters simultaneously for a number of delegates with the *Edit Multiple Delegates dialog (page 2–1197)*.

Interface



Velocity Color—When Show Velocity is on, uses the specified color to draw a vector in the delegate's center during the simulation solution. The vector length indicates the delegate's relative speed. Default=black.

Active—The delegate object is subject to control by a Crowd object. Default=on.

Show Forces—The forces being applied to a delegate by any applicable behaviors are drawn as vectors whose length indicates the extent of the forces and whose orientation shows the direction in which the behavior is influencing the delegate to move. Default=on.

For example, if the delegate is affected by a *Space Warp behavior (page 2–1221)* and a *Wander behavior (page 2–1231)*, the vectors (using default colors) are yellow and blue-green, respectively. These vectors are visible only during solution of the crowd simulation.

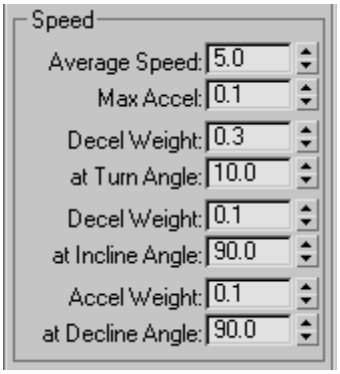
Show Velocity—Uses the Velocity Color (see above) to draw a vector whose length depicts the delegate's relative speed. This vector is visible only during the solution of the crowd simulation. Default=off.

Show Cog Control States—During a solution, a text label appears next to the delegate showing the name of the *cognitive controller (page 2–1170)* state or transition that currently directs its behavior, if any. Default=on.

Constrain to XY Plane—The delegate remains at its initial height (position on the world Z axis) throughout the simulation. When off, the delegate's height can change during the simulation, for example when seeking an object at a different height. Default=on.

Bound Box of Hierarchy—When on, the *Avoid behavior (page 2–1211)* uses the bounding box of the delegate and all of its children to perform its behavior. Default=on.

Speed group



Average Speed—Specifies the delegate's baseline velocity in 3ds Max units (or the current unit type) per frame. Default=5.0.

The delegate's speed can be modified during the simulation by a variety of factors, such as a linked biped's built-in speed and Deviation settings in a behavior.

Max Accel(eration)—Multiplied times Average Speed to determine the maximum acceleration. Default=0.1.

For example, given the defaults of 5.0 for Average Speed and 0.1 for Max Accel, the acceleration or deceleration at any moment can be no greater than 0.5 units/frame/frame.

Decel(eration) Weight—Specifies how much a delegate should slow down when turning.

The higher this setting, the more the delegate slows down when it reaches the turn angle (see following

parameter). A value of 0 specifies no slowdown; a value of 1 tells the delegate to stop. Default=0.3.

The algorithm computes a value, d, which goes linearly from 0 to (1 - Decel Weight) as the turn angle of the delegate goes from 0 to the Turn Angle specified by the user. The speed of the delegate is then multiplied by d. For example, when the delegate turns at the Turn Angle or greater, its speed will be multiplied by (1 - Decel Weight), slowing it down as much as possible based on this parameter. When the delegate is not turning at all, its speed is not affected by the Decel Weight. When the delegate is turning at half the specified Turn Angle, d = Decel Weight / 2, so its speed will be multiplied by (1 - Decel Weight / 2).

As a practical example, take a delegate traveling at 10 units/frame. Decel Weight is set to 0.4, and At Turn Angle is set to 30. When the delegate has turned 15 degrees (half the At Turn Angle), the effective deceleration weight is 0.2. Subtract that quantity from 1 to get 0.8, and then multiply that times the delegate's speed to get 8 units per second halfway into the turn. At the full turn (30 degrees), the delegate travels at 6 units per frame.

At Turn Angle—Specifies the turn angle at which Decel Weight's full slowdown effect is applied. Default=10.0.

If the current turn angle is less than At Turn Angle, the algorithm divides the latter by the former, and then divides the Decel Weight setting by the result to derive the effective deceleration weight.

Decel(eration) Weight—Specifies how much the delegate should slow down when moving at an upward slant. Default=0.1.

See *Decel Weight* for a full explanation.

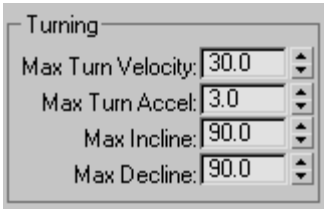
At Incline Angle—Specifies the upward slant angle at which Decel Weight's full slowdown effect is applied. Default=90.0.

Accel(eration) Weight—Specifies how much the delegate should speed up when moving at a downward slant. Default=0.1.

See *Decel Weight* for a full explanation, taking into account that Accel Weight produces a speedup effect rather than a slowdown. Thus, the effective acceleration weight is added to 1, not subtracted from it.

At Decline Angle—Specifies the downward slant angle at which Accel Weight's full speedup effect is applied. Default=90.0.

Turning group



These parameters affect the delegate's turning behavior; that is, how it changes direction in response to forces applied by crowd behaviors.

You can use the *Orientation behavior (page 2–1214)* to affect how a delegate turns and banks, independently of the actual path taken.

Turning applies to objects moving both on the ground and in the air.

Max Turn Velocity—Specifies the maximum number of degrees a delegate can turn per frame. Applies both to heading and pitch. Default=30.0.

Max Turn Accel—Specifies how much the delegate's heading or pitch angle can change per frame. This controls angular acceleration and deceleration. For smooth turns, keep it relatively low. Default=3.0.

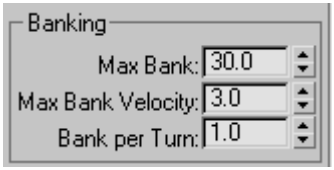
Tip: If a delegate exhibits sluggish turning behavior during a simulation, try increasing Max Turn Velocity, Max Turn Accel, or both.

Max Incline—Specifies the maximum number of degrees a delegate can turn upward at any given frame. For example, most birds can't fly straight up, so you might set this to 45 for a bird. Default=90.0.

Max Decline—Specifies the maximum number of degrees a delegate can turn downward at any given frame. Default=90.0.

For example, for a bird that can't fly straight up but can fly straight down, you might set to Max Incline to **45** and Max Decline to **90**.

Banking group



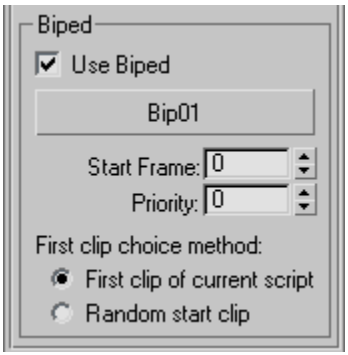
These parameters affect the delegate's banking behavior; that is, how it tilts around its front-back axis as it changes direction while moving. Banking typically applies to objects moving in the air or on water, but can also apply to ground-based objects such as one- or two-wheeled vehicles.

Max Bank—Specifies the maximum number of degrees a delegate can bank. Default=30.0.

Max Bank Velocity—Specifies the maximum number of degrees a delegate can bank per frame. Default=3.0.

Bank per Turn—The number of degrees the delegate will bank as a function of the turn angle at the current frame. For example, if Bank Per Turn=1, the delegate will bank one degree for every degree it is turning at a given frame. Default=1.0.

Biped group



These parameters relate to the use of bipeds associated with delegates. In order to have a biped exhibit character animation as it follows the delegate's course, you must use motion flow methods. For a procedure, see *To use bipeds in a crowd simulation (page 2–1175)*.

Use Biped—Associates the delegate with a biped (specified with the None button), and causes the delegate's speed to be determined by that of the biped's existing motion. Its behavior (for example, seeking another object) remains defined by a crowd object.

Note: This button is not available until a biped is designated via the None button.

None (label)—Click this button and then select a biped to be associated with the delegate's motion.

You can select a biped by clicking it in a viewport, by clicking its center of mass object (e.g., bip01), or with the Select Object dialog, opened by clicking the toolbar button Select By Name (you cannot use the Selection Floater). Thereafter, the name of the biped object appears on the button.

Tip: When using the viewport, the mouse cursor changes to a crosshairs when positioned over a biped's center of mass.

Start Frame—Specifies the frame at which the biped's first clip will begin to play.

If, when several bipeds share the same starting clip, you vary this setting per biped, they won't walk in lockstep formation. This is most useful when you take advantage of the ability of the *Edit Multiple Delegates dialog (page 2–1197)* to randomize the start frame for each delegate.

Priority—Sets the delegate priority, which determines the order of solution in biped/delegate simulations.

For details, see *Priority Rollout (page 2–1235)*.

First clip choice method—Determines which motion clip in the shared motion flow graph Crowd initially uses to animate the biped linked with the delegate.

- **First clip of current script—**Uses the first clip in the biped's motion flow script, if a script exists. If this option is chosen, but there is no script, an error message is generated.
- **Random start clip—**Uses the random start clip or clips specified in the shared motion flow graph, if random start clips have been designated. If this option is chosen, but no random start clips have been designated, an error message is generated.

Crowd Helper Object

Create panel > Helpers > Object Type rollout > Crowd

The Crowd helper object acts as the command center for controlling crowd simulations in **character studio**. In most cases, you won't need more than one Crowd object per scene.

The Crowd object provides the following rollouts:

Setup Rollout (page 2–1188)

Solve Rollout (page 2–1232)

Priority Rollout (page 2–1235)

Smoothing Rollout (page 2–1238)

Collisions Rollout (page 2–1240)

Geometry Rollout (page 2–1240)

Global Clip Controllers Rollout (page 2–1241)

Crowd Behaviors

The Crowd object also lets you add behaviors to the scene, choose the current behavior from a list, and provides a rollout for modifying that behavior. Behaviors provided with **character studio** include:

Avoid Behavior (page 2–1211)

Orientation Behavior (page 2–1214)

Path Follow Behavior (page 2–1216)

Repel Behavior (page 2–1218)

Scripted Behavior (page 2–1220)

Seek Behavior (page 2–1220)

Space Warp Behavior (page 2–1221)

Speed Vary Behavior (page 2–1222)

Surface Arrive Behavior (page 2–1223)

Surface Follow Behavior (page 2–1226)

Wall Repel Behavior (page 2–1227)

Wall Seek Behavior (page 2–1229)

Wander Behavior (page 2–1231)

See also

Creating Crowd Helpers (page 2–1157)

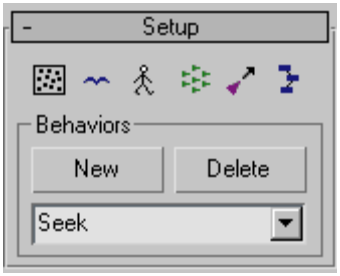
Setup Rollout

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout

Select a Crowd object. > Modify panel > Setup rollout

The Setup rollout of the Crowd helper object contains controls for setting up crowd functions.

Interface



A number of functions are available from a row of buttons at the top of the Setup rollout:

Scatter—Opens the *Scatter Objects Dialog (page 2–1189)*.

Objects/Delegate Associations—Opens the *Object/Delegate Associations dialog (page 2–1196)*.

Biped/Delegate Associations—Opens the *Associate Bipeds With Delegates dialog (page 2–1199)*.

Multiple Delegate Editing—Opens the *Edit Multiple Delegates dialog (page 2–1197)*.

Behavior Assignments—Displays the *Behavior Assignments and Teams dialog (page 2–1200)*.

Cognitive Controllers—Displays the *Cognitive Controller editor (page 2–1206)*.

Behaviors group

Use these controls for adding, removing, and renaming behaviors.

New—Launches the *Select Behavior Type dialog (page 2–1205)*. Choose a behavior and then click OK to add the behavior to the scene. Then use the *Behavior Assignments and Teams dialog (page 2–1200)* to assign the behavior to a delegate or delegates in the scene.

The first time you add a behavior to the scene using this command panel, a new rollout appears for this behavior below the Setup rollout. This rollout lets you change settings for the behavior. To display the rollout for a different behavior in the scene, choose it from the drop-down list in the Behaviors group.

Following is a list of available behaviors:

Avoid Behavior (page 2–1211)

Orientation Behavior (page 2–1214)

Path Follow Behavior (page 2–1216)

Repel Behavior (page 2–1218)

Scripted Behavior (page 2–1220)

Seek Behavior (page 2–1220)

Space Warp Behavior (page 2–1221)

Speed Vary Behavior (page 2–1222)

Surface Arrive Behavior (page 2–1223)

Surface Follow Behavior (page 2–1226)

Wall Repel Behavior (page 2–1227)

Wall Seek Behavior (page 2–1229)

Wander Behavior (page 2–1231)

Delete—Deletes the current behavior.

If the behavior is currently in use, that is, it's assigned to a delegate or team, either directly in the *Behavior Assignments and Teams dialog (page 2–1200)* or indirectly through a cognitive

controller, a small dialog appears asking you to confirm the deletion. If you delete a directly assigned behavior, its assignment is removed from the scene. If you delete a behavior used in a cognitive controller, it is removed from the state to which it was assigned.

Behaviors List—Lists all behaviors in the current scene (added with New). Select a behavior from the list to have its rollout appear below the Setup rollout.

Note that behaviors that appear in this list aren't necessarily assigned to any delegates active in the crowd simulation. Likewise, a behavior whose rollout appears below the Setup rollout isn't necessarily active or assigned. To assign delegates and/or activate behaviors, use the *Behavior Assignments and Teams dialog (page 2–1200)*.

You can rename a behavior by first selecting it from the list, and then clicking its name and entering a new one from the keyboard. It's a good idea to give descriptive names to behaviors; for example, *Avoid Red Team*.

Note: If you add the first behavior in the scene from the *Behavior Assignments and Teams dialog (page 2–1200)*, the text box remains empty and no rollout for the behavior appears. To edit the behavior, choose it from the list.

Scatter Objects Dialog

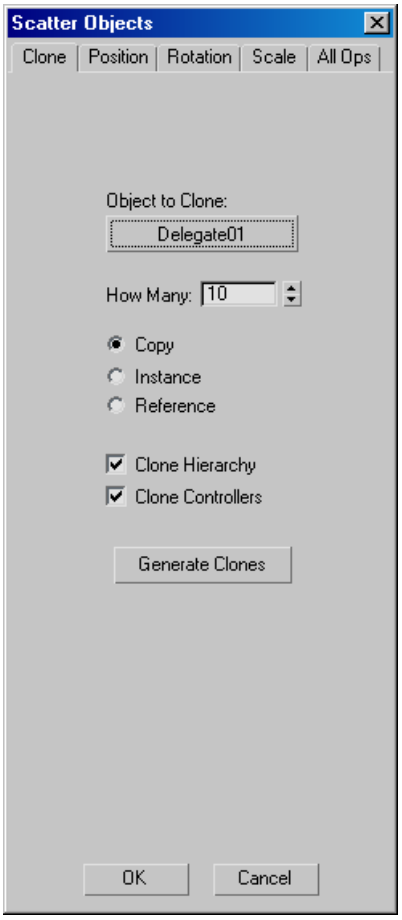
Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > Scatter Objects

Select a Crowd object. > Modify panel > Setup rollout > Scatter Objects

The Scatter Objects dialog of the Crowd helper object includes facilities for creating crowds by cloning objects, such as delegates. It also lets you distribute the clones and other objects within a radial area, along a shape, across a grid object or surface, or within a box or sphere. You can also specify various orientation and scaling options for scatter objects.

Interface

Clone panel



Contains the basic options for cloning an object.

Object to Clone (None)—Click this button, and then select an object in the scene to be cloned.

How Many—Specifies the number of clones to be generated.

Copy/Instance/Reference—Lets you specify how the object is cloned. It can be cloned as a copy, an *instance (page 3–958)*, or a *reference (page 3–1002)*.

Clone Hierarchy—When on, all objects linked to the selected object are cloned as well, with the hierarchical structure retained intact for each clone.

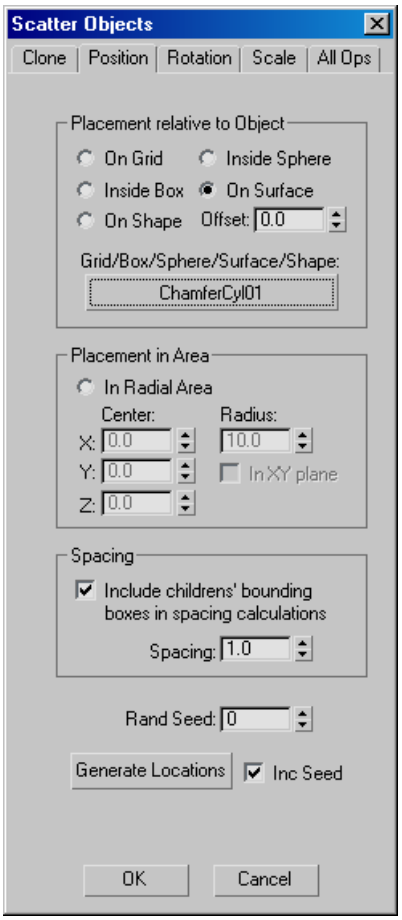
Clone Controllers—When on, any controllers (that is, animation) associated with the selected object are cloned as well.

Generate Clones—Click this button to create the specified number of clones of the object whose name appears on the Object To Clone button.

Cloned objects are all created in the same location; to distribute them, you must set distribution options on the Position panel, and then click the Generate Locations button.

After you generate clones, the original object and its newly created clones are scatter objects; that is, they're selected in the list in the All Ops panel (Select Objects To Transform), and are thus subject to subsequent operations on the Position, Rotation, and Scale panels. You can change this selection via the All Ops panel > *Select Objects to Transform* function or by generating clones again.

Position panel



Contains options for positioning objects using a reference (distribution) object. You can distribute objects randomly over the surface of a grid or other object, along a shape, or within the volume of a box or a sphere. You can choose only one option from the first two group boxes.

Placement Relative to Object group

On Grid/Inside Sphere/Inside Box/On Surface/On Shape—Choose the appropriate item before selecting the reference object.

- On Grid distributes the scatter objects over the surface of a grid object.

- Inside Box and Inside Sphere distribute the scatter objects within the volume of a primitive box or sphere object, respectively.
- On Surface distributes the scatter objects over the surface of any renderable object. For example, you can create a landscape object for use as a distribution surface by applying a Noise modifier to a patch grid.
- On Shape distributes the scatter objects along a shape object: a spline or NURBS curve. If the shape consists of more than one curve, Scatter uses the lowest-numbered curve (typically the first one added).

Offset—When using On Surface, specifies a consistent distance above the surface (using surface normals) for distribution. Available only when On Surface is chosen.

Grid/Box/Sphere/Surface/Shape (None)—Click this button, and then select an object in the scene to be used as a reference object. Listed objects are limited to the chosen category.

Note: With the Grid, Box, and Sphere choices, you can use only a grid helper object, a primitive sphere, or a primitive box as a reference object. A primitive sphere or box that has been converted to a editable mesh object can't be used as a reference object. Also, the appropriate radio button (see the first item in this group) must be chosen before you select a reference object.

Placement in Area group

Contains options for positioning scatter objects in a radial area, without using a reference object.

In Radial Area—Distributes the scatter objects randomly in a spherical or circular arrangement, using the remaining controls in this group box.

Center—Specifies the center of the distribution in world coordinates.

Radius—Specifies the maximum distance from the center within which objects are to be positioned.

In XY Plane—Specifies that objects are to be distributed on the world XY plane only, resulting in a disc-like array.

Spacing group

Include children's bounding boxes in spacing calculations—When on, all of a hierarchical scatter object's sub-objects are considered when determining spacing. When off, only the selected object is considered.

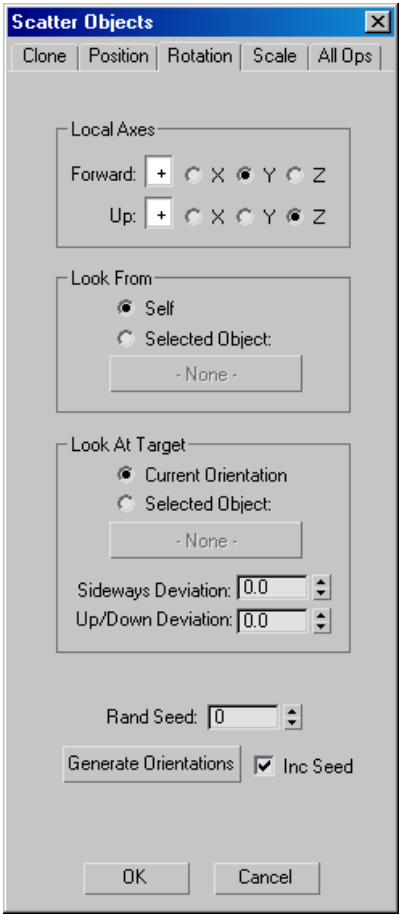
Spacing—Specifies the minimum distance between scatter objects. The Spacing setting is multiplied by the size of the object's bounding sphere to determine how close objects can get. If Spacing is left at 1.0, the default, objects normally cannot be positioned within each others' bounding spheres. If Spacing is set to 2.0, objects are separated by a distance equal to or greater than the size of the bounding sphere.

Rand Seed—Specifies a seed value for randomizing scatter objects' locations. If, for example, you use the same value for clones of a delegate and then a master object, each pair ends up in the same place. If a scene has more than one crowd, each should use a different seed to avoid having identical configurations.

Generate Locations—Click this button to produce a set of locations for all scatter objects; that is, cloned objects or objects selected with the *Select Objects to Transform* button.

Inc(rement) Seed—When on, and you click the Generate Locations button, Scatter adds 1 to the Rand Seed value, and redistributes the objects using the new random seed. Default=on.

Rotation panel



Contains options for orienting scatter objects. You can specify alternative forward and up axes, plus a target object toward which the objects will point. In addition, you can specify a source object; when using both source and target objects, the objects are rotated so they're parallel to the line between the two.

Local Axes group

Use these settings to designate alternative forward and up axes. The default settings match the delegate axes.

Forward—+/-/X/Y/Z: Specifies which axis of the objects points forward, for use with the Look At Target option. When the + button is active, the default condition, the positive chosen axis is used. Click the + button to use the negative axis.

Up—+/-/X/Y/Z: Specifies which axis of the objects points upward; this axis is aligned with the world Z-axis. When the + button is active, the default condition, the positive chosen axis is used. Click the + button to use the negative axis.

Note: You cannot specify the same axis as Local Forward and Local Up simultaneously. If you choose an axis for one that's already chosen for the other, the software switches the other to a different axis.

Look From group

Self/Selected Object—Determines the direction from which the objects look. By default, each object looks from its own position (Self), so that when several objects are looking at a single target, each is oriented differently. To orient each object so that it's parallel to an imaginary line between two objects (the "from" object and the "to" object), choose Selected Object and specify the object with the (None) button.

None (label)—When choosing Selected Object as the Look From object, use this button to specify the "from" direction. Click the button, and then select an object from which the objects are to look.

Look At Target group

Current Orientation/Selected Object—Determines the direction toward which the scatter objects look. By default, each object retains its current orientation. To orient each scatter object so that it's parallel to an imaginary line between two objects (the "look from" object and the "look at target" object), choose Selected Object and specify the object with the (None) button.

None (label)—Use this button to specify the "to" direction. Click the button, and then select an object toward which the scatter objects are to look.

Sideways Deviation—Sets a maximum deviation angle in degrees for the objects' sideways orientation. If the scatter objects should look in an object's general direction but may look at a spot to either side of the target, use Sideways Deviation to set the maximum amount by which they can deviate from the calculated angle. The actual deviation amount for each object is calculated at random, based on the Deviation settings and the Rand Seed setting. Range=0.0 to 180.0.

Up/Down Deviation—Sets a maximum deviation angle in degrees for the objects' up/down orientation.

If scatter objects should look in an object's general direction but may look at a spot above or below the target, use Up/Down Deviation to set the maximum amount by which they can deviate from the calculated angle. The actual deviation amount for each scatter object is calculated at random, based on the Deviation settings and the Rand Seed setting. Range=0.0 to 180.0.

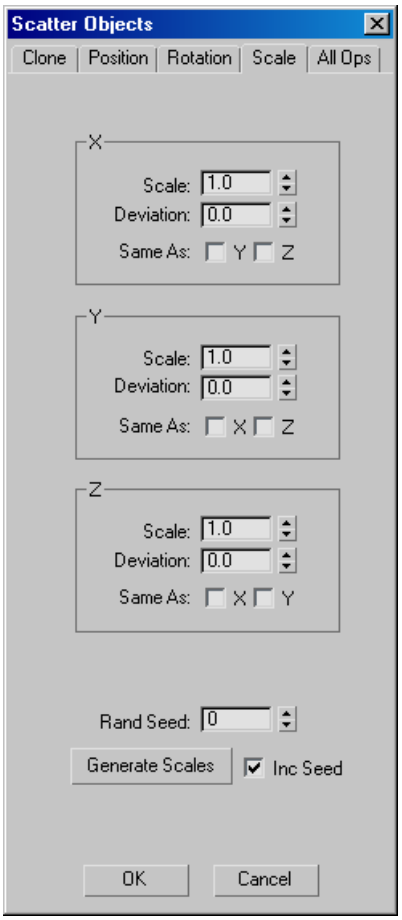
Rand Seed—Specifies a seed value for randomizing the scatter objects' orientations, based on the Deviation settings.

If you use the same value for clones of a delegate and then a master object, each pair ends up with the same orientation. If a scene has more than one crowd, each should use a different seed to avoid having identical configurations.

Generate Orientations—Click this button to produce a set of orientations for all scatter objects; that is, cloned objects or objects selected with the *Select Objects to Transform* button.

Inc(rement) Seed—When on, and you click the Generate Orientations button, Scatter adds 1 to the Rand Seed value, and reorients the scatter objects using the new random seed. Default=on.

Scale panel



Contains options for scaling scatter objects. You can apply uniform or non-uniform scaling, with optional per-axis deviation for scaling variation.

Each axis group has a "Same As" option that lets you scale that axis by the same amount as another. To prevent non-uniform scaling, set two axes to be the same as the third. For example, set scaling in the X group, and then in the Y and Z groups, turn on Same as X.

Warning: These controls can apply non-uniform scaling to objects, which may produce unexpected results when performing other operations within 3ds Max.

X group

Scale—Sets scaling on the X-axis as a multiplier. Default=1.0.

Deviation—Sets the maximum factor for randomization of scaling. For each scatter object, Deviation is multiplied by a random number between 0.0 and 1.0, and then added to the Scale multiplier.

Same As Y/Z—Lets you use the same scaling as on the Y- or Z-axis, whether explicit or randomized. When you specify an axis, the parameters group for that axis becomes unavailable.

Y group

Scale—Sets scaling on the Y-axis as a multiplier. Default=1.0.

Deviation—Sets the maximum factor for randomization of scaling. For each scatter object, Deviation is multiplied by a random number between 0.0 and 1.0, and then added to the Scale multiplier.

Same As X/Z—Lets you use the same scaling as on the X- or Z-axis, whether explicit or randomized. When you specify an axis, the parameters group for that axis becomes unavailable.

Z group

Scale—Sets scaling on the Z-axis as a multiplier. Default=1.0.

Deviation—Sets the maximum factor for randomization of scaling. For each scatter object, Deviation is multiplied by a random number between 0.0 and 1.0, and then added to the Scale multiplier.

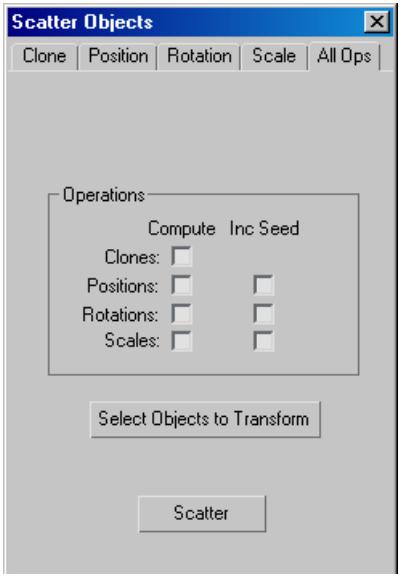
Same As X/Y—Lets you use the same scaling as on the X- or Y-axis, whether explicit or randomized. When you specify an axis, the parameters group for that axis becomes unavailable.

Rand Seed—Specifies a seed value for randomizing the clones' scales, based on the Deviation settings. If you use the same value for clones of a delegate and then a master object, each pair ends up with the same scaling factor.

Generate Scales—Click this button to scale all scatter objects; that is, cloned objects or objects selected with the *Select Objects to Transform* button.

Inc(rement) Seed—When on, and you click the Generate Scales button, Scatter adds 1 to the Rand Seed value, and re-scales the scatter objects using the new random seed. Default=on.

All Ops panel



This panel lets you perform various permutations of cloning and transform operations in a single step, with or without successive randomization.

Operations group

Compute Clones—Turn on to clone the object chosen with the *Object to Clone button*. When you click the Scatter button, the object is cloned,

and then any specified transforms are applied to the clones.

Turning on Clones makes the Select Objects To Transform button unavailable. The object to clone and cloning parameters must be specified on the *Clone panel (page 2–1189)*.

Compute Positions/Rotations/Scales—Any options in this column that are turned on when you click the Scatter button cause the respective transforms to be applied to the current selection (see Select Objects To Transform, below) according to the settings in the *Position panel (page 2–1190)*, *Rotation panel (page 2–1192)*, and *Scale panel (page 2–1193)*.

Inc(rement) Seed Positions/Rotations/Scales—Any options in this column that are turned on cause the respective Rand Seed settings to be incremented by 1 each time you click the Scatter button.

Use this option to experiment with various randomized transform combinations for your clones.

Select Objects to Transform—Lets you designate objects to be affected by clicking the Scatter button.

Clicking this button opens a version of the Select dialog that's unique to Scatter Objects functionality. If you've performed one or more cloning operations during the current session, the results of the most recent cloning are selected by default, including the original cloned object. For example, if you created 10 clones, 11 objects are selected. You can use the Select dialog to alter or replace this selection.

Note: The results of the most recent cloning operation remain selected even if you close and later reopen the Scatter Objects dialog.

Scatter—Performs any cloning and/or transforms that are turned on.

OK—Retains all changes and closes the dialog.

Cancel—Forgets any changes and closes the dialog.

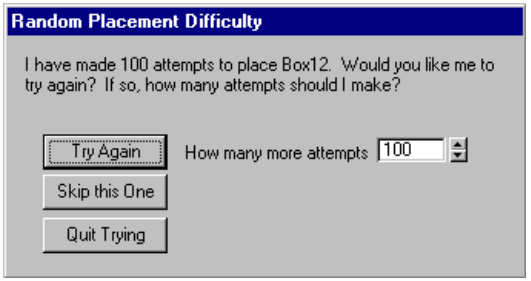
See also

Creating Crowd Helpers (page 2–1157)

Random Placement Difficulty Dialog

This dialog appears when the software encounters difficulty placing cloned objects without overlapping using the *Scatter Objects dialog (page 2–1189)*. The dialog text tells you how many attempts the software has made to place a specific object, and asks you if you want to try again.

Interface



Try Again—Click this button to force the software to make N more attempts, where N is set in the How Many More Attempts field.

Skip This One—Instructs the software to stop trying to place the current object and proceed to the next.

Quit Trying—Aborts the Generate Locations process; no more objects will be placed.

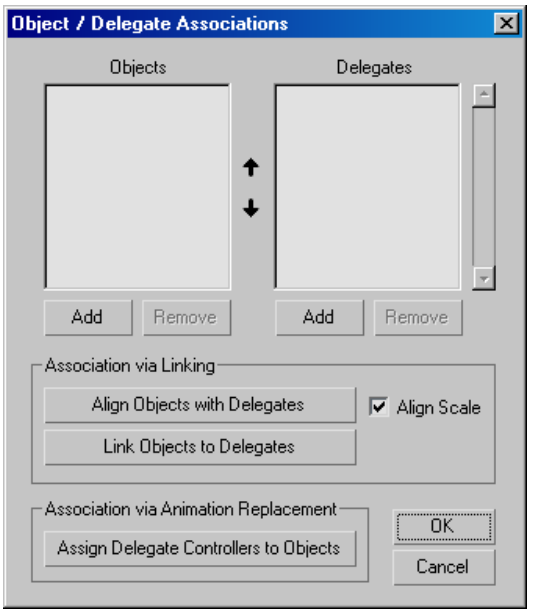
Object/Delegate Associations Dialog

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > Associate Objects with Delegates

Select a Crowd object. > Modify panel > Setup rollout > Associate Objects with Delegates

You can use this dialog to link any number of delegate-object pairs. You can also use this dialog to align objects with delegates, optionally matching scaling factors as well. Lastly, you can specify that the objects use the delegates' controllers.

Interface



Objects—Lists objects available for linking, specified using the Add function (see following item). You can select any number of objects from this list for subsequent removal.

Add—Click this to open the standard 3ds Max Select dialog, which lists all objects in the scene, including delegates. Make your selection, and then click the Select button to add the objects to the Objects list.

Remove—Deletes the highlighted object or objects from the list.

Shift Up/Shift Down—Use the arrow buttons between the two lists to move highlighted items higher or lower in the list.

When you click Align Objects With Delegates, Link Objects To Delegates, or Assign Delegate Controllers To Objects, the software creates associations between pairs of items at matching positions in the lists.

Delegates—Lists delegates available for linking, specified using the Add function (see following item). You can select any number of delegates from this list for subsequent removal.

Add—Click this to open the standard 3ds Max Select dialog, which lists all delegates in the scene. Make your selection, and then click the Select button to add the delegates to the Delegates list.

Remove—Deletes the highlighted delegate or delegates from the list.

Association via Linking group

Align Objects with Delegates—Aligns each object with its corresponding delegate by moving and rotating the object on its local axes. Each delegate's local Y axis points in the forward direction, so aligned objects will be rotated so their Y axes point in the forward direction as well. This option works much like the 3ds Max Align feature.

Align Scale—When on, clicking Align Objects with Delegates sets each object's absolute scaling factor to that of its corresponding delegates. This is useful if, for example, you've randomized delegates' sizes with the Scatter Objects *Scale panel (page 2–1193)*, and want the associated objects to match.

Link Objects to Delegates—Creates a hierarchy for each object-delegate pair, with the delegate as parent.

Association via Animation Replacement group

Assign Delegate Controllers to Objects—Copies each delegate's controllers to the paired objects as an instance.

This is the same as using Track View > Copy Controller from the delegate, and then pasting the controller as an instance to the object. Does not link objects hierarchically with delegates.

Once you've set up the delegate animation the way you want it, if you then want to apply the animation to an object or objects *en masse*, use this function. You can then delete the delegates if you like.

OK—Implements any changes and closes the dialog.

Cancel—Eliminates any changes and closes the dialog.

Note: Clicking OK has no intrinsic effect; in order to implement any of the dialog functions, you must first click at least one of the Align/Link/Assign buttons.

Edit Multiple Delegates Dialog

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > Multiple Delegate Editing

Select a Crowd object. > Modify panel > Setup rollout > Multiple Delegate Editing

The Edit Multiple Delegates dialog lets you define groups of delegates and set parameters for them. You can create and store up to 10 different configurations or settings combinations; each consists of one or more delegates and settings for the delegates. The parameters are mostly the same as those found in the delegate object's *Motion Parameters rollout (page 2–1183)*, with the following exceptions and additions.

- Each setting has an associated SET check box, which lets you determine whether the setting has any effect. When off (the default), the setting has no effect. When on, the setting affects the specified delegate or delegates when you click the Apply Edit button.
- Each numeric parameter has two Value settings and an associated Random check box, which lets you specify a random value within a specified range for each member of the group. By default, Random is off, and the Value 1 setting is applied for all parameters with SET turned on. If you turn on Random for a parameter, its Value 2 setting becomes available. If you then specify a different setting for Value 2, the software calculates a different random number between Value 1 and Value 2 for each delegate in the group.

Tip: To reset a parameter to a specific value for all delegates in a group after it's been set to a random value within a range, turn off Random and turn on SET for the parameter, then set Value1 to the desired value, and then click OK.

Procedure

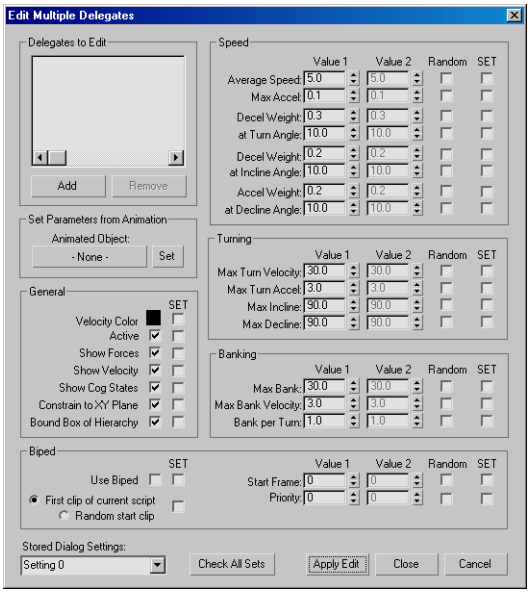
To edit multiple delegates:

1. Select the Crowd object and open the Modify panel.
2.  On the Setup rollout, click the Multiple Delegate Editing button.

   This opens the Edit Multiple Delegates dialog.
3. Optionally, in the dialog (bottom-left corner), choose a stored dialog setting to use from the 10 available settings.

4. If necessary, use the Delegates To Edit group box controls to add delegates to or remove them from the current setting.
5. Modify the remaining parameters as necessary. Be sure to turn on the SET check box for any parameters that are to change.
6. Click the Apply Edit button to make the changes and exit the dialog.

Interface



Delegates to Edit group

This group shows delegates belonging to the current settings combination in a list box and lets you add and delete members.

Add—Click this button, and then choose delegates to add from the Select dialog.

Remove—To remove delegates from the list, first choose the names of those to delete in the list box (drag to choose two or more contiguous names, or use Ctrl-click to choose non-contiguous names), and then click Remove.

Set Parameters from Animation group

Use this function to obtain motion parameters from an animated object and apply them to all specified delegates. It affects only Average Speed, Max Accel, and the Turning parameters.

Because this one animation will set most of the parameters of the delegate, it should be representative of a whole range of motion of the delegate. For example, the object should turn and accelerate. The animation should be somewhat lengthy so that averages are calculated correctly.

Animated Object—Specifies an animated object. Click this button, and then choose the object from the list in the Select dialog.

Set—After specifying the animated object, click this button to apply its parameters to the delegate settings. Also turns on the SET check box for any affected parameters.

General group

Rather than numeric values, the settings in this group are on-off switches, except for the first, Velocity Color. To change Velocity Color, click the color swatch, use the Color Selector dialog to pick a new color, and then turn on the Velocity Color SET check box. To change any other setting in the General group, click the check box to the right of the setting, and then turn on the setting's SET check box.

These settings are the same as those found in the delegate object's *Motion Parameters rollout (page 2–1183)*.

Speed group

These parameters are the same as those found in the delegate object's *Motion Parameters rollout (page 2–1183)*. For an explanation of the Random and SET check boxes, see the introduction to this topic.

Turning group

These parameters are the same as those found in the delegate object's *Motion Parameters rollout (page 2–1183)*. For an explanation of the Random and SET check boxes, see the introduction to this topic.

Banking group

These parameters are the same as those found in the delegate object's *Motion Parameters rollout (page 2–1183)*. For an explanation of the Random and SET check boxes, see the introduction to this topic.

Biped group

These parameters are the same as those found in the delegate object's *Motion Parameters rollout (page 2–1183)*. For an explanation of the Random and SET check boxes, see the introduction to this topic.

Stored Dialog Settings—Use this list to specify up to 10 different combinations of delegates and settings. To store a combination, choose a name from the list, and then specify the delegates and settings. To recall a combination, choose its name from the list. To rename a combination, choose its name from the list, and then highlight the name and edit it using the keyboard.

Check All Sets—Click this button to turn on all SET check boxes. This ensures that any changes you make in the dialog take effect when you click the Apply Edit button.

Apply Edit—Click Apply Edit to implement all changed settings and exit the dialog.

Close—Click Close to remember, but not implement, all changed settings and exit the dialog.

Cancel—Click Cancel to forget all changed settings and exit the dialog.

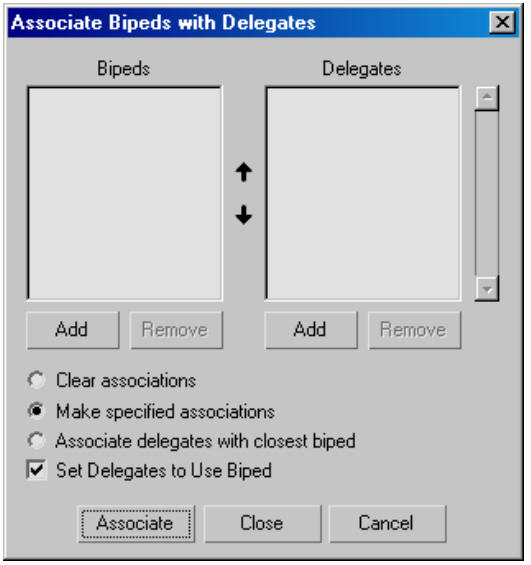
Associate Bipeds With Delegates Dialog

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > Biped/Delegate Associations button

Select a Crowd object. > Modify panel > Setup rollout > Biped/Delegate Associations button

Use this dialog to associate any number of delegates with an equal number of bipeds. Add delegates and bipeds to the two lists, and order them so the desired pairs are across from each other. Then choose Make Specified Associations and click the Associate button. Alternatively, you can remove existing delegate-biped associations, or simply associate each delegate with the biped nearest it in the scene.

Interface



Bipeds—Lists bipeds available for linking, specified using the Add function (see following item). You can select any number of objects from this list for shifting up or down, or deleting.

To clear highlighted items, control-click them.

Add—Click this to open the standard 3ds Max Select dialog, which lists all bipeds in the scene that are not currently listed in the Associate Bipeds with Delegates dialog. Make your selection, and then click the Select button to add the delegates to the Objects list.

Remove—Removes the highlighted biped or bipeds from the list.

Shift Up/Shift Down—Use the arrows between the two lists to move highlighted items higher or lower in the lists.

When Make Specified Associations is chosen and you click the Associate button, the software creates associations between pairs of items at matching positions in the lists.

Delegates—Lists delegates available for linking, specified using the Add function (see following item).

You can highlight any number of objects from this list for shifting up or down, or deleting.

To clear highlighted items, control-click them.

Add—Click this to open the standard 3ds Max Select dialog, which lists all delegates in the scene that are not currently listed in the Associate Bipeds With Delegates dialog. Make your selection, and then click the Select button to add the delegates to the Delegates list.

Remove—Removes the selected delegate or delegates from the list.

Clear Associations—When this is on, and you click the Disassociate button, **character studio** eliminates any delegate-biped associations. Makes the Bipeds list and buttons unavailable.

Make Specified Associations—When chosen, and you click the Associate button, **character studio** associates each parallel delegate-biped pair in the two lists. That is, the first delegate is associated with the first biped, the second delegate with the second biped, and so on.

Associate Delegates With Closest Biped—When chosen, and you click the Associate button, **character studio** calculates the biped nearest each delegate in the scene and links the two together.

Set Delegates to Use Biped—When on, and you click the Associate button, the software turns on the *Use Biped* option for all delegates listed in the dialog.

Associate/Disassociate—Implements specified changes, calculates any random values, and closes the dialog. Button text changes to "Disassociate" when the Clear Associations option is chosen.

Close—Remembers any changed settings and closes the dialog. No new delegate settings are calculated or applied.

Cancel—Ignores any changes and closes the dialog.

Behavior Assignments and Teams Dialog

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > Behavior Assignments

Select a Crowd object. > Modify panel > Setup rollout > Behavior Assignments

The Behavior Assignments and Teams dialog lets you group *delegates* into *teams (page 2–1204)*, and assign *behaviors* and *cognitive controllers (page 2–1206)* to individual delegates and teams. It also lets you modify existing assignments.

The dialog is modeless; while it's open, you can use the Modify panel to adjust behaviors and set up new behaviors, as well as animate assignments' Weight settings.

Note: Crowd doesn't let you use multiple cognitive controllers with a delegate. You can assign them, but when you solve, the software notifies you that it will use only the first assigned cognitive controller.

Note: The Enable Flashing option helps you see which delegates are affected by different actions in this dialog. When it's turned on, and you perform any of the following, the relevant delegates "flash" (highlight briefly) in the viewports:

- Click a delegate or team in the Assignment Design group
- Click an assignment in the Behavior Assignments group
- Add members to a team in the Teams group
- Remove members from a team in the Teams group

Procedures

To group delegates into a team:

1. In the Teams group, click the New Team button.
2. Use the Select Delegates dialog to designate the delegates in the team, and then click the OK button.
3. To change a team name, choose it from the drop-down list at the top of the Teams group, click on its name, and then use the keyboard to edit the text.
4. To remove team members, choose the team from the drop-down list at the top of the Teams group, select the members to remove from the lower list, and then click the Remove Members button.
5. To add team members, choose the team from the drop-down list at the top of the Teams group, click the Add Members button, and then use the Select Delegates dialog to designate the delegates to add.

To create a new behavior assignment:

This procedure gives the basic method for assigning a behavior or cognitive controller to a delegate or team.

1. If you want to assign the same behavior to more than one delegate, use the *Teams group (page 2–1204)* to collect delegates into teams.

   Note: You can still assign behaviors to an individual delegate, even if it belongs to one or more teams.

2. Make sure no existing assignments in the Behavior Assignments group are highlighted. If any are, Ctrl+click them to clear the selections.

   If assignments are highlighted, the software assumes you want to modify the existing assignments

3. In the Assignment Design group, select one delegate or team, and one or more behaviors or one cognitive controller.

   You can select only one item from either side of this group, with the exception of behaviors. If you choose multiple behaviors, the software creates a separate assignment for each.

   Note: When you select a delegate or team, it briefly highlights in the viewports to indicate the affected delegates.

4. Click the New Assignment button. This is the vertical button to the right of the assignment Design group, with five right-pointing arrows.

   This adds the assignments to the list in the Behavior Assignments group.

5. At this point, you can highlight an assignment, and then change its Weight setting, its Active status, delete it, or change the assignee and/or behavior/cognitive controller.
6. Click OK to accept the changes and close the dialog.

To modify an existing behavior assignment or assignments:

1. In the Behavior Assignments group, select the assignment or assignments to change.

   You can select multiple assignments by Ctrl+clicking for non-contiguous items or Shift+clicking for contiguous items, and then change the assignees or behaviors for all of them at once.

2. To change assignees, in the Assignment Design group, select a delegate or team.
3. To change the assigned behaviors, in the Assignment Design group, select a behavior or cognitive controller.
4. Click the Reset Assignment button. This is the vertical button to the right of the assignment Design group, with five right-pointing arrows.
5. Change the Weight setting and Active status as necessary.
6. Click OK to accept the changes and close the dialog.

Interface

Assignment Design group



Lets you set up assignments by choosing a behavior or cognitive controller and a delegate or team to assign it to. Choose one item from the left side (upper or lower list), and one item from right side (upper or lower list). Then click the New/Reset Assignment button immediately to the right of the Assignment Design group (vertical button with five right-pointing arrows).

Tip: With the exception of Behaviors, you can choose only one item from either side of this group. To assign the same behavior to more than one delegate, the most efficient method is to use the *Teams group (page 2–1204)* to gather delegates into teams.

You can select multiple behaviors for a new assignment to a delegate or team. When you click New Assignment, the software creates a separate assignment for each highlighted behavior. For changing assignments, you're still restricted to choosing one behavior at a time. If you choose an existing assignment and multiple behaviors, the Reset Assignment button becomes unavailable.

Delegates—Lists delegates in the scene.

Behaviors—Lists existing behaviors.

To use a behavior that hasn't been added to the scene yet, click the New Behavior button at the bottom of this group.

Teams—Lists teams in the scene.

To create a new team, use the controls in Teams group box.

Cognitive Controllers—Lists existing cognitive controllers.

To create a new controller, click the Setup rollout > *Cognitive Controllers (page 2–1206)* button. You needn't first exit the Behavior Assignments and Teams dialog; when you close the editor, the new controllers are added to the Cognitive Controllers list.

Clear Selections—Deselects all highlighted items in the Assignment Design and Behavior Assignments groups.

Use this before modifying an assignment, to avoid possible confusion.

New Behavior—Opens the *Select Behavior Type dialog (page 2–1205)*, which lets you add a behavior to the scene for use in an assignment.

To modify a new behavior, use the facilities available in the Crowd object's rollouts.

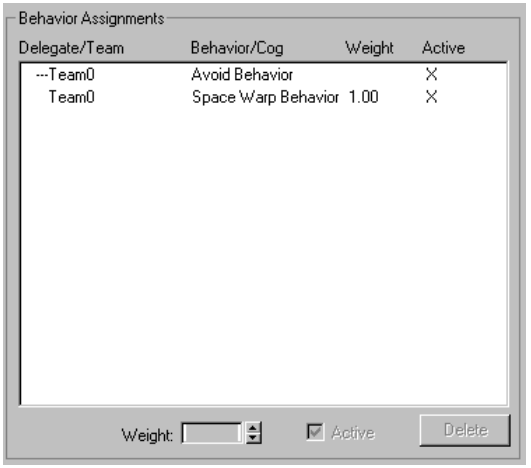
Note: If you add the first behavior in the scene from this dialog, the text box in the Crowd object > Setup rollout remains empty, and no rollout for the behavior appears. To edit the behavior, choose it from the list.

New Assignment/Reset Assignment—Click to assign a behavior or behaviors or a cognitive controller to a delegate or team.



This vertical button with five right-pointing arrows on it is situated between the Assignment Design and Behavior Assignments group. It's available only when two items in the Assignment Design group are highlighted (exception: Multiple behaviors can be highlighted). If no item in the Behavior Assignments group is highlighted, clicking the button creates a new assignment and adds it to the assignments list. If one or more items in the Behavior Assignments group are highlighted, clicking the button sets the highlighted assignments to use the highlighted delegate/team and behavior/cognitive controller combination.

Behavior Assignments group



Lets you create and modify behavior assignments.

List box—Displays all current behavior assignments, including team or delegate name, assigned behavior or cognitive controller, weight setting, and active status.

Items are sorted in alphabetical first by Delegate/Team name, and then by Behavior/Cog name. A dashed line appears before a list entry if it's the first item for that delegate or team.

To modify or delete an assignment, choose it from the list, whereupon the software highlights the assigned components in the Assignment Design group. Make the changes using the remaining controls in this dialog.

Note: You can select multiple assignments from the list by Ctrl+clicking for non-contiguous items or Shift+clicking for contiguous items. To clear an item, Ctrl+click it.

Weight—The relative effect of the assigned behavior or cognitive controller. The higher an assignment's Weight setting is than others', the greater relative effect it will have. This setting is animatable. Default=1.0.

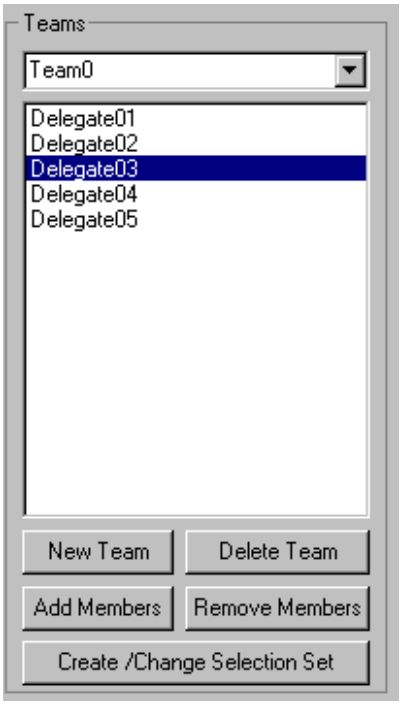
In most cases, you should keep Weight within a range of 0.0 to 1.0. Higher settings are available but shouldn't be used unless absolutely necessary.

Note: The Weight setting is not relevant to the *Avoid (page 2–1211)*, *Orientation (page 2–1214)*, or the *Surface Follow (page 2–1226)* behavior, and is thus unavailable for assignments using any of those three.

Active—When on, the assignment is currently in effect. When off, the assignment has no effect. This check box is animatable. Default=on.

Delete—Deletes the highlighted behavior assignment.

Teams group



Lets you define, modify, and delete teams of delegates.

Note: You can toggle the display of this group box with the No Teams/Teams button below the Behavior Assignments group.

Drop-down list—Displays the name of the current team.

To view a different team, choose it from the list. To change a team name, click in the box and then use the keyboard to edit the text.

List box—Displays delegates in the current team.

New Team—Adds a team to the list, and opens the *Select Delegates dialog (page 2–1205)* to let you specify new team members.

The default team name is "Team," followed by a number, starting with "0" and counting up.

Delete Team—Deletes the current team.

Team members are not deleted from the scene.

Add Members—Lets you add members to the current team. Use the *Select Delegates dialog (page 2–1205)* to specify new team members.

Remove Members—Removes selected members from the team.

Removed members are not deleted from the scene.

Create/Change Selection Set—Adds the current team to the list of selection sets, accessible from the Named Selection Sets list on the Main toolbar.

If the current team already is a selection set, and you subsequently changed the team's makeup, click this button to update the members in the Named Selection Sets list.

Enable Flashing—When on, and you click a list item in the dialog or create/modify a team, the relevant objects highlight briefly in the viewports to indicate which are affected. Default=on.

See the *introductory note* for details.

OK—Click this button to accept all changes and close the dialog.

No Teams/Teams—Toggles display of the Teams group box.

By default, the Teams group is displayed; click the No Teams button to turn it off. When it's off, click the Teams button to turn it on. The state of the toggle persists only during the current session.

Select Behavior Type Dialog

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > Behavior Assignments > New Behavior

Select a Crowd object. > Modify panel > Setup rollout > Behavior Assignments > New Behavior

Use this dialog to select the type of behavior to be added to a Crowd object. The choices are:

Avoid Behavior (page 2–1211)

Orientation Behavior (page 2–1214)

Path Follow Behavior (page 2–1216)

Repel Behavior (page 2–1218)

Scripted Behavior (page 2–1220)

Seek Behavior (page 2–1220)

Space Warp Behavior (page 2–1221)

Speed Vary Behavior (page 2–1222)

Surface Arrive Behavior (page 2–1223)

Surface Follow Behavior (page 2–1226)

Wall Repel Behavior (page 2–1227)

Wall Seek Behavior (page 2–1229)

Wander Behavior (page 2–1231)

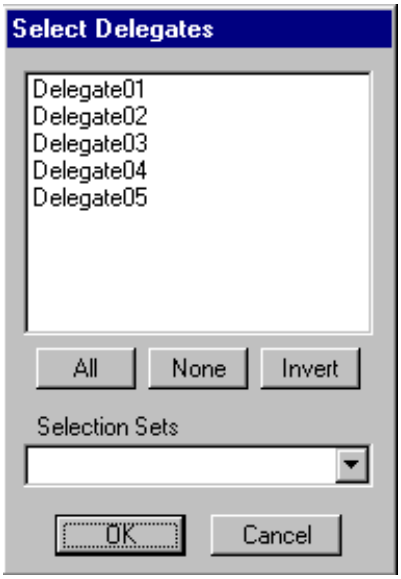
Select Delegates Dialog

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > Behavior Assignments > Choose or add a team. > Add Members

Select a Crowd object. > Modify panel > Setup rollout > Behavior Assignments > Choose or add a team. > Add Members

The Select Delegates dialog lets you designate delegates to be assigned to teams using the *Behavior Assignments and Teams dialog (page 2–1200)* for assigning crowd behaviors.

Interface



List box—Lists all delegates in the scene.

All—Selects all delegates in the list.

None—Deselects all delegates.

Invert—Inverts the current selection.

Selection Sets—Choose an item from this to select all members of the selection set. You must first have created a selection set from one or more delegates.

OK—Closes the dialog and implements changes.

Cancel—Closes the dialog and ignores changes.

Cognitive Controllers

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > Cognitive Controllers

Select a Crowd object. > Modify panel > Setup rollout > Cognitive Controllers

The Cognitive Controller editor lets you combine behaviors into states. More importantly, it lets you sequence different *behaviors (page 2–1159)* and behavior combinations using state diagrams, where conditionals written in MAXScript impose changes in behavior. For example, you can specify that a character or object is to wander aimlessly until it comes within a certain distance of another object, whereupon it heads straight for that object. Or you can specify that one character is to avoid another only when the second character is avoiding the first.

Note: The MAXScript conditionals used in the cognitive controller typically consist of a single line of code. You can load and save them separately as *.ms* files, but they are also stored within the *.max* scene file in which they reside.

The editor interface consists of an icon-based toolbar above a window that contains the state diagram. When you first open the editor, no state diagrams exist. Begin by clicking the New button to create a new state diagram.

Tip: If you find yourself consistently assigning two or more behaviors to delegates or teams, you can save time by combining the behaviors into a single-state cognitive controller, or "behavior module," and assigning that instead. The only disadvantage is that you can't animate the weights of behaviors used in the cognitive controller, but you can work around that by using transitions.

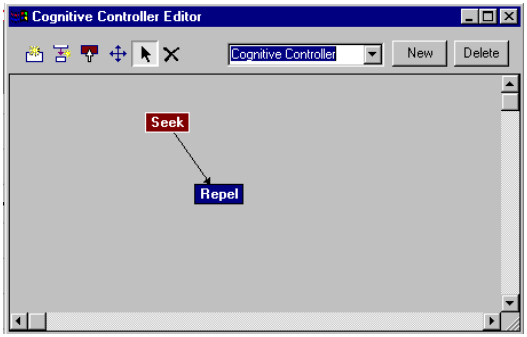
Note: Crowd doesn't let you use multiple cognitive controllers with a delegate. You can assign them, but when you solve, the software notifies you that it will use only the first assigned cognitive controller.

See also

State Dialog (page 2–1207)

State Transition Dialog (page 2–1208)

Interface





Create State—Lets you create new states in the diagram. Click this button, and then click in the state diagram area to add states. A state appears as a rectangular box containing the name of the state.

The first state you add is, by default, the first state in the controller that executes when the simulation is run. This is indicated by its red color; states you add subsequently are colored blue. To set a different state to execute first, use the Set Start State function.

You specify a state's name and behaviors by editing the state. To edit a state, right-click it. This opens the *State dialog (page 2–1207)*.



Create Transition—Lets you link states with transitions. Click this button, and then drag between two states to create the transition, starting

with the earlier state. The transition appears as a black arrow pointing from the first state to the second. Alternatively, if you click a state with the Create Transition tool active, you create a transition that loops back to the state itself.

Right-click a transition to specify its characteristics and conditions by editing the transition. This opens the *State Transition dialog (page 2–1208)*.


Set Start State—Normally the state that executes first in a cognitive controller is the one that was added first. Use this tool to choose a different state to execute first. The start state is red; the rest are blue.

Typically you would use this when you have a circular sequence of states, and you want to change which state executes initially.


Move State—Lets you move states around in the window by dragging them.


Select State/Transition—Lets you select states and transitions for subsequent deletion. Selected states have white outlines, and selected transition lines are white.

You can select multiple states by dragging a box around them. You can select multiple states and transitions by holding the Ctrl key as you click.


Delete State/Transition—Lets you delete one or more states or transitions. First select any combination of states and transitions to delete, and then click this button.

(Name)—Shows the name of the current state diagram. To display and/or edit another, choose it from the list.

To change a state diagram's name, click the name in the box and use the keyboard to edit the text.

New—Adds a new cognitive controller. By default, cognitive controllers are named "Cognitive Controller" followed by a number, but you can change this to anything you like.

Delete—Deletes the current cognitive controller. This is an undoable operation.

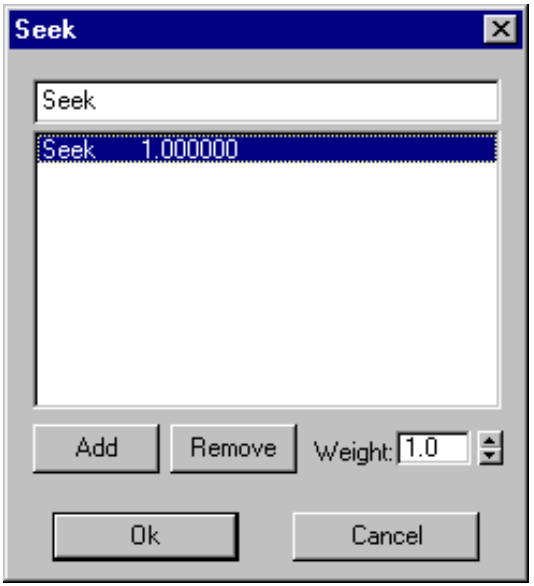
State Dialog

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > Cognitive Controllers > right-click a state icon.

Select a Crowd object. > Modify panel > Setup rollout > Cognitive Controllers > right-click a state icon.

The State dialog lets you add behaviors to, and remove them from, a cognitive controller state in a crowd simulation, and specify weights for each behavior.

Interface



State name—Displays the name of the state.

To change the name, click this text and edit from the keyboard.

List—Displays the names of all behaviors associated with the state.

To remove a behavior or change its weight, click the behavior's name, and then make the appropriate changes using the controls below the list.

Note: The order of behaviors in this list is immaterial; all behaviors execute simultaneously.

Add—Opens the Select Behavior dialog, which displays the names of all *behaviors (page 2–1211)* in the current Crowd object that are not associated with the current state. Select a behavior (use Shift and/or Ctrl to choose multiple behaviors), and then click OK.

Remove—Eliminates the highlighted behavior from the state.

Weight—Specifies the selected behavior's relative influence in the state. Default=1.0.

The higher the weight in relation to other behaviors' weights, the more evident the results of the behavior in the state. In most cases, you should keep Weight within a range of 0.0 to 1.0. Higher settings are available but shouldn't be used unless absolutely necessary.

OK—Closes the dialog and implements changes.

Cancel—Closes the dialog and ignores changes.

State Transition Dialog

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > Cognitive Controllers > Right-click a transition line.

Select a Crowd object. > Modify panel > Setup rollout > Cognitive Controllers > Right-click a transition line.

These settings control how the software effects a transition from one state to another when using a *cognitive controller (page 2–1206)*. For more detailed information, see *To set up and use a cognitive controller (page 2–1170)*.

The Transition Script

The most important element of the transition is the MAXScript conditional script. This is a script associated with the controller that is executed once per frame, and can test any aspect or aspects of the scene and cause a transition or not, depending on whether the result of the test is successful (true, or 1) or unsuccessful (false, or 0).

The software executes scripts once per frame per assigned delegate, so objects and effects can be animated and still let delegates react with accuracy.

All scripts used in transitions use the following structure:

```
fn [FunctionName] del t = (
[MAXScript code]
if [MAXScript conditional]
then 1
else 0 )
```

The opening section contains "fn" (function) followed by the function name, which also must appear in the State Transition dialog, and then the input parameters "del t", and lastly "= (". Following this there can be any MAXScript code, or none.

The closing section contains a necessary MAXScript conditional, and then "then 1 else 0". This means: If the result of the conditional is true, then return 1 (that is, the transition is to take place), or if the result of the conditional is false, then return 0 (that is, the transition is not to take place). You could reverse the order of the numbers 1 and 0 ("then 0 else 1") so that the conditional being true would cause no transition to take place, and vice-versa. Lastly, the function must end with a close parenthesis: ")".

Following are some examples of scripts that can be used in cognitive controllers, along with brief

explanations. These are presented for you to modify and use in your own scenes.

Note: See this topic in the online *User Reference* for sample code to test an object position, test the distance between two objects, and test a modifier parameter.

Testing a Particle System Parameter

This sample script tests the number of particles emitted by particle system Spray01, and returns positive if the number equals 100.

```
fn TestParticles del t = (
if (particleCount $Spray01) == 100
then 1  else 0 )
```

Testing an Atmospheric Property

This sample script tests the Density parameter of a fog effect, and returns positive if it equals 50.

```
fn TestAtmos del t = (
atmos_fog = getAtmospheric 1
print atmos_fog.density -- to:debug
if (atmos_fog.density == 50)
then 1 else 0 )
```

Note the second line, which assigns the fog atmospheric to a variable named "atmos_fog". This is necessary only for atmospheric effects; with most standard objects, you simply use the object name preceded by a $, as in the two previous examples. The "1" following the getAtmospheric command refers to the atmospheric's position in the Rendering Effects dialog > Effects list.

Once you've executed this assignment, you can obtain a list of the atmospheric's properties by entering this command in the MAXScript Listener:

```
ShowProperties atmos_fog
```

Also, the third line in the sample script isn't necessary for the cognitive controller; it simply prints the result of the test in the Listener window for debugging purposes.

Testing Another Delegate's Behavior

You might want to determine in a transition script which behavior is currently influencing a certain delegate. Crowd provides a MAXScript-based method for doing this. You can even check whether a particular delegate is specified as a target within that behavior. An example would be a cocktail party scene in which Betty avoids Harry if Harry is seeking Sally. But if Harry is avoiding Sally, then Betty will seek Harry.

The following example script is taken from the sample file *party.max*, which you can find in the *Biped* download available on this page, under 3ds Max 8 Tutorials – Scene and Support Files > Specialized. The scene uses a more complex scenario than the example described in the previous paragraph. Following is an overview, but to fully understand the setup, you should examine the scene. Study, in particular, the behavior assignments and cognitive controllers, which use a total of eight different transition scripts.

Six delegates are confined in a "room" defined by four grids, using a Wall Repel behavior. Delegates 1, 2, 3, and 5 simply wander at random during the simulation. However, delegate 4 uses a cognitive controller (cc1) that tells it to start wandering, and then switch to one of three Avoid behaviors if members of one of three arbitrary pairs of delegates come within 50 units of each other. Each of the Avoid behaviors targets a different group of three delegates, two of which include delegate 2. Delegate 6 is assigned a second cognitive controller (cc2) that uses the following script to tell it to switch to an Avoid behavior if delegate 4 is avoiding delegate 2. The heart of the script is this line in function *transfunc4*:

Note: See the online *User Reference* for this sample code.

Load the file, press F11 to open the Listener window, and then solve. The Listener window

displays a message whenever delegate 4 is found to be avoiding delegate 2.

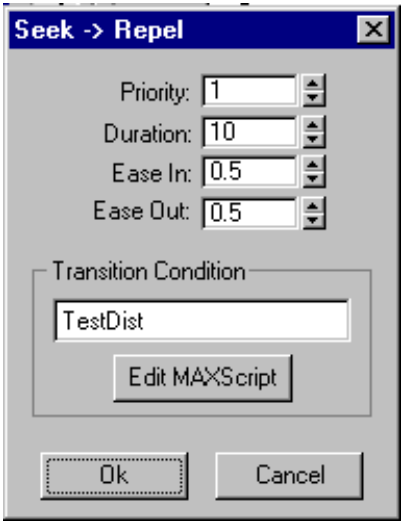
You can use this script as is in your own simulations to check for whether one delegate is avoiding a second by substituting the delegates' names in the above line, and also substituting the names of your Avoid behaviors in the list in *transfunc4*, adding or deleting lines as necessary.

The example script illustrates a second important point: Cognitive controller transition scripts can contain multiple functions. Crowd first executes the function specified in the State Transition dialog > Transition Condition field, and that function calls one or more additional functions in the script, which, of course, can also call functions. In this case, *transfunc4* calls the first function, *isDelegateAvoiding*, passing it three parameters.

Lastly, the script contains a special function, *getBehaviorType*, that compares an input behavior against a list of known behaviors, and on a match, returns the known behavior. In this case, *transfunc4* runs through the list of behaviors currently influencing Delegate04, testing each with *getBehaviorType*, and if an Avoid behavior is in effect, proceeds to check whether Delegate02 is an obstacle of that Avoid behavior. Use of this function is more efficient and flexible than testing for specific behaviors, especially if your scene contains many behaviors of the same type, or you're constantly editing behavior settings. You can see the returned behaviors by removing the comment (double hyphen) from the beginning of the following line in *transfunc4*.

Note: See the online *User Reference* for this sample code.

Interface



Priority—Sets the transition's precedence.

When more than one transition tests true, the software uses the Priority setting to determine which transition occurs. It performs the transition with the lowest Priority setting. Thus, for example, a transition with a Priority setting of 0 takes precedence over one with Priority 1, and so on.

Duration—The number of frames the software takes to effect the transition between states.

Ease In—The rate at which the transition begins. Default=0.5. Range=0 to 1.0.

Lower values cause a more abrupt transition, while higher values cause a more gradual transition.

Ease Out—The rate at which the transition ends. Default=0.5. Range=0 to 1.0.

Lower values cause a more abrupt transition, while higher values cause a more gradual transition.

Note: The sum of the values for Ease In and Ease Out must be less than or equal to 1.0. The software won't let you set a value for either parameter that would cause the sum to exceed 1.0. To increase the value of one parameter when its value equals

1.0 minus the other parameter, decrease the other parameter first.

Transition Condition—The name of the MAXScript function that specifies when/how the transition is to occur.

This name must also appear at the beginning of the main function in the script, after "fn". The script can contain additional functions that are called by the main function and each other.

Edit MAXScript—Opens an editor window for editing, saving, and loading the transition's MAXScript script.

OK—Closes the dialog and implements changes.

Cancel—Closes the dialog and ignores changes.

Behavior Rollout

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button

Select a Crowd object. > Modify panel > Setup rollout > New button

Select a Crowd object. > Modify panel > Setup rollout > Choose a behavior.

The Behavior rollout appears in the Crowd object command panel after you first add a behavior to the scene using the Setup rollout. The rollout's full name (for example: Avoid Behavior or Seek Behavior) depends on the current behavior, displayed in the text box in the Setup rollout. To display the rollout for a different behavior, choose the behavior from the drop-down list.

Note: If you add the first behavior to the scene from the *Behavior Assignments and Teams Dialog (page 2–1200)*, a behavior rollout does not automatically appear in the Crowd command panel. You must first choose the behavior from the drop-down list at the bottom of the Setup rollout.

Use the controls in this rollout to modify the behavior. Following is a list of available behaviors:

Avoid Behavior (page 2–1211)

Orientation Behavior (page 2–1214)

Path Follow Behavior (page 2–1216)

Repel Behavior (page 2–1218)

Scripted Behavior (page 2–1220)

Seek Behavior (page 2–1220)

Space Warp Behavior (page 2–1221)

Speed Vary Behavior (page 2–1222)

Surface Arrive Behavior (page 2–1223)

Surface Follow Behavior (page 2–1226)

Wall Repel Behavior (page 2–1227)

Wall Seek Behavior (page 2–1229)

Wander Behavior (page 2–1231)

For a detailed description of specific behaviors, refer to the above topics. For an overall look at behaviors, see *Assigning Behaviors (page 2–1159)*.

Behaviors

Avoid Behavior

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Avoid Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Avoid Behavior

The Avoid behavior lets you specify any object or objects that delegates must keep away from.

As delegates approach designated objects during the crowd simulation, they steer clear of them while turning and/or braking as necessary. This behavior uses three different methods to let delegates avoid each other and other objects: Steer To Avoid, Repel, and Vector Field.

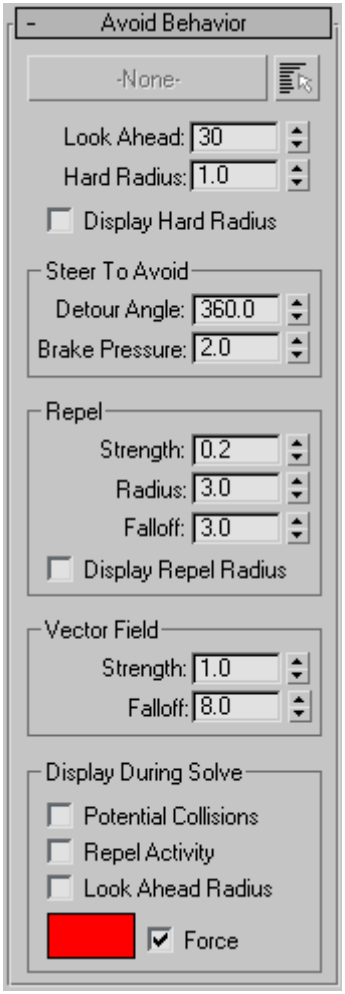
Steer to Avoid behavior is best used for animals that steer around each other at close proximity. Earthbound animals and fish typically do this. Steering motion may be sudden since its action is often engaged for relatively short periods of time.

By contrast, Repel avoidance behavior mimics the continuous action of a repellent magnetic field. Birds, bats and flying insects are best animated with large Repel fields so that they can smoothly avoid each other while maintaining a comfortable margin of error. Repel forces prevent intrusion from all sides, regardless of the direction of travel. Thus even animals that rely mainly on Steer to Avoid will also need some degree of Repel avoidance to maintain spatial separation when they are moving through dense traffic. The forces of Repel avoidance are always directed uniformly outward in a spherical shape.

Use Vector Field avoidance for cases where crowd members must avoid hitting the more complex shapes of arbitrary 3ds Max objects. The outward forces of the Vector Field avoidance can constructed to form the shape of any object. For example, suppose you want to animate a school of fish swimming around a sunken ship. In this case, a vector field can be created so that it extends the shape of the ship into the surrounding space. The field is computed by scan converting the ship's surface normals into a 3D lattice that surrounds the ship. These normals will extend into space as "beacons" in the 3D lattice, telling the fish how to best swim away from the shape of the ship. As the fish enter the space of the vector field lattice, they can be precisely repelled along an avoidance force directed away from ship's surface.

Note: In the explanations that follow, the word "target" is used to refer to the object or objects to be avoided.

Interface



None (label)—Specifies a single target. Click this button, and then click the target object in the viewport. The target name then appears on the button.

If you've selected multiple targets using Multiple Selection (see next item), the word Multiple appears on the button. To see which objects are

designated as targets, click the Multiple Selection button.

Multiple Selection—Opens the Select dialog to let you designate multiple targets. When you have more than one target, you can set delegates to move toward the closest target in the group, or to a computed average of the target positions.

Look Ahead—The number of frames in advance of the current frame that the software looks for potential collisions. Default=30.

Hard Radius—Distance from the target's pivot point, in multiples of the delegate's bounding sphere, where no penetration should occur. Default=1.0.

Tip: Because the hard radius' center is the pivot point, Avoid may not work as expected with target objects whose pivot point is not centered, such as the box primitive. For best results, use the Hierarchy panel > Adjust Pivot rollout controls to center the pivot to the object.

Display Hard Radius—Enables display of a wireframe sphere that depicts the extent of the Hard Radius setting. Default=off.

Steer To Avoid group

Steer To Avoid is used by delegates to steer precisely around anticipated future collisions based on the delegates' current speed and direction. Delegates using this approach can pass very close to one another.

Detour Angle—Maximum necessary turning angle relative to the direction of delegate's goal that delegate will steer to avoid rather than slow down and wait. Default=360. Range=0-360.

Tip: To disable turning for avoidance, thus allowing only braking, set Detour Angle to 0. This forces delegates to remain directed toward their goal so that they must slow down and wait until there is a clearing in front of them, much like an audience queuing to leave through an exit after a concert.

Brake Pressure—Determines how strongly a delegate will react when it encounters an avoided object. Higher values make the delegate more likely to slow down or stop. Lower values will cause delegates to look for a way around the obstacle so they can keep going, sometimes causing delegates to veer off in unexpected directions. Default=2.0.

Repel group

Repel is a general separation force that is based only on the spatial position. Delegates use this to keep from getting into situations where they might side-swipe each other or where they might get so close that Steer To Avoid is too difficult to achieve.

Strength—Determines the strength of the repelling force; higher values result in greater repulsion force. Default=0.2. Range=0.0 to 1.0.

Radius—Maximum distance from delegate's bounding sphere within which "repel" avoidance is sensed and carried out. Default=3.0.

Falloff—The rate at which the strength diminishes between the Repel radius and the hard radius. A value of 1.0 indicates a linear falloff. Higher values cause the strength to fall off to zero more rapidly with distance, thus focusing its effect closer to the delegate's hard radius. Lower values reduce the rate of diminishment, with a Falloff setting of 0.0 indicating that the strength is the same at the Radius distance as it is at the Hard Radius. Default=3.0.

Display Repel Radius—Enables display of a wireframe sphere that depicts the extent of the Repel setting. Default=off.

Vector Field group

If you've applied a *Vector Field space warp (page 2–1241)* to an object in your scene, you can specify the vector field as an object to avoid. The distinction is this: When used with the *Space Warp behavior (page 2–1221)*, delegates use the vector

field to steer around the object by being guided to travel perpendicular to the field's vectors. When used with the Avoid behavior, the delegate simply moves away in the direction of the vectors.

Tip: Sometimes when using Avoid with a vector field, the behavior might seem to be "fighting" with other behaviors (such as Seek) over delegate movement, causing a halting and/or wavering motion. In such cases, try reducing Brake Strength and/or increasing Falloff.

Strength—Higher values result in more powerful influence. Delegates will be directed to move perpendicular to the field. Default=1.0. Range=0.0 to 1.0.

Falloff—Higher values cause vector field influence to fall off to zero more rapidly with distance, thus focusing its effect closer to the delegate's hard radius. Default=8.0.

Display During Solve group

Use these switches for debugging a crowd simulation. During the solve, they display information about the simulation using graphical metaphors for different aspects of the Avoid behavior.

Potential Collisions—Displays a green line from the delegate to the location of a potential collision. Default=off.

Repel Activity—Displays a white line between the delegate and target when the repel force is in effect. Default=off.

Look Ahead Radius—Displays a sphere that shows the current distance used to check for potential collisions.

Color Swatch—Shows the color used to draw the Avoid force vector during the solution. Click the box to choose a different color. Default=red.

Force—When on, force exerted on one or more delegates by the Avoid behavior is drawn in the viewports as a colored line during the simulation solution. Default=on.

Orientation Behavior

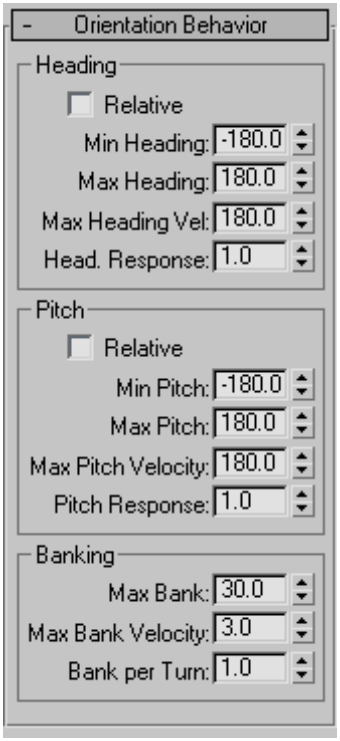
Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Orientation Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Orientation Behavior

The Orientation behavior lets you control whether and how delegates rotate, independent of their direction of motion. Normally, a delegate always faces in the direction it's moving. You can use the Orientation behavior to specify limits to the delegate's rotational activity without affecting its path, which is generated by other behaviors. Use these settings, for example, to keep delegates facing in one direction while moving in another.

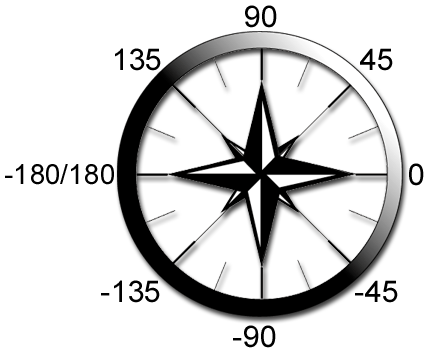
Note: These settings do not affect the path a delegate takes, which is produced by other behaviors such as Seek and Avoid. They influence only the direction it faces as it traverses the path.

Interface



Heading group

Use these controls to affect how delegates turn on the vertical axis. By default, heading is absolute, with 0 specifying the positive X axis in World coordinates. Thus, -90 would specify the negative Y axis, 90 the positive Y axis, and 180 or -180 the negative X axis.



Absolute heading as viewed from the top

For example, if you wanted a delegate to be able to turn between the positive X axis and the positive Y axis, you would set Max Heading to 0 and Min Heading to 90.

You can also specify heading limits in amounts relative to the delegate's heading at the time that the Orientation behavior takes effect by turning on the Relative check box.

Relative—When on, Heading settings are applied relative to the delegate's heading at the time the behavior takes effect. When off, settings are absolute. Default=off.

Min Heading—The minimum permissible heading. This number should be lower than the Max Heading value. Default=180. Range=-180 to 180.

Max Heading—The maximum permissible heading. This number should be higher than the Min Heading value. Default=180. Range=-180 to 180.

Max Heading Vel(ocity)—Specifies how much the delegate's heading can change per frame. This controls angular acceleration and deceleration. Default=180.

Head. Response—Determines how quickly the heading follows the direction the object is

moving in. A value of 1.0 indicates maximum responsiveness, and will point in the direction the delegate is moving (within the limits) while a lower value means that it is less responsive. Default=1. Range=0 to 1.

Pitch group

Use these controls to affect how delegates turn on the left-right axis.

Relative—When on, Pitch settings are applied relative to the delegate's pitch at the time the behavior takes effect. When off, settings are absolute. Default=off.

Min Pitch—The minimum number of degrees a delegate can incline or decline. This number should be lower than the Max Pitch value. Default=-180. Range=-180 to 180.

Max Pitch—The maximum number of degrees a delegate can incline or decline. This number should be higher than the Min Pitch value. Default=180. Range=-180 to 180.

Max Pitch Velocity—Specifies how much the delegate's pitch can change per frame. This controls angular acceleration and deceleration. Default=180.

Pitch Response—Determines how quickly the pitch follows the direction the object is moving in. A value of 1.0 indicates maximum responsiveness, so that will point in the direction the delegate is moving (within the limits) while a lower value means that it is less responsive. Default=1. Range=0 to 1.

Banking group

Use these controls to affect how delegates turn on the in-out axis. Banking is primarily a result of heading changes.

Max Bank—The maximum number of degrees the delegate can bank. Default=30.0.

Max Bank Velocity—The maximum number of degrees the delegate's bank angle can change per frame. This controls angular acceleration and deceleration. Default=3.0.

Bank per Turn—The number of degrees the delegate will bank as a function of the turn angle at the current frame. For example, if Bank per Turn=1, the delegate will bank one degree for every degree it is turning at a given frame. Default=1.0.

Path Follow Behavior

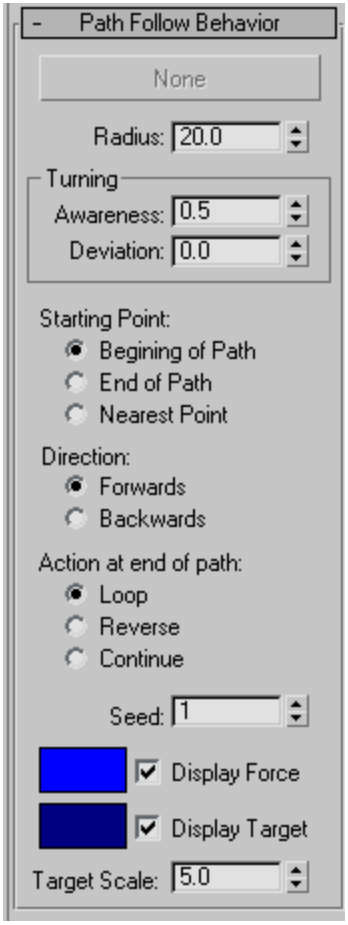
Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Path Follow Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Path Follow Behavior

The Path Follow behavior lets you direct delegates to traverse a specified path during a crowd simulation. Delegates can move forward or backward along paths, and when they reach the end, they can loop back to the start or reverse direction, or even continue in the same general direction.

If the delegate's start position isn't on the path at the start of the simulation, it moves to the path before following the path. During the solution, **character studio** intermittently displays an optional *target* icon to show the delegate's immediate goal; this changes as the simulation proceeds.

Interface



None (label)—Click this button, then select a path object. Suitable path objects include splines and NURBS curves. If a path object contains more than one spline or curve, **character studio** uses the lowest-numbered element (usually the earliest created one).

Note: You can assign a path object only from the Modify panel.

Radius—The radial distance from the path, in units, within which the delegate stays while traversing the path. Default=20.0. Range=0.0 to 99,999.0.

Turning group

These parameters determine how delegates turn while following the path. Awareness determines how well a delegate anticipates turns in the path as it moves; you can apply random variation to Awareness with the Deviation setting.

Awareness—Specifies how "intelligent" the delegate is while traversing this path. A high Awareness setting means that it takes into account the curve of the path while moving and will try to anticipate changes. A low value for Awareness, on the other hand, means that the delegate notices the path only when leaving it. Default=0.5. Range=0.0 to 1.0.

Note: You can randomize awareness behavior with the Deviation and Seed settings.

Deviation—Specifies the maximum amount by which Awareness should vary. **character studio** takes a random number between the negative and positive values of the Deviation setting, multiplies it by the Awareness setting, and adds the result to Awareness. Default=0.5. Range=0.0 to 1.0.

Note: You can vary behaviors among different Path Follow behaviors that use the same Awareness and Deviation settings by changing the Seed value.

Starting Point

Determines where on the path the delegate begins to follow the path. The default choice is Beginning of Path.

Hint: To see a selected spline path's start point, open the Modify panel and turn on any sub-object level; the start point is represented with a unique indicator. Also in the Modify panel, with closed curves, you can see the vertex ordering at any sub-object level by turning on Selection rollout > Display group > Show Vertex Numbers. To see a NURBS curve's start point, go to the Curve sub-object level; the start point is indicated by a small green circle.

Beginning of Path—The delegate first moves to the start of the path before following it.

End of Path—The delegate first moves to the end of the path before following it. With closed curves, this is the same point as the beginning of the path.

Nearest Point—The delegate first moves to the closest point on the path and then follows the path from there.

Direction

Determines the direction the delegate takes initially when following the path. The default choice is Forwards.

Forwards—The delegate moves along path vertices in ascending order.

Backwards—The delegate moves along path vertices in descending order.

Action at End of Path

Determines what the delegate does when it reaches the path end. The default choice is Loop.

Loop—The delegate loops around the path, even if it isn't closed. If Beginning of Path or End of Path is chosen, it returns to the path's start or end point each time it finishes traversing the path. If Nearest Point is chosen, it returns to an arbitrary point determined by its position and the path shape.

Reverse—The delegate reverses direction at the end of the path. Use this choice to simulate a back-and-forth "patrol" behavior.

Continue—The delegate continues moving in the same direction it faced at the end of the path until the simulation ends or it's acted upon by another force or behavior.

Seed—Specifies a seed value for randomizing Awareness. Default=1.

Color Swatch—Shows the color used to draw the Path Follow force vector during the solution. Click the box to choose a different color. Default=blue.

Display Force—When on, force exerted on the delegate(s) by the Path Follow behavior is drawn in the viewports as a vector during the simulation solution. Default=on.

Color Swatch—Shows the color used to draw the target icon. Default=dark blue.

Display Target—Enables display of the target icon, which appears during the solution when a new interim goal is calculated for the delegate. Default=on.

Target Scale—Specifies the overall size of the target icon. Default=5.0.

Repel Behavior

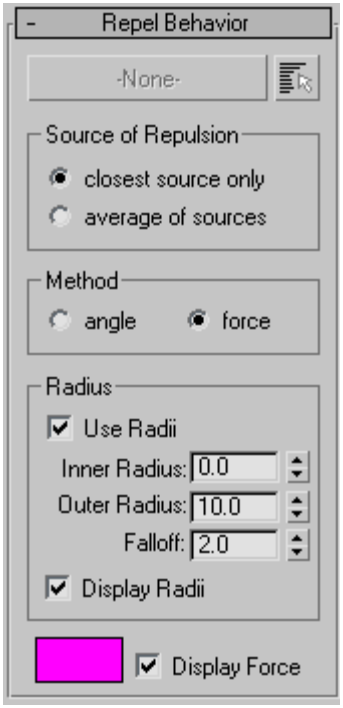
Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Repel Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Repel Behavior

The Repel behavior lets you specify any object or objects (sources) that will force delegates to move away from them. This is basically the opposite of the *Seek behavior (page 2–1220)*. If you want delegates to back away from an object, as opposed to turning to face the direction they're moving, use Repel in conjunction with the *Orientation behavior (page 2–1214)*.

Note: Repel is set by default to work only within a specific radius around the source. If you want it to work at any distance, turn off Radius group > Use Radii.

Interface



None (label)—Specifies a single source. Click this button, and then click the target object in the viewport. The target name then appears on the button.

If you've selected multiple sources using Multiple Selection (see next item), the word Multiple appears on the button. To see which objects are designated as sources, click the Multiple Selection button.

Multiple Selection—Opens the Select dialog to let you designate multiple sources. When you have more than one source, you can set delegates to move toward the closest target in the group, or to a computed average of the source positions.

Source of Repulsion group

Determines repel activity when the behavior uses multiple sources. The default choice is Closest Source Only.

Closest Source Only—Each delegate is repelled by the closest of the assigned sources. Use this to have delegates assigned a single Repel behavior move away from sources in different directions.

Average Of Sources—All delegates move away from a common point determined by averaging all sources' locations.

Method group

Determines whether delegate direction as influenced by the behavior is calculated by an angular method or a force method. Default=Force.

Angle—Applies a force to the delegate based on the angle between the delegate's current direction and the direction it would need to take in order to be moving directly away from the source.

The magnitude of the force is greatest when the delegate is moving directly towards the source, and needs to turn around. It can be as little as 0 when the delegate is moving directly away from the source.

Force—Always applies a force directly away from the source. The magnitude of the force is constant.

Radius group

Use the Radius settings to activate the Repel behavior only when the delegates are within a specific distance from the target. The relative strength of the force increases from 0 percent at the outer radius to 100 percent at the inner radius.

Use Radii—When on, the behavior applies only to delegates closer to the target than the Outer Distance value. Default=on.

Inner Radius—The distance from the target at which the force is applied at full strength. Default=0.0.

Outer Radius—The distance from the target at which the force begins to be applied. Default=10.0.

Falloff—Default=2.0.

Display Radii—The radii are displayed when the force is active.

Color Swatch—Shows the color used to draw the Repel force vector during the solution. Click the box to choose a different color. Default=violet.

Display Force—When on, force exerted on the delegate(s) by the Repel behavior is drawn in the viewports as a vector during the simulation solution.

Scripted Behavior

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Scripted Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Scripted Behavior

The Scripted behavior lets you create custom behaviors using MAXScript. A scripted behavior can incorporate one of three behavior types: Force, Constraint, or Orientation.

For detailed information on scripting behaviors, see the MAXScript reference, available from the Help menu.

Interface



Behavior Type—Choose Force, Constraint, or Orientation.

Script Context Name—Specify a name for the script.

Edit MAXScript—Click to open an editor window.

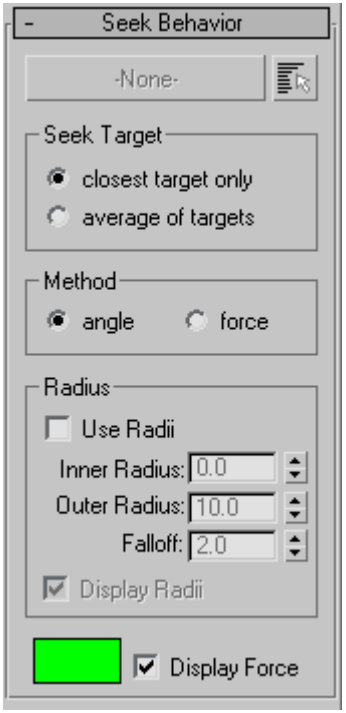
Seek Behavior

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Seek Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Seek Behavior

The Seek behavior lets you specify any object or objects as a stationary or moving target for delegates. Delegates move toward the target during the crowd simulation while turning as necessary.

Interface



None (label)—Specifies a single target. Click this button, and then click the target object in the viewport. The target name then appears on the button.

If you've selected multiple targets using Multiple Selection (see next item), the word Multiple appears on the button. To see which objects are designated as targets, click the Multiple Selection button.

 **Multiple Selection—**Opens the Select dialog to let you designate multiple targets. When you have more than one target, you can set delegates to move toward the closest target in the group, or to a computed average of the target positions.

Seek Target group

Determines seek activity when the behavior uses multiple targets.

Closest Target Only—Each delegate seeks the closest of the assigned targets. Use this to have delegates assigned a single Seek behavior move in different directions.

Average Of Targets—All delegates move toward a common point determined by averaging all targets' locations.

Method group

Determines whether delegate direction as influenced by the behavior is calculated by an angular method or a force method. Default=Angle.

Angle—Applies a force to the delegate based on the angle between the delegate's current direction and the direction it would need to take in order to be moving directly toward the target.

The magnitude of the force is greatest when the delegate is moving away from the target, and needs to turn around. It can be as little as 0 when the delegate is directly approaching the target.

Force—Always applies a force directly towards the target. The magnitude of the force is constant.

Radius group

Use the optional radius settings to activate the Seek behavior only when the delegates are within a specific distance from the target. The relative strength of the Seek behavior increases from 0 percent beyond the outer radius to 100 percent at the inner radius.

Use Radii—When on, the Seek behavior applies only to delegates less than the Outer Radius distance from the target. Default=off.

Inner Radius—The distance from the target at which Seek is applied at full strength. Default=0.0.

Outer Radius—The distance from the target at which Seek begins to be applied. Default=10.0.

Falloff—Default=2.0.

Display Radii—The radii are displayed when the force is active.

Color Swatch—Shows the color used to draw the Seek force vector during the solution. Click the box to choose a different color. Default=green.

Display Force—When on, force exerted on the delegate(s) by the Seek behavior is drawn in the viewports as a vector during the simulation solution. If Use Radii is turned on, the radii are also displayed when the force is active.

Space Warp Behavior

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Space Warp Behavior

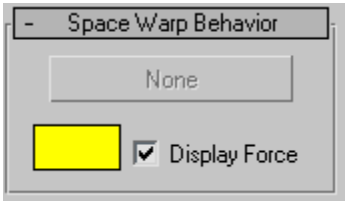
Select a Crowd object. > Modify panel > Setup rollout > New button > Space Warp Behavior

The Space Warp behavior lets space warps, such as wind or gravity, influence a crowd simulation. You can use it to apply space warps in the Forces

and Particles & Dynamics categories to crowd members.

In particular, use the Space Warp behavior to tie delegates to a *Vector Field space warp (page 2–1241)*, so that they avoid penetrating an irregularly shaped object while following its contours.

Interface



None (label)—Click this button, then select a space warp object.

Color Swatch—shows the color used to draw the Space Warp force vector during the solution. Click the box to choose a different color.

Display Force—When on, force exerted on the delegate(s) by the Space Warp behavior is drawn in the viewports as a vector during the simulation solution.

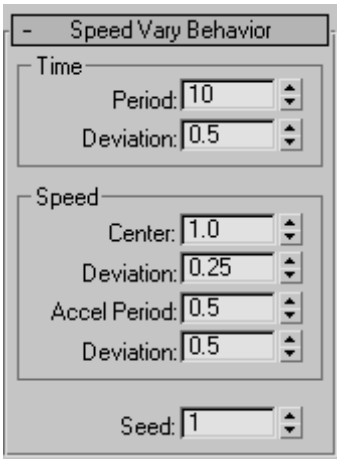
Speed Vary Behavior

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Speed Vary Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Speed Vary Behavior

The Speed Vary behavior is useful for objects whose velocity changes at random as they move, such as sightseeing tourists. Its parameters let you specify how often a delegate should change speed, what speed range it should look at for a new speed, and how long it should accelerate to get to that new speed.

Interface



Time group

Period—Specifies how many frames should elapse before a new speed is chosen.

Deviation—Specifies the maximum amount by which Period should vary.

Each time a period ends, **character studio** takes a random number between the negative and positive values of the Deviation setting, multiplies it by the Period setting, and adds the result to Period. Default=0.5. Range=0.0 to 1.0.

Speed group

The mathematical formula for the new speed is as follows: new speed = (delegate's Average Speed*Center)*(1 + RN*Center Deviation), where RN is a random number between -1 and 1.

Center—Specifies the speed the delegate should change to. Center is a multiplier: A value of 0.0 means to stop, a value of 1.0 means to move at its *average speed*, and a value greater than 1.0 means to move faster than its average speed. Default=1.0. Range=0.0 to 99,999.0.

Deviation—Specifies the maximum amount by which the delegate's calculated speed (Average Speed*Center) should vary.

Each time a period ends, **character studio** takes a random number between the negative and positive values of the Deviation setting, multiplies it by the calculated speed, and adds the result to the calculated speed. Default=0.25. Range=0.0 to 99,999.0.

Accel Period—Specifies the rate at which the delegate's speed should change in relation to the period length.

A value of 0.0 means that the transition to the new speed will proceed as quickly as possible, and a value of 0.5 means the transition will take half the period. A value of 1.0 means the transition will take the entire period. Default=0.5. Range=0.0 to 1.0.

Deviation—Specifies the maximum amount by which acceleration should vary.

Each time a period ends, **character studio** takes a random number between the negative and positive values of the Deviation setting, multiplies it by the Acceleration setting, and adds the result to Acceleration. Default=0.5. Range=0.0 to 1.0.

Seed—Specifies a value for randomizing the Speed Vary behavior.

Surface Arrive Behavior

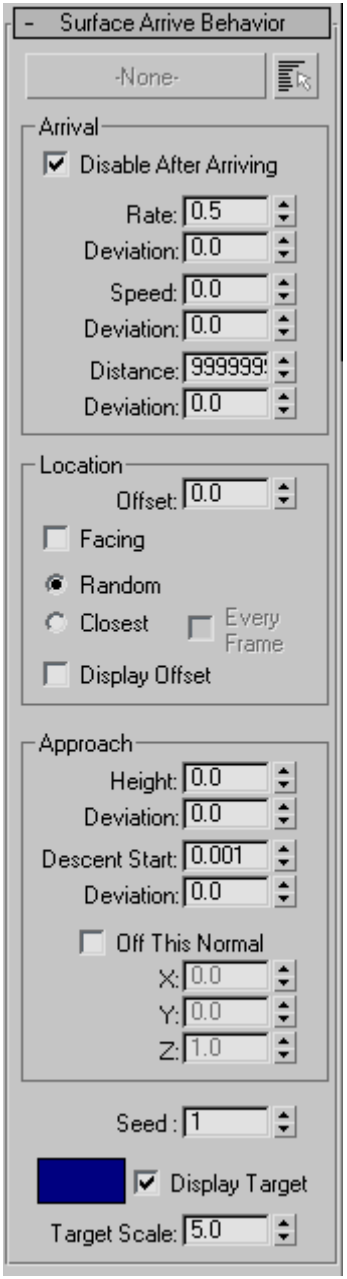
Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Surface Arrive Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Surface Arrive Behavior

The Surface Arrive behavior is similar to the *Seek behavior (page 2–1220)* in that it lets you specify an object or objects as a stationary or moving target for delegates. The principal difference is that Surface Arrive can cause delegates to stop when they reach the target. You can also specify, to some degree, where the delegate will stop on the object, and how it will approach the target before stopping.

An example of a use for this behavior would be birds flying over a row of telephone poles, then each one landing on top of a different pole.

Interface



None (label)—Specifies a single target. Click this button, and then click the target object in the viewport. The target name then appears on the button.

If you've selected multiple targets using Multiple Selection (see next item), the word Multiple appears on the button. To see which objects are designated as targets, click the Multiple Selection button.

Multiple Selection—Opens the Select dialog to let you designate multiple targets.

When you have more than one target, you can set delegates to move toward the closest target in the group, or to a computed average of the target positions.

Arrival group

Specifies three aspects of the Surface Arrive behavior: Rate, Speed, and Distance.

Disable After Arriving—When on, turns off the Surface Arrive behavior after the delegate arrives at the surface. Default=on.

Rate—A multiple of the delegate's *Max Accel* setting that specifies the acceleration with which it will try to arrive. A value of 1.0 means to use the full acceleration of the delegate. Default=0.5.

Deviation—Adds random variation to the to the Rate setting.

The actual deviation is calculated by multiplying the Deviation setting by a random number between -1 and 1, and then multiplying the result by the Rate setting. Default=0.0.

Speed—The speed at which to arrive, relative to the speed of the target. Default=0.0.

Deviation—Adds random variation to the Speed setting. Default=0.0.

The actual deviation is calculated by multiplying the Deviation setting by a random number between -1 and 1, and then multiplying the result by the Speed setting.

Distance—The maximum radial distance from the target within which the behavior will be active. Default=9999999.0.

Until the delegate is within this radius, the behavior has no influence.

Deviation—Adds random variation to the to the Distance setting.

The actual deviation is calculated by multiplying the Deviation setting by a random number between -1 and 1, and then multiplying the result by the Distance setting. Default=0.0.

Location group

Offset—Specifies a consistent distance from the calculated arrival point, based on the surface normal, for the delegate to use. Default=0.0.

Facing—When on, the delegate will try to arrive only at points on triangles on the surface that are facing it. Default=off.

Random—The software chooses a random point on the target surface as the arrival point.

When using the Random option, the software chooses arrival points for delegates once, at the beginning of the simulation. This is the default choice

Closest—The software chooses the closest point on the target surface as the arrival point.

If Closest is chosen, but Every Frame is off, the software chooses arrival points for delegates once, at the beginning of the simulation.

Every Frame—When on, the software chooses arrival points for delegates at every frame. Available only when Closest is chosen. Default=off.

Every Frame is useful when the target object is rotating during the animation, but requires more time for calculation.

Display Offset—When on, shows the Offset distance as lines emanating from each vertex in the surface object, perpendicular to the surface.

Approach group

The Height and Descent settings together specify the path the delegate will take for its arrival. They allow for a wide range of behavior, from soft, gradual landings to direct helicopter-type descents.

In both cases, the actual deviation is calculated by multiplying the Deviation setting by a random number between -1 and 1, and multiplying the result by the relevant value.

Height—Specifies a distance from the arrival point along its face normal.

This is the point that the delegate will go to first before descending to the arrival point.

Deviation—Adds random variation to the to the Height setting.

The actual deviation is calculated by multiplying the Deviation setting by a random number between -1 and 1, and then multiplying the result by the Height setting.

Descent Start—Specifies the distance between the delegate and the arrival point at which the descent should start.

Note: Be careful that Descent Start is set high enough that the delegate won't overshoot when descending because its speed is too high and deceleration too low, compared to when it should start descending.

Deviation—Adds random variation to the to the Descent Start setting.

The actual deviation is calculated by multiplying the Deviation setting by a random number between -1 and 1, and then multiplying the result by the Descent Start setting.

Off This Normal—When on, lets you set an approach vector to specify the angle at which the final approach occurs. Default=off.

X/Y/Z—Use these settings to specify the final approach vector in world coordinates. For example, the vector specified by the default settings of X=0, Y=0, Z=1 means that the delegates will approach the target along the vertical world axis.

Seed—Affects the random numbers used to calculate the Deviation settings.

For similar randomization among different Surface Arrive behaviors, use the same Seed value.

Color Swatch—Shows the color used to draw the target icon. Default=dark blue.

Display Target—Enables display of the target icon, which appears during the solution when a new interim goal is calculated for the delegate. Interim goals are created when using the Approach group settings. Default=on.

Target Scale—Specifies the overall size of the target icon. Default=5.0.

Surface Follow Behavior

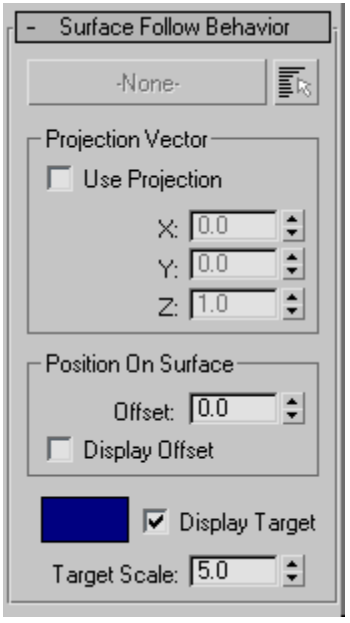
Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Surface Follow Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Surface Follow Behavior

The Surface Follow behavior moves delegates with respect to object surfaces, which can be still or animated. For example, you can apply an animated Noise modifier to a patch grid to simulate a choppy water surface, and objects guided by Surface Follow will stay on top.

Note: By default, a delegate influenced by Surface Follow picks a direction to move in at any given frame based on its current facing and the plane of the face it's currently over. Thus objects moving up a hill, while seeking a point at the bottom of the other side of the hill, tend to turn left or right to skirt the hill, rather than following the upward slope. You can override this with the Projection Vector option.

Interface



None (label)—Specifies a single "target" object to use as a surface. Click this button, and then click the target object in the viewport. The target name then appears on the button.

If you've selected multiple targets using Multiple Selection (see next item), the word Multiple appears on the button. To see which objects are designated as targets, click the Multiple Selection button. The Select dialog appears with designated targets highlighted.

Multiple Selection—Opens the Select dialog to let you designate multiple targets.

When you have more than one target, delegates initially move toward the closest target in the group, and then move over its surface until they encounter another target, at which point they switch to its surface, and so on.

Projection Vector group

These controls let you override the default direction calculated by the Surface Follow behavior by describing a virtual plane along which the delegate is to move. You do this by specifying a vector, in world coordinates, that's perpendicular to the desired virtual plane.

For example, if you want the delegate, when it encounters a hill, to keep moving forward, straight up and over the hill, instead of skirting it, you would use the default Projection Vector settings: X=Y=0, Z=1. This vector is aligned with the world Z (vertical) axis, so it specifies a plane parallel to the world XY plane. Thus, the delegate always moves straight ahead while following the surface.

Use Projection—When on, Surface Follow calculates delegate direction from the specified vector, rather than using the default.

X/Y/Z—Specifies a vector using world coordinates. Default=X=Y=0, Z=1. Range=-1.0 to 1.0.

If only one of these settings is not 0, then the projection vector is aligned with the non-zero axis. Combine non-zero settings to create angled planes for Surface Follow. For example, to create a virtual plane that's rotated 45 degrees clockwise about world Y axis, set X=Z=1 and Y=0. Also, while you can set all three axes to 0, that specifies no vector, and so effectively turns off Use Projection.

Position on Surface group

Offset—Specifies the delegate's distance above the surface, using the surface normal. Recalculated at each frame.

Display Offset—When on, shows the Offset distance as lines emanating from each vertex in the surface object, perpendicular to the surface.

Color Swatch——Shows the color used to draw the Surface Follow target (see Display Target, next) during the solution. Click the box to choose a different color.

Display Target——When on, the interim goal for each delegate influenced by the Surface Follow behavior is drawn in the viewports as a wireframe sphere during the simulation solution.

If the delegate starts out away from the surface to be followed, the target is most visible before the delegate reaches the surface; the target is then positioned along the surface edge. While the delegate is actually following the surface, the target is usually coincident with the delegate, because Surface follow sets a new destination only a frame or two ahead.

Target Scale——Sets the target size.

Wall Repel Behavior

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Wall Repel Behavior

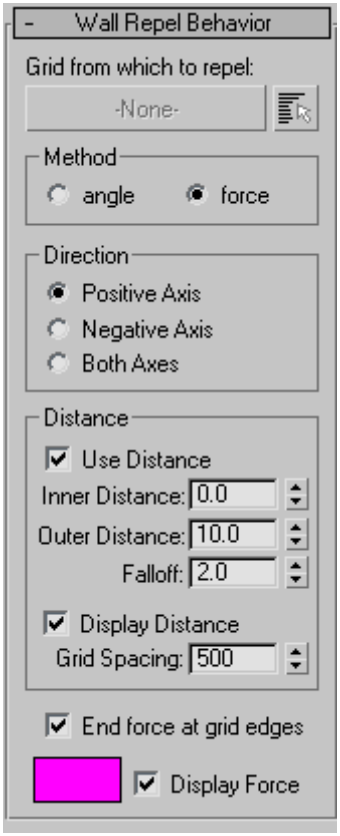
Select a Crowd object. > Modify panel > Setup rollout > New button > Wall Repel Behavior

The Wall Repel behavior uses one or more grid objects to repel delegates. When influenced by the Wall Repel force, delegates turn until they're heading away from the grid. It's useful for keeping objects inside a straight-sided enclosure, such as a room in a building.

You can set the grids to repel from either side or both sides, and optionally specify a maximum distance for repelling. You can also set the behavior to act as though each grid extends infinitely along its plane.

Note: The Wall Repel behavior simply applies a force on the delegate in the direction opposite the wall; it does not guarantee that the delegate won't go through the wall. If the delegate does go through the wall, adjust settings such as Inner and Outer Distance, and, in particular, try reducing Falloff.

Interface



Grid from which to repel——Set a repelling grid ("source") by clicking the None button and then selecting the grid. Thereafter, the grid's name appears on the button.

To set multiple source grids, click Multiple Selection and use the Select dialog. With multiple source grids, the word "Multiple" appears on the large button.

Method group

Determines whether delegate direction as influenced by the behavior is calculated by an angular method or a force method. Default=Force.

Angle—Applies a force to the delegate based on the angle between the delegate's current direction and the direction it would need to take in order to be moving directly away from the source.

The magnitude of the force is greatest when the delegate is moving directly towards the source, and needs to turn around. It can be as little as 0 when the delegate is moving directly away from the source.

Force—Always applies a force directly away from the source. The magnitude of the force is constant.

Direction group

Determines whether the grid repels from its positive-axis side, its negative-axis side, or both.

To determine which is the positive-axis side, select the grid, and then set the reference coordinate system to Local (the default system is View). Look at the grid in a viewport in which it appears edge-on. The side with the axis sticking out is the positive-axis side, and the opposite side is the negative-axis side.

Tip: For ease of setup, when building a "room" from grids, make sure they all point inward (or outward).

Positive Axis—The grid repels only from the positive-axis side.

Negative Axis—The grid repels only from the negative-axis side.

Both Axes—The grid repels from both sides.

Distance group

Use the distance settings to activate the Wall Repel behavior only when the delegates are within a

specific distance from the target. The relative strength of the behavior increases from 0 percent at the outer radius to 100 percent at the inner radius.

Use Distance—When on, the behavior applies only to delegates closer to the target than the Outer Distance value. Default=on.

Inner Distance—The distance from the target at which the force is applied at full strength. Default=0.0.

Outer Distance—The distance from the target at which the force begins to be applied. Default=10.0.

Falloff—The rate at which the repelling force diminishes between the Inner Distance and the Outer Distance. Default=2.0.

A value of 1.0 indicates a linear falloff. Higher values cause the strength to fall off to zero more rapidly with distance, thus focusing its effect closer to the Inner Distance. Lower values reduce the rate of diminishment, with a Falloff setting of 0.0 indicating that the strength is the same at the Outer Distance as it is at the Inner Distance.

Display Distance—Shows the inner and outer distance settings as grids offset from the target grid in the viewports. The Inner Distance grid is light blue, while the Outer Distance grid is blue-white. Default=on.

Grid Spacing—Alters the spacing of grid lines used to draw the Inner/Outer Distance grids. The lower the value, the closer the spacing. Default=500.

End force at grid edges——When on, the force emanates only from the grid object. When off, the force emanates from an imaginary infinite grid created by extending the grid plane in all directions.

Color Swatch——Shows the color used to draw the Wall Repel force during the solution. Click the box to choose a different color. Default=violet.

Display Force——The force, when activated, is drawn in the viewports as a wireframe rectangle during the simulation solution. Default=on.

Wall Seek Behavior

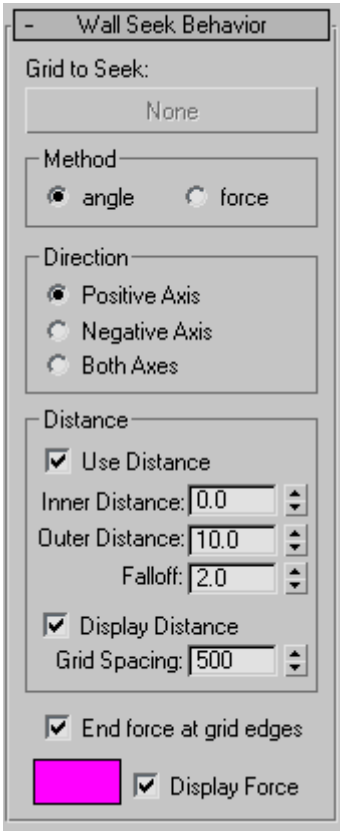
Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Wall Seek Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Wall Seek Behavior

The Wall Seek behavior uses a grid object to attract delegates. When influenced by the Wall Seek force, delegates turn until they're heading toward the grid. It's useful for moving objects toward a rectangular area, such as a doorway.

You can set the grid to attract from either side or both sides, and optionally specify a maximum distance for attraction. You can also set the behavior to act as though the grid extends infinitely along its plane.

Interface



Grid to seek—Set the target grid by clicking the button (initially labeled "None") and then selecting the grid. Thereafter, the grid's name appears on the button.

Method group

Determines whether delegate direction as influenced by the behavior is calculated by an angular method or a force method. Default=Angle.

Angle—Applies a force to the delegate based on the angle between the delegate's current direction and the direction it would need to take in order to be moving directly toward the target.

The magnitude of the force is greatest when the delegate is moving away from the target, and needs to turn around. It can be as little as 0 when the delegate is directly approaching the target.

Force—Always applies a force directly towards the target. The magnitude of the force is constant.

Direction group

Determines whether the grid attracts from its positive-axis side, its negative-axis side, or both.

To determine which is the positive-axis side, select the grid, and then set the reference coordinate system to Local (the default system is View). Look at the grid in a viewport and check the direction of the Z-axis arrow. This arrow determines the positive-axis direction. The opposite direction is the negative-axis direction.

Positive Axis—The grid attracts only from the positive-axis side.

Negative Axis—The grid attracts only from the negative-axis side.

Both Axes—The grid attracts from both sides.

Distance group

Use the distance settings to activate the Wall Seek behavior only when the delegates are within a specific distance from the target. The relative strength of the behavior increases from 0 percent at the outer distance to 100 percent at the inner distance.

Use Distance—When on, the behavior applies only to delegates closer to the target than the Outer Distance value. Default=on.

Inner Distance—The distance from the target at which the force is applied at full strength. Default=0.0.

Outer Distance—The distance from the target at which the force begins to be applied. Default=10.0.

Falloff—The rate at which the attracting force diminishes between the Inner Distance and the Outer Distance. Default=2.0.

A value of 1.0 indicates a linear falloff. Higher values cause the strength to fall off to zero more rapidly with distance, thus focusing its effect closer to the Inner Distance. Lower values reduce the rate of diminishment, with a Falloff setting of 0.0 indicating that the strength is the same at the Outer Distance as it is at the Inner Distance.

Display Distance—Shows the inner and outer distance settings as grids offset from the target grid in the viewports. The Inner Distance grid is light blue, while the Outer Distance grid is blue-white. Default=on.

Grid Spacing—Alters the spacing of grid lines used to draw the Inner/Outer Distance grids. The lower the value, the closer the spacing. Default=500.

End force at grid edges—When on, the force emanates only from the grid object. When off, the force emanates from an imaginary infinite grid created by extending the grid plane in all directions.

Color Swatch—Shows the color used to draw the Seek force vector (and, if used, the radii) during the solution. Click the box to choose a different color. Default=violet.

Display Force—When on, force exerted on the delegates by the Seek behavior is drawn in the viewports as a vector during the simulation solution. If Use Radii is turned on, the radii are also displayed when the force is active.

Wander Behavior

Create panel > Helpers > Object Type rollout > Crowd > Setup rollout > New button > Wander Behavior

Select a Crowd object. > Modify panel > Setup rollout > New button > Wander Behavior

The Wander behavior imparts a random motion to delegates, letting you simulate meandering activity in which delegates move and turn in a haphazard manner. It works by randomly picking a new direction, and then turning and moving in that direction. You can specify how often to pick a new direction, how far to turn, and how fast or slow to turn while moving.

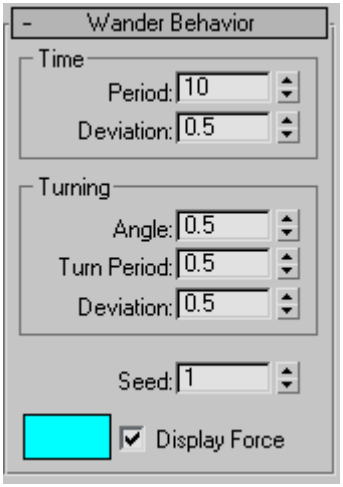
Tip: The trajectory calculated by the Wander behavior varies significantly for each object to which it's applied. To create a group of wandering delegates all using the same trajectory, apply the Wander behavior to a delegate, solve the simulation, and then replicate the delegate using the *Scatter Objects (page 2–1189)* facility with Clone Controllers turned on.

Procedure

To use the Wander behavior:

1. Add a Wander behavior to the Crowd object.
2. Change the default settings as desired. Probably the most important is Period, which sets the number of frames between changes of direction.
3. Use *Behavior Assignments (page 2–1200)* to assign the behavior to a delegate or team.

Interface



Time group

Period—Specifies how many frames should elapse before a new direction is chosen. Default=10.

Deviation—Specifies the maximum amount by which Period should vary. Each time a period ends, **character studio** takes a random number between the negative and positive values of the Deviation setting, multiplies it by the Period setting, and adds the result to Period. Default=0.5. Range=0.0 to 1.0.

Turning group

Angle—Specifies how far to turn when changing direction. A small value means to change direction by only a small amount, while as the value approaches 1.0 the delegate will turn randomly in any direction. Default=0.5. Range=0.5 to 1.0.

Turn Period—Specifies how long over the current period it takes to turn. A value of 0.0 means that the delegate will rotate as quickly as possible to face a direction and then travel in that direction, while a value of 1.0 means the delegate will take the entire period to rotate in that direction. Default=0.5. Range=0.5 to 1.0.

Deviation—Specifies the maximum amount by which Angle should vary. Each time a period ends, **character studio** takes a random number between the negative and positive values of the Deviation setting, multiplies it by the Angle setting, and adds the result to Angle. Default=0.5. Range=0.0 to 1.0.

Seed—Specifies a seed value for randomizing the Wander behavior.

Color swatch—Shows the color used to draw the Wander force vector during the solution. Click the box to choose a different color.

Display Force—When on, force exerted on the delegates by the Wander behavior is drawn in the viewports as a vector during the simulation solution.

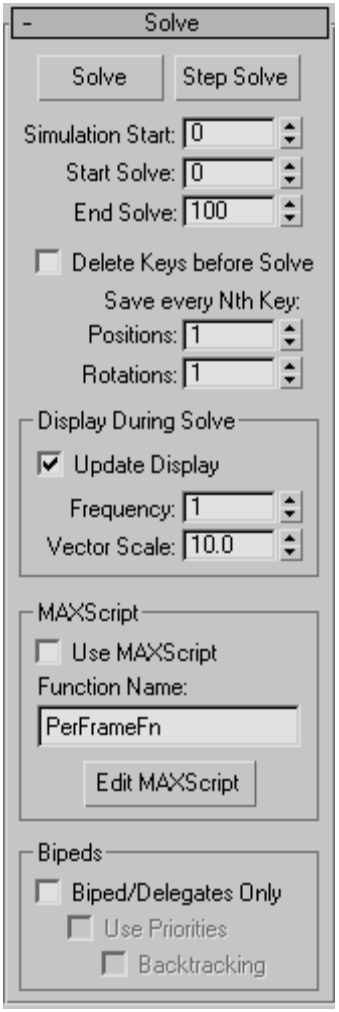
Solve Rollout

Create panel > Helpers > Object Type rollout > Crowd > Solve rollout

Select a Crowd object. > Modify panel > Solve rollout

Once you've set up the crowd simulation, use this rollout to set solution parameters and to solve the simulation. You can solve continuously or a frame at a time, starting at any frame.

Interface



Solve—Runs the crowd simulation continuously, applying all specified behaviors to delegates to which they are assigned. Solving a simulation overwrites any previous solutions.

To abort a solution in progress and save all keys generated up to that point, press the Esc key. Alternatively, with complex simulations, you can save time by pressing Shift+Esc to abort a solution without saving keys.

Step Solve—Runs the crowd simulation one frame at a time, starting at the current frame as specified by the time slider position. Press the spacebar to advance one frame.

To abort a solution in progress and save all keys generated up to that point, press the Esc key. When you do so, the software disregards any non-default settings for Save Every Nth Key. Alternatively, with complex simulations, you can save time by pressing Shift+Esc to abort a solution without saving keys.

Note: Step Solve always starts at the current frame; it disregards the Simulation Start setting.

Simulation Start—The first frame of the simulation. Default=0.

To make the solution repeatable, set this and keep it the same.

Start Solve—The frame at which you begin solving. Default=0.

This value must be greater than or equal to that of Simulation Start. If greater, the solve will begin in the middle of the simulation.

Start Solve should equal Simulation Start the first time you solve, so that when you set Start Solve to start in the middle of the simulation, the simulation up to that point will be correct.

Note: If you set Start Solve to a frame number lower than the first frame of the active time segment, **character studio** changes the first frame of the time segment to the Start Time value.

End Solve—Specifies the last frame considered for the solution. Default=100.

Note: If you set End Solve to a frame number higher than the last frame of the active time segment, **character studio** changes the last frame of the time segment to the End Solve value.

Delete Keys before Solve—Deletes the keys of active delegates in the range over which the solution takes place. Default=off.

This option leaves the first two keys so that the delegate doesn't end up with no keys and then pop to its current position. This is a useful feature for biped crowds; it lets you watch each biped compute, one after another, without the ones not yet computed still performing their old animation.

Save every Nth Key—Lets you specify the number of position and rotation keys saved after the solution.

Positions/Rotations—The frequency with which keys are saved for delegate positions and rotations. If 0, no keys are saved. If 1, a key is saved every frame. If 2, a key is saved every other frame, and so on. Defaults=1.

Display During Solve group

Update display—When on, motion produced during solution of a crowd simulation appears in the viewports. Default=on.

Frequency—How often the display is updated during the solution. If 1, the update occurs every frame. If 2, the update occurs every other frame, and so on. Default=1.

Vector Scale—Globally scales all force and velocity vectors that are displayed during the simulation. Default=10.0.

Scaling vectors up helps to see them better when they are very small. It does not effect the simulation.

MAXScript group

This feature lets you execute a MAXScript script at each frame. Its primary purpose is for working with bipeds; specifically, to take advantage of available MAXScript calls to Biped that let you specify which clip the biped will be likely to choose next when using the shared motion flow feature. With this feature, you can write a script that dynamically selects the biped's next clip during the Crowd simulation, based on which clip is currently being used, the frame number, the proximity of other bipeds, or anything else that you can find out in a script. Of course, this scripting feature can be used for other purposes as well.

Use MAXScript—When on, a user-specified script is executed at each frame during the solution. Default=off.

Function Name—The name of the function to be executed. This name must also specified in the script.

Edit MAXScript—Click this button to open a MAXScript window for displaying and modifying the script.

Bipeds group

When solving simulations that use bipeds linked to delegates, it is strongly recommended that you use all three options in this group.

Biped/Delegates Only—When on, only biped/delegates are included in the computation. Also, the options to use priorities and backtracking become available. These options are available only for biped-only computations. Default=off.

Use Priorities—When on, biped/delegates are computed one delegate at a time, in order of their Priority values, from lowest to highest. Also, backtracking becomes available, and Step Solve becomes unavailable. Default=off.

Backtracking—Turns on backtracking functionality when solving a crowd simulation that uses bipeds. Default=off.

When Backtracking is on during the solution, in the case of an impending collision between bipeds, the Crowd system will back up the simulation to the beginning of the current clip, and then

try a different traversal of the lower-priority delegate/biped's motion flow graph. If necessary, the system will back up two or more clips.

Priority Rollout

Create panel > Helpers > Object Type rollout > Crowd > Priority rollout

Select a Crowd object. > Modify panel > Priority rollout

The Crowd system uses the Priority rollout settings when solving a simulation involving bipeds associated with delegates. The Priority parameter is a positive integer assigned by the user to a delegate. When priorities are used, the Crowd simulation computes one biped at a time, based on its priority setting *from lowest to highest*; that is, a lower Priority setting means a higher priority. If the priorities of two biped/delegates are the same, the computation order of those two biped/delegates is randomly determined.

This topic describes the six different ways of setting a delegate's priority, and how priorities might be put to use in different situations.

Using Priorities

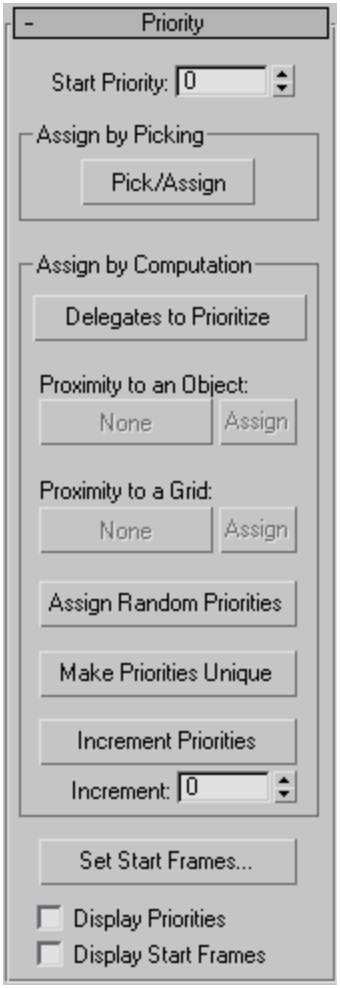
If you have a large crowd all going in one direction, you would typically want the delegates in front to solve first. In that case, using Proximity To An Object or Proximity To A Grid would be useful in setting priorities.

In a case where you start with a circle of bipeds, and you want them all to wander and mingle, you might not care about the bipeds' priorities. You could let them all have the same priority and let the system decide which goes first. However, it might be better to assign random priorities or make priorities unique, so that you are guaranteed the same order each time, and you can read the priority numbers to know what will happen next. This also lets you change the order if you need to.

Suppose you have two groups of bipeds, all of which are assigned random priorities. If you wanted to keep the priority relationships within each group, but make one group start after or before the other, you could use Increment Priorities to increment or decrement all the priorities in one group.

If none of the algorithms applies to your situation, you need some way to set the priorities by hand. It's useful to be able to set them visually. That's what the Assign By Picking method is for.

Interface



Start Priority—Sets the initial priority value. Applies to the first four methods of setting priorities: Assign By Picking, Proximity To An Object, Proximity To A Grid, and Assign Random Priorities. Default=0.

Note: Priority is assigned in increasing order. Thus, a delegate with Priority value 0 goes before a delegate with Priority 1, 1 goes before 2, and so on.

Assign by Picking group

Pick/Assign—Lets you assign successively higher Priority values to any number of delegates by selecting each in turn in the viewport. The first delegate you select is assigned the Start Priority value. The Priority value assigned to each succeeding delegate you select is incremented by one.

To stop assigning priorities, right-click in a viewport or click the Pick/Assign button again.

Delegate priorities appear in viewports as black numerals attached to each delegate; they're usually most easily seen in Wireframe views.

Note: You can undo and/or redo assignments during the process.

Tip: It's possible to assign two or more delegates the same priority value using this method. In such a case, for more predictable behavior, use Make Priorities Unique so that delegates don't share priorities.

Assign by Computation group

This group provides five different methods for assigning priorities to delegates, plus a button for selecting delegates to be affected by these methods.

Delegates to Prioritize—Lets you use the Select dialog to specify delegates to be affected by subsequent use of other controls within this group. Select the delegates with the Select dialog, and then click Select to exit the dialog. This selection applies only to Proximity assignments (that is, Proximity To An Object and Proximity To A Grid).

Proximity to an Object—Lets you assign priorities based on delegates' distance from a specific object. To specify the object, click the None button, and then select the object on which priorities are to be based. Lastly, click the Assign button to compute and assign priorities. The delegate closest to the object is assigned the Start Priority value, and each

successively farther delegate is assigned the next highest priority.

For any delegates that are equidistant from the object, the software assigns priorities randomly.

Proximity to a Grid—Lets you assign priorities based on delegates' distance from an infinite plane defined by a specific grid object. To specify the grid object, click the None button, and then select the grid object on which priorities are to be based. Lastly, click the Assign button to compute and assign priorities. The delegate closest to the grid object is assigned the Start Priority value, and each successively farther delegate is assigned the next highest priority.

For any delegates that are equidistant from the plane, the software assigns priorities randomly.

Assign Random Priorities—Assigns random priorities to the selected delegates. The range of priority values assigned lies between the Start Priority value and that value plus the number of selected delegates.

Make Priorities Unique—Ensures that all delegates have unique priority values. If two delegates share the same priority, one of them will be given a new priority value that differs from the rest.

Increment Priorities—Increments the priorities of all selected delegates by the Increment value.

Increment—Sets the value by which the Increment Priorities button adjusts delegate priorities. Use a negative Increment value to decrement priorities. Default=0.

Set Start Frames—Opens the *Set Start Frames dialog (page 2–1237)*, for setting start frames based on assigned priorities.

Display Priorities—Enables the display of assigned priority values as black numerals attached to the delegates. Default=off.

Display Start Frames—Enables the display of assigned start frame values as black numerals attached to the delegates. Default=off.

When both Display Priorities and Display Start Frames are on, the two values are displayed side by side, separated by a slash symbol (/). The priority appears to the left of the slash and the start frame appears to the right.

Set Start Frames Dialog

Create panel > Helpers > Object Type rollout > Crowd > Priority rollout > Set Start Frames button

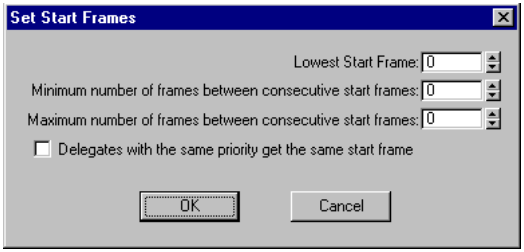
Select a Crowd object. > Modify panel > Priority rollout > Set Start Frames button

When you include bipeds in a crowd simulation using a shared motion flow, you usually don't want them all walking in lockstep formation. You can avoid this by setting different start frames to vary the animation frame at which each biped starts moving using its initial motion clip. In most cases, you would set start frames in the same order as priority, so you don't get bipeds with earlier start frames stuck behind bipeds with later start frames. The Set Start Frames dialog lets you automatically assign start frames to delegates in the same order as priority.

This dialog assigns start frames to the delegates selected via the *Delegates To Prioritize* control. In general, you should first use a non-random method to assign priorities in the *Priority rollout (page 2–1235)*, such as the Pick/Assign button or one of the Proximity options, and then assign start frames with this dialog. You can use the dialog to set a start frame for the highest-priority delegate (that is, the delegate with the lowest Priority setting), and then have start frames incremented randomly between two limits for the remaining delegates, based on priority order.

For instance, if you set Lowest Start Frame to 0, Minimum Number ... to 1, and Maximum Number to 3, you could get a start frame sequence like this: 0, 2, 5, 6, 8, and so on. If you want more regularly staggered start frames, set Minimum and Maximum to the same value. For example, with Lowest Frame=0 and Minimum/Maximum both set to 2, you'd get: 0, 2, 4, 6, 8, etc.

Interface



Lowest Start Frame—Sets the start frame assigned to the delegate with the lowest Priority setting.

Minimum number of frames between consecutive start frames—The smallest value the software will use to increment assigned start frames.

Maximum number of frames between consecutive start frames—The largest value the software will use to increment assigned start frames.

Delegates with the same priority get the same start frame—The software assigns the same start frame to any delegates with identical Priority settings. When off, the software randomizes the order of same-priority delegates' start frames. Default=on.

OK—Assigns start frame values to selected delegates based on the dialog settings and priority order and closes the dialog.

Cancel—Closes the dialog without changing start frame values.

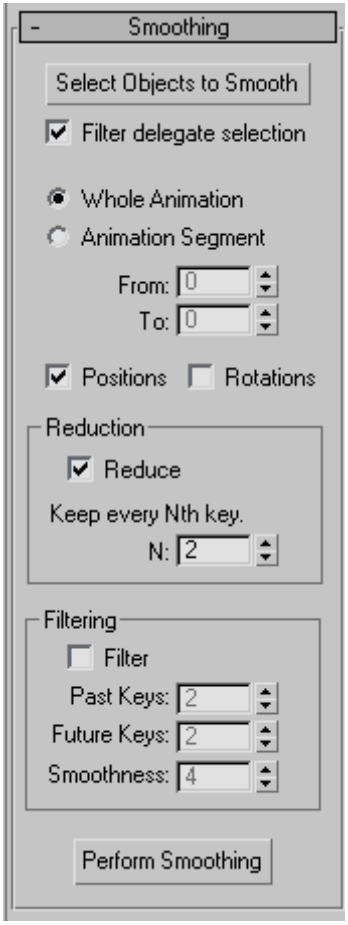
Smoothing Rollout

Create panel > Helpers > Object Type rollout > Crowd > Smoothing rollout

Select a Crowd object. > Modify panel > Smoothing rollout

Smoothing works on existing animation keys (that is, a solved simulation) to create more natural-looking animation. Use these controls if solving a crowd simulation results in abrupt position or rotation changes of animated objects. By default, smoothing works by reducing keys. As an option, it can also filter (average) existing animation keys to make changes more gradual, resulting in more natural motion.

Interface



Select Objects to Smooth—Opens the Select dialog, which lets you specify which objects' positions and/or rotations to smooth.

Filter delegate selection—When on, the Select dialog opened by the Select Objects To Smooth button shows only delegates. When off, it shows all scene objects. Default=on.

Whole Animation—Smoothes all animation frames. This is the default option.

Animation Segment—Smoothes only the frame ranges specified in the From and To fields.

From—When Animation Segment is chosen, specifies the first animation frame for smoothing.

To—When Animation Segment is chosen, specifies the last animation frame for smoothing.

Positions—When on, selected objects' animation paths generated via the simulation are smoothed after the simulation has finished. Default=on.

Rotations—When on, selected objects' rotations generated via the simulation are smoothed after the simulation has finished. Default=on.

Reduction group

Reduce—Reduces the number of keys by keeping only every Nth key.

Keep every Nth key—N: Limits the amount of smoothing by keeping every other key (N=2), or every third key (N=3), and so on. Default=2.

Filtering group

Smoothing works by averaging the delegate's current position and/or orientation with those several keyframes ahead and behind. All keyframes used in the calculation can be affected. These settings let you control the number of keyframes used and the extent of smoothing performed.

Filter—When on, smoothing is performed using the remaining settings in this group.

Past Keys—The number of keys prior to the current frame used for averaging position and/or rotation. Default=2.

Future Keys—The number of keys after the current frame used for averaging position and/or rotation. Default=2.

Smoothness—Determines the degree to which smoothing is performed. The higher the setting, the closer all keys involved in the calculation are moved toward the average value. Default=4.

The highest available Smoothness value is always 6. The lowest available value depends on the Past Keys and Future Keys settings.

Perform Smoothing—Click this button to carry out the smoothing operation.

Note: If neither the Reduce nor the Filter check box is turned on, no smoothing is performed.

Collisions Rollout

Create panel > Helpers > Object Type rollout > Crowd > Collisions rollout

Select a Crowd object. > Modify panel > Collisions rollout

During a crowd simulation, you can use this rollout to keep track of collisions defined by Avoid behaviors. A delegate whose hard radius as defined by the Avoid behavior intersects with the hard radius of anything it is avoiding is marked as having collided at that frame. If too many collisions occur, the simulation might not provide satisfactory results; in such cases, you might want to alter the simulation setup.

Interface



Hilite Colliding Delegates—When on, delegates that collide are highlighted in the collision color.

Only during collisions—Colliding delegates are highlighted only in frames in which they actually collide.

Always—Colliding delegates are highlighted in frames in which they collide and all subsequent frames.

Collision Color—The color swatch indicates the color used to highlight colliding delegates. To change the color, click the swatch and use the Color Selector dialog to set a new color.

Clear Collisions—Clears collision information from all delegates.

If you re-solve part of the simulation, new collision information for the recomputed frames will be computed. However, if you move a delegate manually, its collision information will remain the same, and may be incorrect. In such cases, use Clear Collisions to correct the collision information.

Geometry Rollout

Create panel > Helpers > Object Type rollout > Crowd > Geometry rollout

Select a Crowd object. > Modify panel > Geometry rollout

Use this parameter to modify the crowd object's size.

Interface

Icon Size—Determines the size of the Crowd helper object's icon.

This setting is primarily for visibility; it has no effect on the crowd simulation.

Global Clip Controllers Rollout

Create panel > Helpers > Object Type rollout > Crowd > Global Clip Controllers rollout

Select a Crowd object. > Modify panel > Global Clip Controllers rollout

Use global clip controllers when assigning non-biped animated objects (such as a bird flapping its wings) to delegates in a crowd simulation. Applications include synthesis of animation activity based on a variety of criteria, such as an object's speed, acceleration, and pitch. This group of controls replicates the Track View controls for the same functionality. For an in-depth discussion of global clip controllers and related topics, see:

Non-Biped Crowds (page 2–1179)

Synthesis Dialog (page 2–1246)

ClipState Dialog (page 2–1253)

For a procedure covering clip controller usage, see:

To use Motion Synthesis with non-bipedal creatures: (page 2–1180)

Interface



(List)—Lists objects designated as Global Objects, whose controllers can be used as animation clips to control other objects (typically clones).

To designate an object as a Global Object, click the New button, and then select the object in the Select dialog.

New—To designate a Global Object and add it to the list, click this button, and then select the object in the Select dialog.

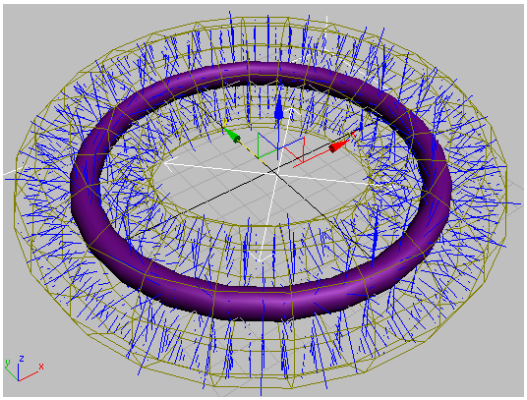
Edit—To modify a Global Object's properties, click its name in the list, and then click this button. This opens the *Synthesis dialog (page 2–1246)*.

Load—Loads a previously saved Global Motion Clip (*.ant*) file from disk.

Save—Stores the current Global Motion Clip settings on disk in the *.ant* file format.

Vector Field Space Warp

Create panel > Space Warps > Particles & Dynamics > Object Type rollout > Vector Field



A vector field is a special type of space warp that crowd members use to move around irregular objects such as curved, concave surfaces. The Vector Field gizmo, a box-shaped lattice, is placed and sized so that it surrounds the object to be avoided. The vectors are generated from the lattice intersections. These vectors are, by default,

perpendicular to the surface of the object to which the field is applied; if necessary, you can smooth them out with a blending function. The crowd members move around the object by traveling perpendicular to the vectors.

You can use the vector field as a *Space Warp behavior (page 2–1221)*, as the source object for an *Avoid behavior (page 2–1211)*, or both. When used together, delegates slow down when they approach a complex object, and then go around it. This guarantees that delegates will not pass through the obstacle's surface.

The vector field includes settings for strength, falloff, and a push/pull effect, as well as options to display the lattice, the effective range, and the vectors.

The Vector Field space warp provides these rollouts:

Create Method Rollout (page 2–1242)

Lattice Parameters Rollout (page 2–1242)

Obstacle Parameters Rollout (page 2–1242)

Create Method Rollout

Create panel > Space Warps > Particles & Dynamics > Object Type rollout > Vector Field > Create Method rollout

The Create Method rollout for the Vector Field space warp lets you specify whether to create the vector field using the cube or box method.

Interface

Cube—Forces length, width, and height to be equal.

Creating a cube-shaped space warp is a one-step operation. Starting at the center of the cube, drag in a viewport to set all three dimensions simultaneously. You can change the individual dimensions in the Lattice Parameters rollout.

Box—Creates a standard box-shaped space warp from one corner to the diagonally opposite corner, with different settings for length, width, and height.

Lattice Parameters Rollout

Create panel > Space Warps > Particles & Dynamics > Object Type rollout > Vector Field > Lattice Parameters rollout

Select a Vector Field space warp. > Modify panel > Lattice Parameters rollout

Use these parameters to specify the Vector Field lattice size and number of segments.

Interface

Length/Width/Height—Specify the dimensions of the lattice. The lattice should be larger than the Vector Field object.

Length Segs/Width Segs/Height Segs—Specify the resolution of the Vector Field lattice. The greater the resolution, the higher the accuracy of the simulation.

Obstacle Parameters Rollout

Create panel > Space Warps > Particles & Dynamics > Object Type rollout > Vector Field > Obstacle Parameters rollout

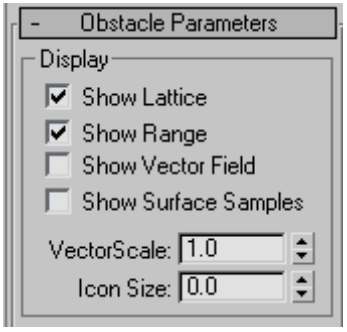
Select a Vector Field space warp. > Modify panel > Obstacle Parameters rollout

A *vector field (page 2–1241)*, generated around an *obstacle object*, allows *crowd (page 2–1187)* members to avoid that object in a scene. The field consists of a three-dimensional array of vectors which guide delegates or other objects around the obstacle. The settings in this rollout help determine how the vectors are generated and displayed, and how they affect other objects.

Note: Objects are subject to a vector field's forces only if they are bound to the field with a Crowd object. For general usage guidelines, see *To use a Vector Field space warp (page 2–1166)*.

Interface

Display group



The check boxes in this group let you enable and disable display of four different elements of the Vector Field space warp.

Note: You can select a Vector Field space warp in the viewport by clicking any of its visible elements *except* the range representation. If you've turned off display of all four elements, you can still select the space warp by clicking its gizmo, which, when viewed from the "top" (the orthogonal viewport in which the warp was created), resembles a pair of crossed double-headed arrows in the shape of an X.

Show Lattice—Displays the vector field lattice, a yellow wireframe box. Default=on.

The vectors are generated at lattice intersections inside the vector field range.

Show Range—Displays the volume about the obstacle object within which vectors are generated, as an olive-colored wireframe. Default=on.

The range starts out the same shape and size as the obstacle object, and is typically enlarged with the Compute Vectors group > Range setting.

Show Vector Field—Displays vectors, which appear as blue lines emanating outward from lattice intersections within the range volume. Default=off.

Show Surface Samples—Displays short green lines emanating from sample points on the surface of the obstacle object. Default=off.

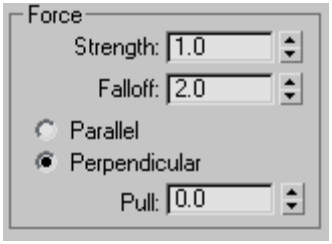
These appear only after you've *computed* the vector field. See *Sample Resolution* for more information.

VectorScale—Scales the vectors so they're more visible or less obtrusive. Default=1.0.

Note: This setting does not affect the strength of the vectors; only their visibility.

Icon Size—Adjusts the size of the Vector Field space warp icon, a pair of crossed double-headed arrows. Increase the size for easier viewport selection. Default=size originally drawn in viewport.

Force group



These parameters determine how the vector field affects objects within its volume.

Changing any of the Force group settings does not require that you recalculate the vector field.

Note: Using a vector field does not guarantee that delegates or particles remain at a particular distance from the obstacle. In some cases you might have to animate the Strength, Falloff, and/or Pull settings to keep delegates within the vector field.

Strength—Sets the degree of effect the vectors have on the movement of an object entering the

vector field. If *Show Vector Field* is on as you adjust Strength, you can see the vector lines change size in the viewports in real time. Default=1.0.

Note: Sometimes, after changing strength, vectors will be too large or too small. In such cases, adjust the *VectorScale* parameter so that they display properly.

Falloff—Determines the rate at which the strength of the vectors falls off with distance from the surface of the object. Default=2.0.

A value of 0 will make all the vectors the same size. A value greater than 0 will make them get smaller as they get farther away from the object surface. A value less then 0 will make them get larger as they get farther away.

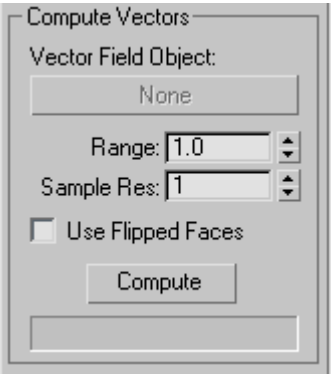
Parallel/Perpendicular—Sets whether the force generated by the vectors works parallel or perpendicular to the vector field. Default=Perpendicular.

Because the vectors are perpendicular to the object surface, and you typically would want delegates to travel parallel to the surface, you would normally use a perpendicular force.

Pull—Adjusts objects' position relative to the field. Available only when Perpendicular is chosen. Default=0.0. Range=-1.0 to 1.0.

Objects moving perpendicular to a vector field sometimes tend to drift away from it, due to lack of subsampling. The Pull parameter helps to pull objects back. Pull values greater than 0 create a pulling force towards the source of the vector field vector. Values less than 0 pull the force towards the direction in which the vector field's vector is pointing. A value of 0.0 produces a force perfectly perpendicular to the vector field's vector.

Compute Vectors group



Vector Field Object—Lets you designate the obstacle object. Click this button, and then select the object around which the vector field is to be generated. Thereafter the object's name appears on the button (which is intially labeled "None").

Note: You can only use primitives and unmodified editable mesh objects as obstacles. Also, the object should be fully enclosed in the Vector Field lattice.

Range—Determines the volume within which vectors are generated. Default=1.0.

The range is represented in viewports as an olive-colored wireframe that starts out the same size and shape as the obstacle object. Increasing the Range setting moves the wireframe away from the obstacle object in the direction of its surface normals.

In crowd simulations, the Range outline is where the delegates start to "see" the obstacle object, and begin to turn to avoid it. If your crowd members are penetrating the obstacle, or even just coming too close to it before turning, increase the Range setting. Also try increasing the Vector Field lattice resolution and/or the Sample Res setting.

Sample Res(olution)—Acts as a multiplier of the effective sampling rate used on the obstacle object's surface to calculate vector directions in the field. Default=1.

The basic sampling rate is determined by the program from the size of the lattice and the size of each polygon.

Use Flipped Faces—Causes flipped normals to be used during the computation of the vector field. Default=off.

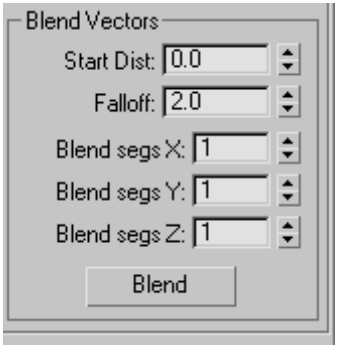
By default, vectors are generated in the same direction as the obstacle object's face normals, so that assuming these face normals point outward, objects move around its exterior in a crowd simulation. However, if you want objects to remain within an object's interior, turn on Use Flipped Faces.

Tip: If you run a crowd simulation inside an object that is also being viewed from the inside, such as a room, you'll probably want the object's faces to point inward. In that case, use Editable Mesh/Edit Mesh > Surface Properties to flip the normals, and don't turn on Use Flipped Faces.

Compute—Calculates the vector field.

Always recalculate the vector field after changing any parameters except those in the Force group.

Blend Vectors group



Use the Blend Vectors parameters to reduce abrupt changes in the angles of neighboring vectors. For example, if you have a wavy surface, you might get wavy vectors very far out from the surface, which could adversely affect the simulation. Use Blend Vectors to correct this condition.

Start Dist—The distance from the object at which you want to start blending the vectors. Default=0.0.

Falloff—The falloff of the blend of the surrounding vectors. Default=2.0.

Blend segs X—The number of adjacent lattice points to blend on the X-axis. Default=1.

Blend segs Y—The number of adjacent lattice points to blend on the Y-axis. Default=1.

Blend segs Z—The number of adjacent lattice points to blend on the Z-axis. Default=1.

Blend—Click this button to implement the blending.

Motion Synthesis

The Crowd system in **character studio** can use two different types of *motion synthesis*, which lets the software adjust the simulation results dynamically to account for differing conditions. This section deals with the application of motion synthesis to non-bipedal crowds, such as a flock of birds.

You use two dialogs to set up non-bipedal motion synthesis:

Synthesis Dialog (page 2–1246)

ClipState Dialog (page 2–1253)

For information on using motion synthesis with bipeds, see *Biped Crowds (page 2–1172)*.

Synthesis Dialog

Select a Crowd object. > Modify panel > Global Clip Controllers rollout > New > Choose a GlobalClip object. > Select the object in the list. > Edit > Synthesis dialog

Open Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog

The Synthesis dialog is where you set up motion synthesis for non-bipedal crowd members. It uses three panels to split up the workflow. On the Motion Clips panel, you specify the global object from which the motion clips are to be derived; you also set up the motion clips here. Controls on the State panel let you set up states and link clips to states. The Synthesis panel controls let you blend clips and synthesize the motion for some or all crowd members.

The panels are:

Motion Clips Panel (page 2–1246)

State Panel (page 2–1248)

Synthesis Panel (page 2–1250)

See also

ClipState Dialog (page 2–1253)

Procedure

Usage of the Synthesis dialog is closely associated with usage of the ClipState dialog. For a procedure that combines both, see *To use Motion Synthesis with non-bipedal creatures: (page 2–1180)*.

Motion Clips Panel

Select a Crowd object. > Modify panel > Global Clip Controllers rollout > New > Choose a GlobalClip object. > Select the object in the list. > Edit > Synthesis dialog > Motion Clips panel

Open Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > Motion Clips panel

On the Motion Clips panel, you specify the global object from which the motion clips are to be derived. You also set up the motion clips here.

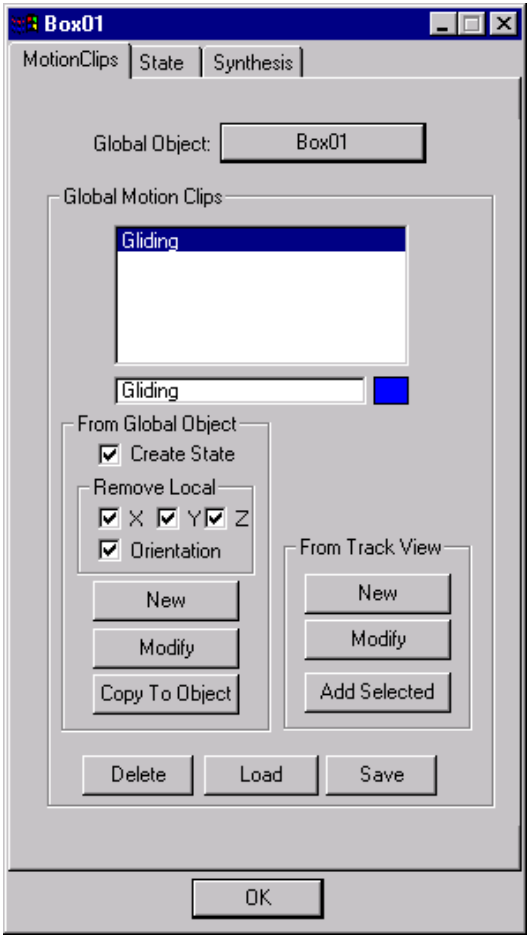
See also

Synthesis Dialog (page 2–1246)

State Panel (page 2–1248)

Synthesis Panel (page 2–1250)

Interface



Global Object—Click Global Object and pick the object that contains the animation (all the clips) in the Select Global Objects To Copy dialog.

Global Motion Clips group

List Window—The list of motion clips. The clips that you create appear in this list.

Edit Window—Rename or change the color for the selected motion clip.

From Global Object group

Create State—Create a new state with parameters specific to the motion clip, such as speed, heading, acceleration, and so on.

character studio evaluates the motion and orientation of the object and creates a new state with parameters set accordingly.

Remove Local X, Y, Z, Orientation—Turn on any or all of these options to strip out transformation and orientation data from the motion clip.

The idea is to animate the global object with full lateral motion. **character studio** then creates states based on the actual motion of the global object. After the states are created, the animation is stripped from the object. When delegates linked to clones of the global object come close to the actual motion recorded in the state, then the appropriate state is used to trigger the motion clip.

This technique is used to minimize sliding feet. If you animate a creature with many legs, you should animate lateral motion as well a leg motion. Then you create clips that record and then strip out the lateral motion that you created. When the delegate approaches the speed and direction you created originally on the global object, then the leg motion clip will be activated by the state; leg motion will be accurate relative to speed, preventing sliding feet.

New—Lets you set up a new motion clip. Displays the *MotionClip Parameters dialog (page 2–1252)* where you can specify the name and duration of the clip.

This selects all active controllers.

Modify—Lets you modify a motion clip's parameters. Highlight the clip to modify, and then click Modify. Displays the *MotionClip Parameters dialog (page 2–1252)*, where you can specify the name and duration of the clip.

Also, if you've changed the animation on the Global Object, you can re-apply it to a clip simply

by clicking Modify and then clicking OK to close the dialog. The change will affect all the objects being synthesized.

Copy to Object—The keys from the highlighted motion clip are copied back to the Global Object. Deletes any existing animation keys in that part of the Global Object's animation.

You can use this function in several ways. First, if you've changed the animation on the Global Object, you can restore it from a stored motion clip with Copy To Object. Also, if you've edited the clip in Track View (Global Tracks > Block Control > Global Motion Clip:: ...), you can apply it to the Global Object with Copy To Object.

From Track View group

New—Lets you create a motion clip from an animation track in Track View.

Specify a track in the *Track View Pick dialog (page 2–1252)*. The *MotionClip Parameters dialog (page 2–1252)* displays upon exiting the Track View Pick dialog.

Modify—Lets you modify a motion clip in the list by changing the Track View track from which it's derived, as well as the name and duration.

Specify a track in the *Track View Pick dialog (page 2–1252)*. The *MotionClip Parameters dialog (page 2–1252)* displays upon exiting the Track View Pick dialog.

Add Selected—Creates a motion clip from only the selected tracks in Track View.

Delete—Deletes the highlighted clip in the list.

Load—Loads a motion clip file (*.clp*). Displays the Load Motion Clip dialog. Use the Save command to create a *.clp* file.

Save—Saves a motion clip into a *.clp* file. Displays the Save Motion Clip dialog.

OK—Accepts changes and closes the dialog.

State Panel

Select a Crowd object. > Modify panel > Global Clip Controllers rollout > New > Choose a GlobalClip object. > Select the object in the list. > Edit > Synthesis dialog > State panel

Open Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > State panel

Controls on the State panel let you set up states for non-bipedal motion synthesis and link motion clips to states.

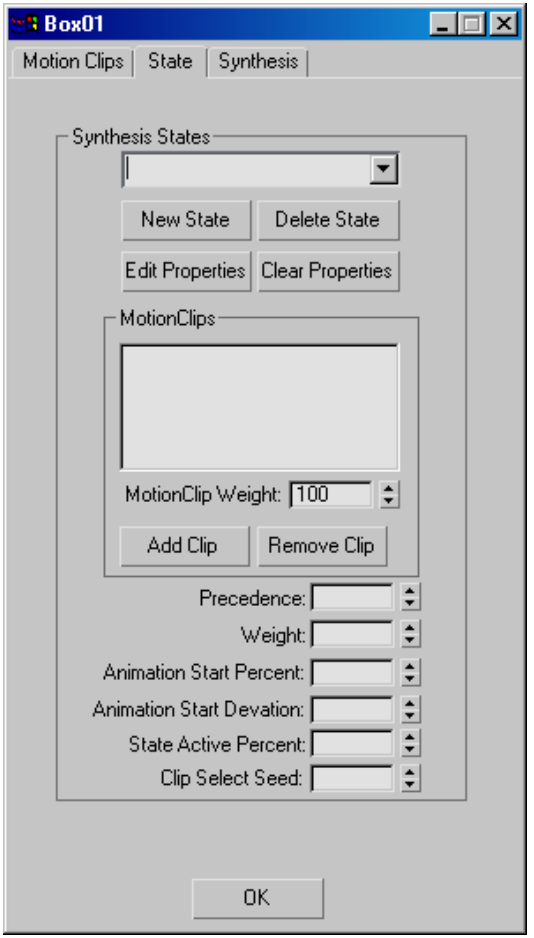
See also

Synthesis Dialog (page 2–1246)

Motion Clips Panel (page 2–1246)

Synthesis Panel (page 2–1250)

Interface



Synthesis States drop-down list—Displays the current state. Choose a state to modify from the list. You can change the state name by editing the text in the list window.

New State—Creates a new state and adds it to the list.

The default name is "ClipState," optionally followed by a number. You can change the state name by editing the text in the window.

Delete State—Deletes the current state. This is undoable.

Edit Properties—Lets you modify the current state. Displays the *ClipState dialog (page 2–1253)* for the current state.

Clear Properties—Returns the state to the default settings and removes clips from the MotionClips window.

MotionClips group

MotionClips window—Displays motion clips used by the current state.

Use Add Clip to associate clips with the state. If you assign more than one clip to a state, the software picks the clips randomly during synthesis based on their respective weights.

MotionClip Weight—Determines the chance that a clip will be chosen during synthesis. Range=0 to 1000.

Highlight the clip in the list, and then set its weight. A higher weight means that a clip is more likely to be chosen.

Add Clip—Displays the Select MotionClip dialog. Highlight a clip and click OK to add a clip to the current state.

Remove Clip—Removes the highlighted clip from the current state.

Precedence—Sets the precedence for the current state. Range=0 to 1000.

If multiple states qualify for activation based on their properties, **character studio** uses the clip associated with the one with the higher precedence first. If more than one state has the same precedence, then the one with the greater weight is played first.

Weight—Specifies a weight value for a state.

If two states have the same precedence, the state with a greater weight will be give higher priority considered during random selection.

Animation Start Percent—Specifies where in the clip's animation you want it to start playing when the state is active.

At the default value of 0, the animation will start once the state is active. If the value is 66, the animation will start playing two-thirds of the way in once the state becomes active. You can also randomize where the animation starts by specifying a Animation Start Deviation value other than 0.

Animation Start Deviation—You can randomize where the animation starts by specifying a Animation Start Deviation value other than 0. Range=0.0 to 1.0.

State Active Percent—Specifies the percentage of time the state needs to be valid over its interval in order for it to be selected. Range=0 to 100. Default=50.

Clip Select Seed—Changes how the random selections occur.

If the value stays the same, you are guaranteed to get the same random selections for that state.

OK—Accepts changes and closes the dialog.

Synthesis Panel

Select a Crowd object. > Modify panel > Global Clip Controllers rollout > New > Choose a GlobalClip object. > Select the object in the list. > Edit > Synthesis dialog > Synthesis panel

Open Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > Synthesis panel

Controls in the Synthesis panel are for adding objects to be synthesized, selecting blend transition points, and performing the synthesis. When you synthesize, different sequences of motion clips are applied to each specified object based on the behavioral motion of its delegate.

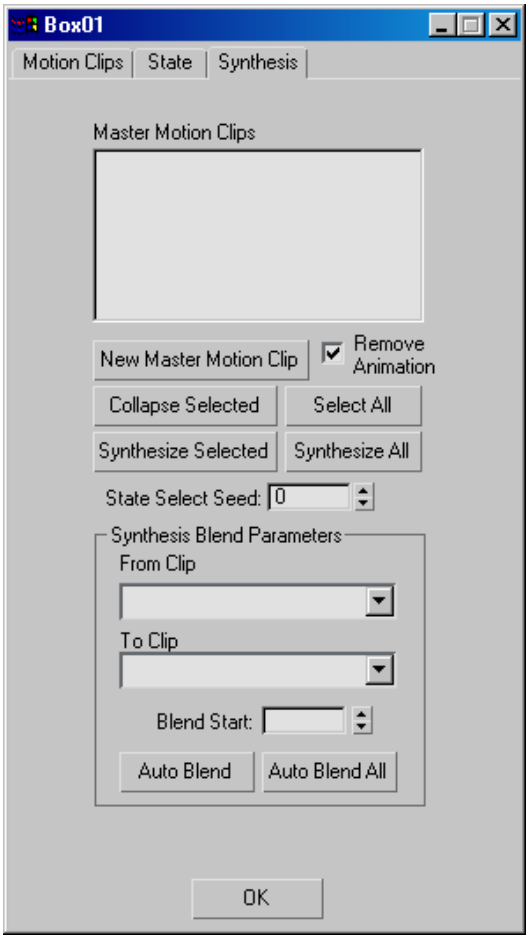
See also

Synthesis Dialog (page 2–1246)

Motion Clips Panel (page 2–1246)

State Panel (page 2–1248)

Interface



Master Motion Clips list—Displays the objects to which the synthesized motion will be applied.

You can highlight any of these by clicking them, and select multiple items with Ctrl+click and Shift+click. Highlight all objects with the Select All button.

New Master Motion Clip—Displays the Select Object To Copy dialog. Use this to specify the objects to which the synthesized motion will be applied.

These objects must all be structurally identical to the Global Object. In effect, they should be clones.

Remove Animation—Strips the animation from the clones.

After making clones of the original animated object, you can strip the animation from the clones. During synthesis, motion is applied based on which state is active.

Collapse Selected—Collapses motion clips to keys on the highlighted objects.

This deletes the Master Motion Clip for that object. You can then edit keys and make changes to the animation manually in the scene.

Synthesize Selected—Analyzes the motion of the delegates linked to the highlighted objects in the Master Motion Clips list, determines which state definitions fit that motion throughout the animation, and applies the corresponding motion clips.

Select All—Selects all of the objects in the Master Motion Clips list.

Synthesize All—Analyzes the motion of the delegates linked to all objects in the Master Motion Clips list, determines which state definitions fit that motion throughout the animation, and applies the corresponding motion clips.

A progress bar displays during synthesis. When synthesis is completed, the Synthesis dialog reappears, and you can view the calculated minimum, maximum, and average values of delegate motion for the different state properties in the ClipState dialog. You can use these values to fine-tune the state properties.

State Select Seed—Sets a seed value for random state selection.

During synthesis, it's possible that several states qualify for activation at the same time, in which case one state is chosen at random. The seed is used to modify the random value selected when determining which state to select.

Synthesis Blend Parameters group

During synthesis, clips aren't sequenced from end to end; instead, they partially overlap, with key blending or averaging occuring during the overlap intervals. This allows for smooth transitions between clips. During the blending or overlap period, weighting gradually shifts between the "from" and "to" clips, so the former's keys predominate at the beginning of the blend, and the latter's at the end.

You can specify explicitly how blending occurs between clips, or you can let the software calculate blending parameters automatically. To specify blending, use the drop-down lists to choose a pair of clips to blend between, and then set Blend Start to the frame in the "from" clip at which blending should begin. Alternatively, use Auto Blend or Auto Blend All to have **character studio** determine the best blend points.

If you use Auto Blend, you can then see the calculated blend start point for each pair of clips by choosing the clips from the drop-down lists.

From Clip list—Lets you select the starting clip to blend; the clip to blend *from*.

To Clip list—Lets you select the ending clip to blend; the clip to blend *to*.

Blend Start—Displays the frame in the From clip at which the transition is to begin, whether the

default, calculated, or set manually, and lets you change the start frame.

Auto Blend—Automatically sets the Blend Start frame for the current From and To clips.

Auto Blend All—Automatically sets the Blend Start frames for all possible pairs of clips.

OK—Accepts changes and closes the dialog.

MotionClip Parameters Dialog

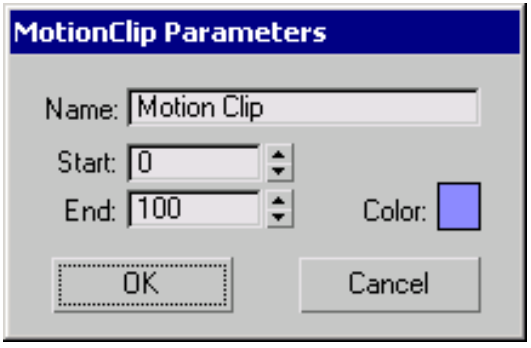
Open Track View. > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > Motion Clips panel > Click New (in the From Global Object group). > Motion Clip Parameters dialog

Select a Crowd helper. > Modify panel > Global Clip Controllers rollout > New > Choose a GlobalClip object. > Select the object in the list. > Edit > Synthesis dialog > Motion Clips panel > Click New (in the From Global Object group). > Motion Clip Parameters dialog

When setting up a crowd simulation with non-bipedal characters using motion synthesis, use this dialog, accessed from the *Synthesis Dialog (page 2–1246)*, to define motion clips from the Global Object's animation.

Enter a clip name, frame range, and color using controls in the MotionClip Parameters dialog. The software determines which clips to trigger by states that you define on the *State panel (page 2–1248)*.

Interface



Name—Enter a name for the new motion clip.

Start—Specifies the first frame of the animation clip from the Global Object.

End—Specifies the last frame of the animation clip from the Global Object.

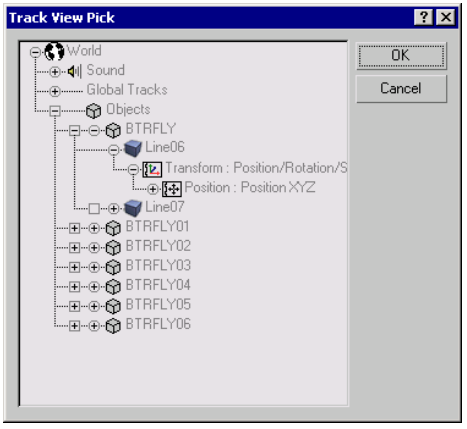
Color—Click the color swatch and choose a color in the color picker. This color is used in Track View when displaying the clip.

Track View Pick Dialog

Select a Crowd helper. > Modify panel > Global Clip Controllers rollout > New > Choose a GlobalClip object. > Select the object in the list. > Edit > Synthesis dialog > Motion Clips panel > Track View group > New > Track View Pick dialog

Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > Motion Clips panel > Track View group > New > Track View Pick dialog

Open the hierarchy and pick a track to use in the motion clip.



ClipState Dialog

Select a Crowd helper. > Modify panel > Global Clip Controllers rollout > New > Choose GlobalClip object. > Select object in list. > Edit > Synthesis dialog > State panel > New State > Edit Properties > Clip State dialog

Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > State panel > New State > Edit Properties > ClipState dialog

In **character studio**, in crowd animation with non-bipedal motion synthesis, a *state* is a particular property or set of properties of a delegate's animation; for example, the period during which it is pitched upwards and is decelerating. After determining the delegate's state, the motion synthesis engine chooses a motion clip for animating the object, or character, linked to the delegate; for example, a bird. When preparing a simulation that uses motion synthesis, you use the ClipState dialog to define states and associate states with motion clips.

You can define a state with any combination of these properties: speed, acceleration, pitch, pitch velocity, and heading velocity (plus a script). For each active property, you can specify a range or a unique value that triggers the clip for its respective state.

When you use a range, be sure to set the Min setting lower than the Max setting. For example, when using a negative range such as -180 to -10, enter the number with the larger absolute value (-180) as the Min setting.

You can see the delegates' actual ranges and average values for all properties after *synthesizing* the clips.

The dialog panels are:

Speed Panel (page 2–1254)

Acceleration Panel (page 2–1256)

Pitch Panel (page 2–1257)

Pitch Velocity Panel (page 2–1259)

Heading Velocity Panel (page 2–1261)

Script Panel (page 2–1262)

Note: The dialog contains more tabs than can fit across its top. To see additional tabs, click the left or right arrow button near the upper-right corner of the dialog.

You can enable and disable each state property individually. For example, you can tell the motion synthesis to consider only speed and pitch when analyzing delegate motion. For each of a state's active properties, you can designate a range of values between which a state can be activated.

Alternatively, you can specify a single, unique motion value to be used when a state is active, such as a speed of 50 units per frame. In addition, you specify the In and Out values of the parameter as it approaches and then passes through that unique value. These values are analogous to the tangents of a curve. For each, you can pick Anything, Decreasing, Increasing, or Constant. For example, an animation of an object that has just landed after flying might have a Unique speed value of 0.0 (the object isn't moving laterally), a decreasing In value (the object decelerated before landing) and a constant Out value (the object remains stationary). Taking off, on the other hand, would have an increasing Out value.

If you don't specify state parameters, the state is a default state. For example, if you want the motion synthesis engine to choose random motion clips for an object, and you don't care about what the speed or pitch is, you create default states, and the synthesis engine will randomly pick which clip is active.

Tip: When setting up states, it's useful to know the extents of state property values for delegate motion in the scene, such as maximum and minimum speed. To determine these, set up the clips and states you want to use, set the ranges to any values,

and then synthesize (see the procedure below). After you synthesize, the different property panels of the ClipState dialog will show the actual ranges and average values of delegate motion. These values are saved with the scene.

Note: The default name of this dialog is "ClipState" because that's the default name of the first state created on the Synthesis dialog > State panel. It's the dialog invoked by clicking the Edit Properties button. If you rename the state, the dialog assumes the changed name.

Automatic State Creation

character studio can create states automatically if you turn on Create State on the Synthesis dialog > Motion Clips panel. You can use this feature if you created a creature that contains lateral motion as well as looping motion.

If you turn on all the options in the Motion Clips panel > Remove Local group, then **character studio** creates a state that reflects the actual heading, speed, and acceleration of your creature. When a delegate approaches the heading and speed contained in this state, it triggers the appropriate motion clip. This method prevents sliding feet when animating multi-legged creatures.

See also

Synthesis Dialog (page 2–1246)

Speed Panel

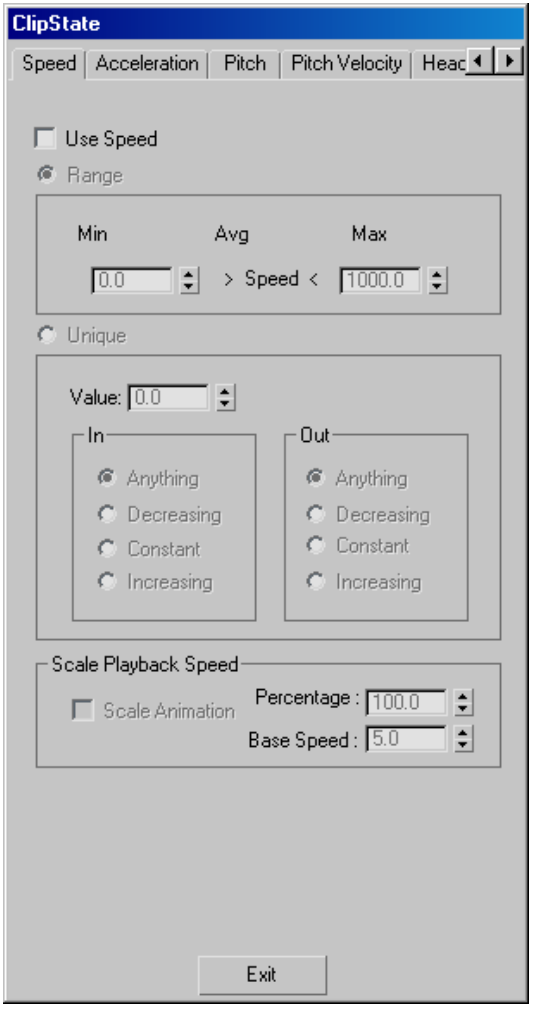
Select a Crowd helper. > Modify panel > Global Clip Controllers rollout > New > Choose GlobalClip object. > Select object in list. > Edit > Synthesis dialog > State panel > New State > Edit Properties > Clip State dialog > Speed Panel

Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > State panel > New State > Edit Properties > ClipState dialog > Speed Panel

To have delegate speed considered for motion synthesis, turn on Use Speed, choose Range or Unique, and then make the appropriate settings.

Speed is measured in units per animation frame, where the unit is the current 3ds Max system unit. By default, this is Generic Units.

Interface



Use Speed—Turn on to have the motion synthesis engine consider delegate velocity in determining whether to use the clip.

Range—Choose Range to have the motion synthesis engine activate the clip when the delegate's speed falls inside the specified range.

Range Display—After you *synthesize* the Master Motion Clips, displays delegates' minimum, average, and maximum speed.

Min—Set the minimum speed value for the range.

Max—Set the maximum speed value for the range.

Unique—Choose Unique to have the motion synthesis engine activate the clip when the delegate's speed matches a specific value, optionally with a rising, falling, or constant value before or after the specified value.

Value—Set a unique speed value.

In/Out—These radio buttons let you specify the behavior of the parameter before and after the unique value is met. For example, with a flying object, to trigger a landing animation clip you might set the unique value to 0, the In setting to Decreasing, and the Out setting to Constant. This means that the object has slowed down and then stopped, and is now stationary.

- **Anything—**Speed before or after the target value is not relevant.
- **Decreasing—**Speed decreases before or after it reaches the target value.
- **Constant—**Speed before or after the target value is constant.
- **Increasing—**Speed increases before or after it reaches the target value.

Scale Playback Speed group

These settings let you scale the rate at which the animation is played depending upon the speed of the delegate. By default, when Scale Animation is off, the object animation will always play at its normal rate. If you turn it on, you can then set a Base Speed, at which the animation should play at its normal rate, plus a percentage to specify how

much the animation rate is modified by the actual delegate speed.

Scale Animation—Scales the clip's animation based on speed.

For example, as a bird increases its velocity, its wings beat more rapidly. Scaling an animation scales the keys of the animation.

Percentage—Specify how much to alter the playback speed based upon the difference between the delegate's speed and the Base Speed setting.

The formula used is this:

Animation Speed change % = (current speed/Base Speed -1) x Percentage %

For example, if a delegate is moving 50 percent faster than the base speed, and the Scale Percentage value is 50, then the playback speed is scaled up by 25 percent.

Base Speed—Specifies the delegate speed at which the animation should be played back at its normal rate.

Acceleration Panel

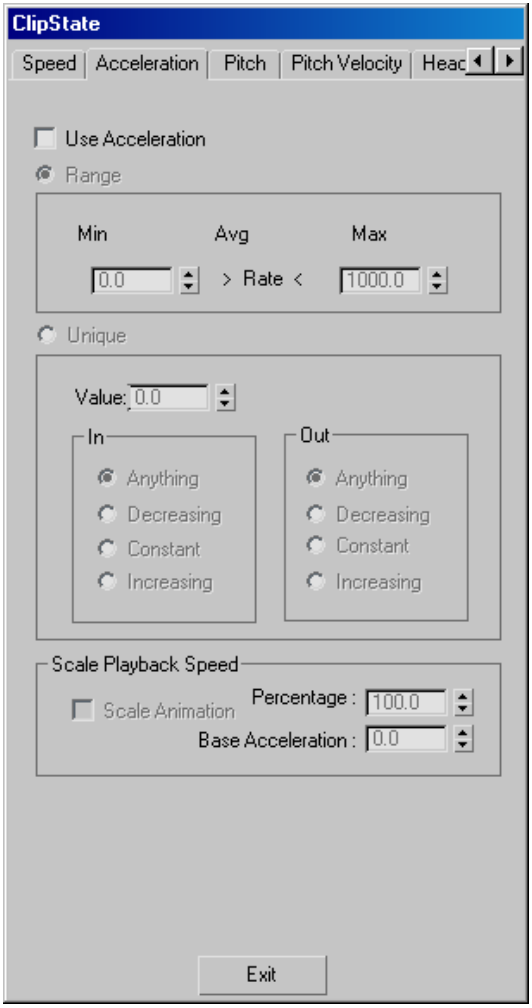
Select a Crowd helper. > Modify panel > Global Clip Controllers rollout > New > Choose GlobalClip object. > Select object in list. > Edit > Synthesis dialog > State panel > New State > Edit Properties > Clip State dialog > Acceleration Panel

Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > State panel > New State > Edit Properties > ClipState dialog > Acceleration Panel

To have delegate acceleration considered for motion synthesis, turn on Use Acceleration, choose Range or Unique, and then make the appropriate settings.

Interface

Acceleration is measured in units per frame per frame, where the unit is the current 3ds Max system unit. That is, the acceleration is determined by the rate at which the speed changes per frame.



Use Acceleration—Turn on to have the motion synthesis engine consider delegate acceleration in determining whether to activate the state.

Range—Choose Range to have the motion synthesis engine activate the clip when the delegate's acceleration falls inside the specified range.

Range Display—After you *synthesize* the Master Motion Clips, displays delegates' minimum, average, and maximum acceleration.

Min—Set a minimum acceleration value for the range.

Max—Set a maximum acceleration value for the range.

Unique—Choose Unique to have the motion synthesis engine activate the clip when the delegate's acceleration matches a specific value, optionally with a rising, falling, or constant value before or after the specified value.

Value—Set a unique acceleration value.

In/Out—

These radio buttons let you specify the behavior of the parameter before and after the unique value is met.

Anything—Acceleration before or after the target value is not relevant.

Decreasing—Acceleration decreases before or after it reaches the target value.

Constant—Acceleration before or after the target value is constant.

Increasing—Acceleration increases before or after it reaches the target value.

Scale Playback Speed group

These settings let you scale the rate at which the animation is played, depending upon the acceleration of the delegate. By default, when Scale Animation is off, the object animation will always play at its normal rate. If you turn it on, you can then set a Base Acceleration, at which the animation should play at its normal rate, plus a percentage to specify how much the animation rate is modified by the actual delegate acceleration.

Scale Animation—Scale the clip's animation based on acceleration.

For example, as a bird accelerates, its wings beat more rapidly. Scaling an animation scales the keys of the animation.

Percentage—Specify how much to alter the playback speed based upon the difference between the delegate's acceleration and the Base Acceleration setting.

The formula used is this:

Animation Speed change % = (current speed/Base Acceleration -1) x Percentage %

For example, if a delegate is accelerating 50 percent faster than the base acceleration, and the Scale Percentage value is 50, then the playback speed is scaled up by 25 percent.

Base Acceleration—Specifies the delegate acceleration at which the animation should be played back at its normal rate.

Pitch Panel

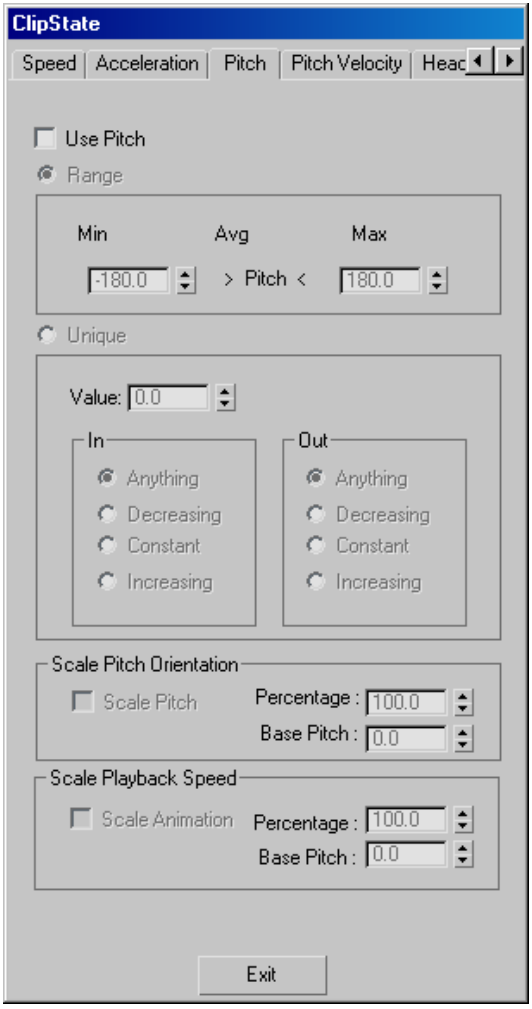
Select a Crowd helper. > Modify panel > Global Clip Controllers rollout > New > Choose GlobalClip object. > Select object in list. > Edit > Synthesis dialog > State panel > New State > Edit Properties > Clip State dialog > Pitch Panel

Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > State panel > New State > Edit Properties > ClipState dialog > Pitch Panel

To have delegate pitch considered for motion synthesis, turn on Use Pitch, choose Range or Unique, and then make the appropriate settings.

Pitch is determined by the angle in degrees of the delegate about the world X-axis. In world coordinates, the delegate's pitch is positive when it's aimed upward, negative when it's aimed downward, and 0 when it's aimed parallel to the home grid.

Interface



Use Pitch—Turn on to have the motion synthesis engine consider delegate pitch in determining whether to activate the state.

Range—Choose Range to have the motion synthesis engine activate the clip when the delegate's pitch falls inside the specified range.

Range Display—After you *synthesize* the Master Motion Clips, displays delegates' minimum, average, and maximum pitch.

Min—Set a minimum pitch for the range.

Max—Set a maximum pitch for the range.

Unique—Choose Unique to have the motion synthesis engine activate the clip when the delegate's pitch matches a specific value, optionally with a rising, falling, or constant value before or after the specified value.

Value—Set a unique pitch value

In/Out—These radio buttons let you specify the behavior of the parameter before and after the unique value is met.

- **Anything**—Pitch before or after the target value is not relevant.
- **Decreasing**—Pitch decreases before or after it reaches the target value.
- **Constant**—Pitch before or after the target value is constant.
- **Increasing**—Pitch increases before or after it reaches the target value.

Scale Pitch Orientation group

These settings let you scale the pitch at which the animation is played, depending upon the pitch of the delegate. By default, when Scale Pitch is off, the animated object will pitch normally. If you turn it on, you can then set a Base Pitch at which the animation should play, plus a percentage to specify how much the Base Pitch is modified by the actual delegate pitch. If the Percentage is 100, then only the Base Pitch will be used; anything else will scale between the base pitch and the delegate pitch. This is useful for when birds and other objects elevate, but don't pitch up.

Scale Animation—Scales the clip's pitch based on delegate pitch.

Percentage—Specify how much to alter the pitch based upon the difference between the delegate's pitch and the Base Pitch setting.

The formula used is this:

Pitch change % = (current pitch/Base Pitch -1) x Percentage %

Base Pitch—Specifies the delegate pitch at which the animation should be played back at its normal rate.

Scale Playback Speed group

These settings let you scale the rate at which the animation is played, depending upon the pitch of the delegate. By default, when Scale Animation is off, the object animation will always play at its normal rate. If you turn it on, you can then set a Base Pitch, at which the animation should play at its normal rate, plus a percentage to specify how much the animation rate is modified by the actual delegate pitch.

Scale Animation—Scales the clip's animation based on pitch.

Percentage—Specify how much to alter the playback speed based upon the difference between the delegate's pitch and the Base Pitch setting.

The formula used is this:

Animation Speed change % = (current speed/Base Pitch -1) x Percentage %

For example, if a delegate's pitch velocity is 50 percent above its base pitch rate, and the Scale Percentage value is 50, then the playback speed is scaled up by 25 percent.

Base Pitch—Specifies the delegate pitch at which the animation should be played back at its normal rate.

Pitch Velocity Panel

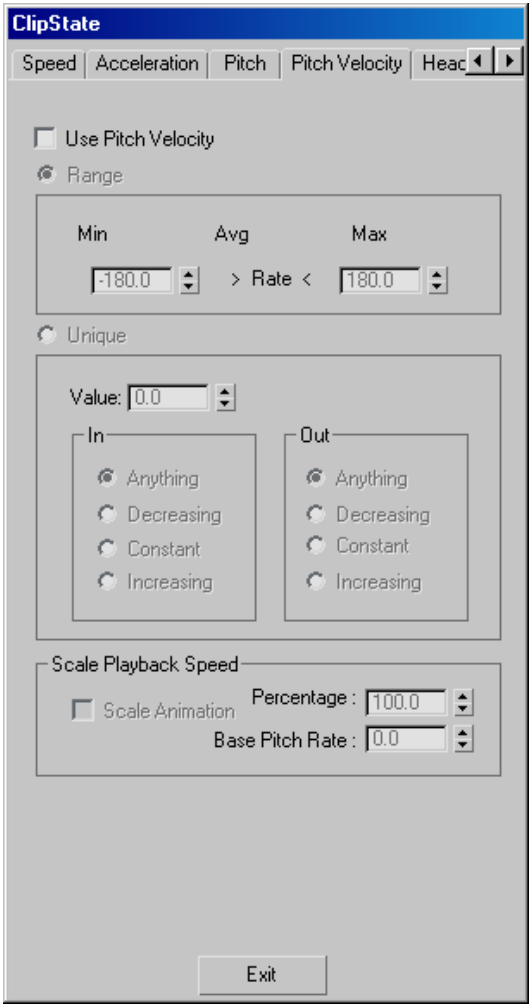
Select a Crowd helper. > Modify panel > Global Clip Controllers rollout > New > Choose GlobalClip object. > Select object in list. > Edit > Synthesis dialog > State panel > New State > Edit Properties > Clip State dialog > Pitch Velocity Panel

Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > State panel > New State > Edit Properties > ClipState dialog > Pitch Velocity Panel

To have delegate pitch velocity considered for motion synthesis, turn on Use Pitch Velocity, choose Range or Unique, and then make the appropriate settings.

Pitch velocity is determined by the rate of change in degrees per frame of the angle of the delegate about the world X-axis. In other words, pitch velocity measures how fast the delegate is changing its *pitch*.

Interface



Use Pitch Velocity—Turn on to have the motion synthesis engine consider pitch velocity in determining whether to activate the state.

Range—The motion synthesis engine activates the clip when the delegate's pitch velocity falls inside the specified range.

Range Display—After you *synthesize* the Master Motion Clips, displays delegates' minimum, average, and maximum pitch velocity.

Min—Set a minimum pitch velocity for the range.

Max—Set a maximum pitch velocity for the range.

Unique—Choose Unique to have the motion synthesis engine activate the clip when the delegate's pitch velocity matches a specific value, optionally with a rising, falling, or constant value before or after the specified value.

Value—Set a unique pitch velocity value.

In/Out—These radio buttons let you specify the behavior of the parameter before and after the unique value is met.

- **Anything**—Pitch velocity before or after the target value is not relevant.
- **Decreasing**—Pitch velocity decreases before or after it reaches the target value.
- **Constant**—Pitch velocity before or after the target value is constant.
- **Increasing**—Pitch velocity increases before or after it reaches the target value.

Scale Playback Speed group

These settings let you scale the rate at which the animation is played, depending upon the pitch velocity of the delegate. By default, when Scale Animation is off, the object animation will always play at its normal rate. If you turn it on, you can then set a Base Pitch Rate, at which the animation should play at its normal rate, plus a percentage to specify how much the animation rate is modified by the actual delegate pitch velocity.

Scale Animation—Scales the clip's animation based on pitch velocity.

Percentage—Specify how much to alter the playback speed based upon the difference between the delegate's pitch velocity and the Base Pitch Rate setting.

The formula used is this:

Animation Speed change % = (current speed/Base Pitch Rate -1) x Percentage %

For example, if a delegate's pitch velocity is 50 percent above its base heading rate, and the Scale Percentage value is 50, then the playback speed is scaled up by 25 percent.

Base Pitch Rate—Specifies the delegate pitch velocity at which the animation should be played back at its normal rate.

Heading Velocity Panel

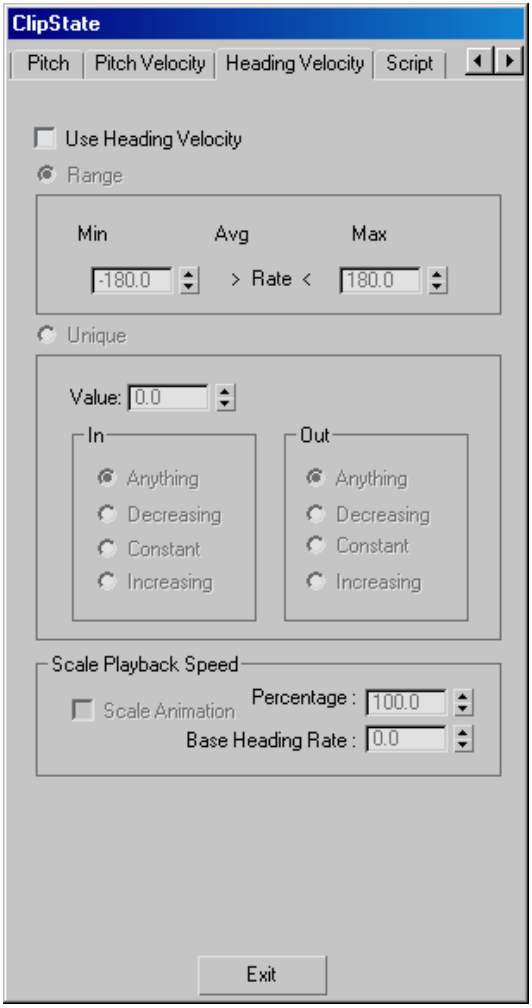
Select a Crowd helper. > Modify panel > Global Clip Controllers rollout > New > Choose GlobalClip object. > Select object in list. > Edit > Synthesis dialog > State panel > New State > Edit Properties > Clip State dialog > Heading Velocity Panel

Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > State panel > New State > Edit Properties > ClipState dialog > Heading Velocity Panel

To have delegate heading velocity considered for motion synthesis, turn on Use Heading Velocity, choose Range or Unique, and then make the appropriate settings.

Heading velocity is determined by the rate of change in degrees per frame of the angle of the delegate about the world Z-axis. In other words, heading velocity measures how fast the delegate is changing its heading.

Interface



Use Heading Velocity—Turn on to have the motion synthesis engine consider heading velocity in determining whether to activate the state.

Range—Choose Range to have the motion synthesis engine activate the clip when the delegate's heading velocity falls inside the specified range.

Range Display—After you *synthesize* the Master Motion Clips, displays delegates' minimum, average, and maximum heading velocity.

Min—Set a minimum heading rate value for the range.

Max—Set a maximum heading rate value for the range.

Unique—Choose Unique to have the motion synthesis engine activate the clip when the delegate's heading rate matches a specific value, optionally with a rising, falling, or constant value before or after the specified value.

Value—Set a unique heading rate value.

In/Out—These radio buttons let you specify the behavior of the parameter before and after the unique value is met.

- **Anything—**Heading velocity before or after the target value is not relevant.
- **Decreasing—**Heading velocity decreases before or after it reaches the target value.
- **Constant—**Heading velocity before or after the target value is constant.
- **Increasing—**Heading velocity increases before or after it reaches the target value.

Scale Playback Speed group

These settings let you scale the rate at which the animation is played, depending upon the heading velocity of the delegate. By default, when Scale Animation is off, the object animation will always play at its normal rate. If you turn it on, you can then set a Base Heading Rate, at which the animation should play at its normal rate, plus a percentage to specify how much the animation rate is modified by the actual delegate heading velocity.

Scale Animation—Scales the clip's animation based on heading velocity.

Percentage—Specify how much to alter the playback speed based upon the difference between the delegate's heading velocity and the Base Heading Rate setting.

The formula used is this:

Animation Speed change % = (current speed/Base Heading Rate -1) x Percentage %

For example, if a delegate's heading velocity is 50 percent above its base heading rate, and the Scale Percentage value is 50, then the playback speed is scaled up by 25 percent.

Base Heading Rate—Specifies the delegate heading velocity at which the animation should be played back at its normal rate.

Script Panel

Select a Crowd helper. > Modify panel > Global Clip Controllers rollout > New > Choose GlobalClip object. > Select object in list. > Edit > Synthesis dialog > State panel > New State > Edit Properties > Clip State dialog > Script Panel

Track View > Global Tracks > Block Control > GlobalClip Properties > Synthesis dialog > State panel > New State > Edit Properties > ClipState dialog > Script Panel

The script state option lets you create a MAXScript script that takes two parameters: node and time. The script typically tests one or more values, and then returns 1 if the condition (the result of the test) is true or 0 if it's false. This result determines whether or not the state is to be activated.

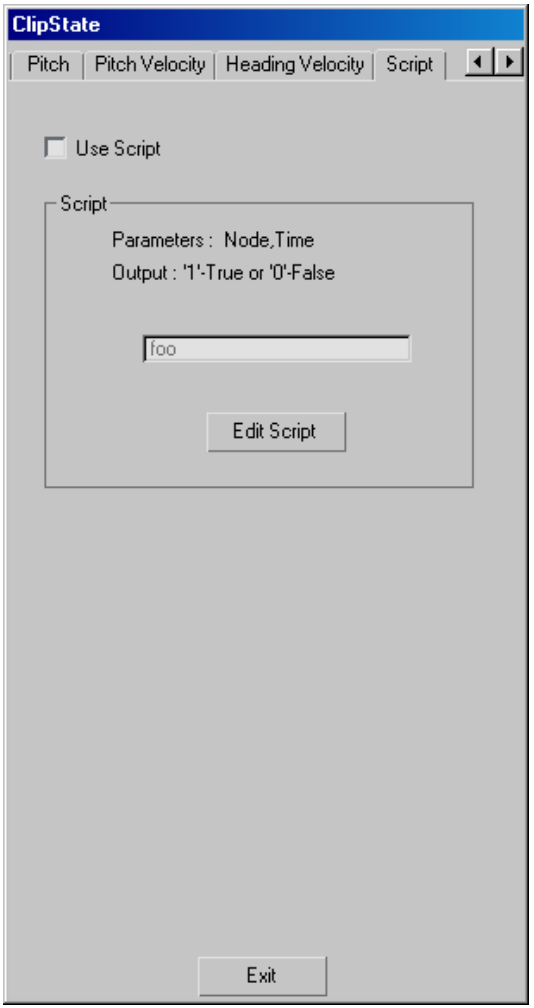
Scripts used by the clip controller are similar to those used by the *cognitive controller (page 2–1206)*, with the exception that a special time-related statement is required.

In the following sample script, "del" is the delegate's node, and "t" is the time. The name of the scripted function, "stoppedScript," would also need to be entered into the name field in the Script panel of the ClipState dialog. Unlike cognitive

controller scripts, the statement "at time t" needs to be invoked because the animation is not running when the synthesis takes place.

```
fn stoppedScript del t = (
at time t
if del.pos.z < 65 then 1 else 0
)
```

Interface



Use Script—Turn this on to use a MAXScript script to control a clip.

Script—Enter the name of the function defined by the script, also found at the start of the script.

Edit Script—Opens a MAXScript editor window that lets you edit the script.

File Formats and Index of Procedures

character studio File Formats

character studio uses several proprietary file formats for various functions.

BIP—Motion file. Includes biped motion and the size of the original biped so motion can be accurately loaded to other bipeds. Can be loaded directly onto a biped, or used to combine motions with Motion Flow or the Motion Mixer. See *Loading and Saving BIP Animation (page 2–920)*.

FIG—Figure file. Contains the sizes and poses of each biped body part, and overall biped size. Use this file to save and load a biped pose in Figure mode. See *Saving and Loading FIG Files (page 2–855)*.

CPY—Copy file. Contains postures, poses and tracks you have copied with the Copy/Paste rollout. See *Copying and Pasting Postures and Poses (page 2–910)* and *Copying and Pasting Tracks (page 2–926)*.

MFE—Motion Flow Graph file. Contains the motion flow graph, and a motion flow script if one was present for the biped upon saving. See *Saving, Loading, and Appending Motion Flow Graphs (page 2–1032)*.

MIX—Motion Mixer file. Contains clip data, tracks and trackgroups, transitions between clips, track/clip weights, and scale, trim and time

warp information for edited clips. See *Exporting Animation to the Biped (page 2–624)*.

PHY—Physique file. Contains information on envelope sizes and vertex weights set up with the Physique modifier. See *Saving and Loading Physique Data (page 2–1098)*.

STP—Step file. Contains footstep data. See *Loading and Saving STP Files (page 2–924)*.

BVH, CSM, MNM—These file formats contain motion-capture data. They are not specific to **character studio**. See *BVH File (page 3–920)*, *CSM File (page 3–925)*, or *MNM File (page 3–973)* in the glossary.

Procedures

These procedures appear elsewhere in this *User Reference*, in their associated topics. Here, they are organized by feature.

Biped and Physique

Creating a Biped

To delete a biped: (page 2–854)

To create a biped on a surface: (page 2–845)

To change the biped name upon creation: (page 2–847)

To change the biped name after creation: (page 2–847)

To create a Named Selection for the biped center of mass: (page 2–847)

Loading and Viewing Motion

To save animation you have created or imported to a BIP file: (page 2–921)

To load biped motion from an existing BIP file: (page 2–920)

To preview biped motion using biped playback stick figure: (page 2–929)

To preview biped motion using the full biped model: (page 2–929)

Animating the Biped with Freeform

To use pivots: (page 2–907)

To make a hand or foot follow an object: (page 2–908)

To anchor a hand or foot: (page 2–909)

Animating the Biped with Footsteps

To create footsteps automatically: (page 2–862)

To create multiple footsteps: (page 2–988)

To create footsteps beginning at the current frame: (page 2–864)

To make the biped walk up or down stairs: (page 2–992)

To make the biped walk in place: (page 2–992)

To make the biped walk backward: (page 2–992)

To prepare for manual footstep creation: (page 2–864)

To create footsteps manually, beginning at the current frame: (page 2–989)

To append footsteps onto the existing footsteps: (page 2–989)

To create footsteps using AutoGrid: (page 2–864)

To make the biped speed up as it walks: (page 2–992)

To activate footsteps: (page 2–865)

To deactivate footsteps: (page 2–866)

Editing Footsteps

To save footstep data: (page 2–925)

To move selected footsteps in time: (page 2–871)

To change the duration of a footstep: (page 2–871)

To scale keys in time: (page 2–871)

To prevent keys from changing when active footsteps are edited: (page 2–872)

To turn off Scale Stride mode: (page 2–922)

To display footstep tracks: (page 2–889)

To display footstep keys in Track View: (page 2–871)

To change the display of numbers of footstep keys: (page 2–871)

To convert an airborne period to freeform mode: (page 2–884)

To convert between footsteps and freeform animation: (page 2–885)

To create a purely freeform animation: (page 2–887)

To create a freefrom animation from a footstep animation: (page 2–887)

To save footstep data: (page 2–925)

To copy footsteps to the buffer: (page 2–873)

To edit the footstep buffer: (page 2–873)

To splice the footstep buffer: (page 2–873)

Editing Biped Keys

To use In Place mode to adjust keyframes: (page 2–930)

To set the balance factor: (page 2–878)

To set ballistic tension: (page 2–882)

To set the IK Blend value of a key: (page 2–901)

To locate vertical center of mass keys: (page 2–881)

To set a Body Vertical key: (page 2–881)

To change TCB for a biped arm: (page 2–894)

To select and rotate multiple links: (page 2–895)

To rotate all links in the spine, neck, or tail: (page 2–895)

To change Dynamics Blend for multiple Body Vertical keys: (page 2–882)

To bend a biped trajectory: (page 2–915)

To edit keys on the Center of Mass trajectory: (page 2–915)

To edit biped trajectory keys in a viewport: (page 2–931)

To use In Place mode to adjust keyframes: (page 2–930)

Motion-Capture Data

To import a motion-capture file: (page 2–1067)

To import a marker file: (page 2–1063)

To use Fit To Existing to import a motion-capture file: (page 2–1071)

To use Convert From Buffer: (page 2–1066)

To compare raw and filtered trajectories: (page 2–1066)

To use Show Buffer: (page 2–1066)

Motion Flow

To create clips in the Motion Flow Graph: (page 2–1027)

To create multiple clips in the Motion Flow Graph: (page 2–1028)

To create transitions using From > To or To < From method: (page 2–1029)

To create transitions using Create All Transitions: (page 2–1029)

To optimize transitions in the Motion Flow Graph: (page 2–1030)

To create a Motion Flow script: (page 2–1031)

To save Motion Flow Editor files: (page 2–1032)

To load Motion Flow Editor files: (page 2–1033)

To append Motion Flow Editor files: (page 2–1033)

To manually customize transitions between two clips: (page 2–1034)

To automatically customize transitions between two clips: (page 2–1035)

To create a random script for one biped: (page 2–1036)

To customize transition "weighting" and setting a start clip: (page 2–1037)

Example: To create a unified motion: (page 2–1038)

Example: To share a random motion flow among multiple bipeds: (page 2–1039)

To apply a shared motion flow to a different group of bipeds: (page 2–1041)

To load a Motion Flow scene from another system: (page 2–1042)

Posing the Biped

To link a mechanical character to the biped (without Physique): (page 2–855)

To fit biped legs to the skin: (page 2–850)

To fit the spine to the skin torso: (page 2–850)

To fit both arms using copy/paste: (page 2–851)

To pose both arms simultaneously: (page 2–851)

To create a symmetrical pose by copying one side of the biped to the other: (page 2–851)

To copy a posture: (page 2–912)

To paste the posture: (page 2–912)

To copy a track from one biped to another: (page 2–927)

To copy the entire pose of a biped: (page 2–912)

To attach a mesh to a bones hierarchy using Physique: (page 2–1106)

Applying Physique

To attach a mesh to a biped using Physique: (page 2–1106)

To adjust default envelope shape: (page 2–1087)

To adjust envelopes around the biped's pelvis: (page 2–1087)

To add a bone after you've already used Attach To Node: (page 2–1107)

To troubleshoot bulges and tendons: (page 2–1109)

To select and edit cross sections: (page 2–1088)

To copy an envelope and its settings to a mirrored link (for example, from one thigh to the other): (page 2–1088)

To copy all Bulge angles from one link to its opposite: (page 2–1143)

To use Select Nearest Bulge Angle: (page 2–1144)

To remove a link's influence on vertices: (page 2–1151)

To check vertex assignments: (page 2–1151)

To remove deformable vertices from a link's influence: (page 2–1151)

To override vertex assignments manually: (page 2–1151)

To make vertices rigid: (page 2–1152)

To have a spline influence a mesh: (page 2–1110)

To apply Physique to an FFD to animate the entire mesh: (page 2–1105)

To use an FFD to complement the effects of Physique on a portion of a character mesh: (page 2–1105)

To make compressible bones (bones with end effectors): (page 2–1081)

To add a bone after Physique is applied using Reinitialize: (page 2–1081)

To add a bone after Physique is applied using Add (Add Bone): (page 2–1082)

To optimize skin objects: (page 2–1100)

To reinitialize a scaled mesh: (page 2–1100)

Bulges

To create a new bulge angle using the Bulge Editor: (page 2–1115)

To create a new bulge angle on a selected link: (page 2–1142)

To add a cross section: (page 2–1116)

To change the shape of a cross section: (page 2–1117)

To make a cross section the active cross section: (page 2–1116)

To choose a specific bulge angle for editing: (page 2–1116)

To change a bulge angle value: (page 2–1116)

To select multiple cross sections: (page 2–1116)

To move cross sections along the link: (page 2–1117)

To copy and paste cross sections: (page 2–1117)

To change the Profile view orientation: (page 2–1117)

To delete a bulge angle: (page 2–1116)

To delete a cross section: (page 2–1116)

Links and Tendons

To adjust joint intersection parameters: (page 2–1140)

To create and attach a tendon: (page 2–1097)

To attach a tendon to another link: (page 2–1097)

To delete a tendon: (page 2–1097)

Saving and Reusing Physique Information

To save Physique data: (page 2–1107)

To load Physique data: (page 2–1107)

To merge a skinned biped: (page 2–922)

To clone a skinned biped: (page 2–924)

Figure Mode

To work in Figure mode: (page 2–849)

To adjust the biped center of mass with Rubber Band: (page 2–853)

To rubber band an arm or leg link: (page 2–853)

To Scale a Biped and a Physique Mesh: (page 2–984)

To save a biped's figure information to a file: (page 2–856)

To load a biped figure: (page 2–856)

Layers

To increment all keys using layers (global offset): (page 2–914)

To increment an interval of keys with an envelope (blended offset): (page 2–914)

To increment an interval of keys without an envelope (nonblended offset): (page 2–914)

To use layers to reposition a biped with freeform animation: (page 2–928)

Props

Example: To animate a biped swinging a prop with one hand: (page 2–899)

Example: To animate the biped switching the prop to the other hand: (page 2–899)

To collapse the transforms of a prop: (page 2–900)

Controllers and Biped

To add a controller to a biped object: (page 2–896)

To collapse a controller: (page 2–897)

To add controllers to multiple biped parts at the same time: (page 2–897)

To animate the weights of an added controller: (page 2–897)

Miscellaneous

To load a file created with different system units: (page 2–1099)

To hide the finger, toe, and head dummy objects: (page 2–923)

To change biped display: (page 2–854)

To select from the screen: (page 2–889)

To reposition a freeform animation (with no IK attachment) or a footstep animation: (page 2–927)

To reposition a biped with limbs attached to an Object Space object (IK attachment): (page 2–928)

To reposition a biped with limbs attached to world space (IK attachment): (page 2–928)

To reposition a biped animated with a motion flow script: (page 2–928)

To scale a biped that has a mesh attached to it by using Physique: (page 2–1100)

Facial Animation (page 2–1100)

To isolate lip vertices from influence by inappropriate links: (page 2–1103)

Crowd System

Crowd and Delegate Helpers

To create a Crowd helper object: (page 2–1158)

To create a Delegate helper object: (page 2–1158)

To clone and scatter delegates: (page 2–1158)

To link objects to delegates: (page 2–1169)

To associate bipeds with delegates: (page 2–1169)

To edit multiple delegates: (page 2–1197)

To group delegates into a team: (page 2–1201)

Behaviors

To create a new behavior assignment: (page 2–1201)

To use assign behaviors to delegates: (page 2–1162)

To modify an existing behavior assignment or assignments: (page 2–1202)

To use the Avoid behavior: (page 2–1165)

To use the Orientation behavior: (page 2–1167)

To use the Path Follow behavior: (page 2–1163)

To use the Repel behavior: (page 2–1166)

To use the Seek behavior: (page 2–1163)

To add a Vector Field space warp: (page 2–1166)

To use a Vector Field space warp with delegates: (page 2–1166)

To use the Speed Vary behavior: (page 2–1168)

To use the Surface Arrive behavior: (page 2–1164)

To use the Surface Follow behavior: (page 2–1164)

To use the Wall Repel behavior: (page 2–1165)

To use the Wall Seek behavior: (page 2–1163)

To use the Wander behavior: (page 2–1231)

Solving a Crowd Simulation

To solve a simulation: (page 2–1168)

To speed up the solution time: (page 2–1168)

To troubleshoot the simulation: (page 2–1168)

Cognitive Controllers

To set up and use a cognitive controller: (page 2–1170)

Testing a Particle System Parameter (page 2–1209)

Testing an Atmospheric Property (page 2–1209)

Testing Another Delegate's Behavior (page 2–1209)

Motion Synthesis

To use Motion Synthesis with non-bipedal creatures: (page 2–1180)

Procedures: Using Bipeds in a Crowd Simulation (page 2–1175)

Motion Mixer

General

To access the Motion Mixer from the Motion panel: (page 2–606)

To display the Motion Mixer in a viewport: (page 2–607)

To remove the Motion Mixer from a viewport: (page 2–607)

To add a biped to the Motion Mixer: (page 2–608)

To add a non-biped object to the Motion Mixer: (page 2–608)

To see the motions in the Motion Mixer on the biped in the scene: (page 2–611)

Loading Clips

To import BIP clips from a file: (page 2–610)

To import clips from the Reservoir: (page 2–610)

To import clips from a motion flow script: (page 2–611)

To replace one clip with another: (page 2–612)

To replace a clip with motion from a biped in the scene: (page 2–612)

To clone a clip: (page 2–612)

Clip Timing

To move a clip in time within a track: (page 2–611)

To move all clips horizontally on one track: (page 2–611)

To move a clip to another track: (page 2–611)

To change the length of a clip without changing its speed: (page 2–615)

To change the speed of an entire clip: (page 2–616)

To prepare to warp a clip's time: (page 2–621)

To warp a clip's time: (page 2–621)

To add multiple time warps to a clip: (page 2–622)

Trackgroups

To add a new trackgroup for the biped: (page 2–614)

To filter and name a trackgroup: (page 2–614)

To convert a clip track from one type to another: (page 2–608)

To add a Transition or Layer track to a trackgroup: (page 2–608)

Transitions

To create a transition between two clips: (page 2–618)

To change the focus on a foot-based transition with Biped animation: (page 2–619)

To adjust a weight curve: (page 2–620)

Biped Balance

To adjust balance using the balance track: (page 2–623)

To fine-tune balance compensation on the pelvis and spine: (page 2–624)

Saving Mixer Data

To perform a mixdown and copy it to the biped: (page 2–625)

To save Motion Mixer data to a MIX file: (page 2–626)

Reservoir

To access the Reservoir: (page 2–627)

To add clips to the Reservoir: (page 2–627)

To replace a clip in the Reservoir: (page 2–627)

To save an instanced clip to a new clip: (page 2–628)

To save multiple instanced clips from the Reservoir: (page 2–628)

To remove clips not used in the Motion Mixer from the Reservoir: (page 2–628)

Workbench

General

To display a biped body part curve in the Workbench, do one of the following: (page 2–1016)

To hide or unhide the Tab panel: (page 2–1016)

Workbench Panels

To analyze a curve: (page 2–1017)

To fix an individual error: (page 2–1020)

To fix multiple errors: (page 2–1020)

To remove keys from curves automatically: (page 2–1021)

To filter a track: (page 2–1023)

15 Lights and Cameras



"Little Village Far, Far Away"
Copyright 2000 Eni Oken

Lights and cameras are scene objects that simulate their real-world counterparts.

Lights provide lighting for the geometry of a scene: they can light the scene from "offstage," or (with a little extra work), they can appear within the scene itself. Standard lights are simple and easy to use. Photometric lights are more complex, but provide a physically accurate model of real-world lighting. The Daylight and Sunlight systems create outdoor lighting that simulates sunlight based on location and time of day, month and year. You can animate the time of day to create shadow studies.

Cameras frame the scene, providing a controllable point of view. You can animate camera movement. Cameras can simulate some aspects of real-world photography, such as depth-of-field and motion blur.

See also

Lights (page 2–1272)

Cameras (page 2–1365)

Lights

Create panel > Lights

Create menu > Lights



Nighttime scene with artificial lighting

Lights are objects that simulate real lights such as household or office lamps, the light instruments used in stage and film work, and the sun itself. Different kinds of light objects cast light in different ways, emulating different kinds of real-world light sources.

When there are no lights in a scene, the scene is shaded or rendered with default lighting. You add lights to give the scene a more realistic appearance. Lighting enhances the clarity and three-dimensionality of a scene. In addition to general lighting effects, lights can be used to project images. (See *Advanced Effects Rollout (page 2–1341)*.)

Light objects *replace* the default lighting. As soon as you create a light, the default lighting is turned off. If you delete all the lights in the scene, default lighting is turned back on. The default lighting consists of two invisible lights: one is above and to the left of the scene, and the other is below and to the right.

Tip: One way to begin your work on lighting a scene is to convert the default lighting into light objects by using the command *Add Default Lights To Scene (page 1–49)*.

Note: A scene's lighting is also affected by the Ambient Light setting on the Environment And Effects dialog > *Environment panel (page 3–272)*.

Types of Lights

3ds Max provides two types of lights: standard and photometric. All types are displayed in viewports as light objects. They share many of the same parameters, including shadow generators.

Standard Lights

Standard lights (page 2–1288) are computer-based objects that simulate lights such as household or office lamps, the light instruments used in stage and film work, and the sun itself. Different kinds of light objects cast light in different ways, simulating different kinds of light sources. Unlike photometric lights, Standard lights do not have physically-based intensity values.

Photometric Lights

Photometric lights (page 2–1301) use *photometric (page 3–993)* (light energy) values that enable you to more accurately define lights as they would be in the real world. You can set their distribution, intensity, color temperature, and other characteristics of real-world lights. You can also import specific photometric files available from lighting manufacturers to design lighting based on commercially available lights.

Tip: Combine Photometric lights with the *radiosity solution (page 3–51)* to generate physically accurate renderings or perform lighting analysis.

Tips

- You can animate not only the location of a light, but also its color, intensity and some other

creation parameters. See *Animating Lights (page 2–1282)*.

- You can use the *Place Highlight (page 1–467)* command to change a light's position. See the Procedures in *Working with Lights (page 2–1274)*.
- A *Light viewport (page 3–750)* can be a useful way to adjust lights other than omni lights.
- To simulate sunlight, use a *daylight or sunlight system (page 1–418)*, which allows you to set the date, time, and geographic location of your model. The daylight system is photometric, while the sunlight system uses a standard directional light.

  Note: The standard *Skylight (page 2–1296)* light is distinct from the photometric daylight lights. The Skylight light is for use with *light tracing (page 3–44)*.

See also

Working with Lights (page 2–1274)

Using Lights (page 2–1274)

Light Include/Exclude Tool (page 2–1283)

Light Lister (page 2–1285)

Animating Lights (page 2–1282)

Name and Color Rollout (Lights)

Create panel > Lights > Create a light. > Name and Color rollout

Create menu > Photometric Lights > Create a light. > Modify panel > Name and Color rollout

Create a light. > Modify panel > Name and Color rollout

The Name And Color rollout lets you change the name and geometry color of a light. Changing the color of the light geometry can be useful when working with many lights. For example, in a scene with many different types of lights, you could make all spotlights red, and all omni lights blue to easily distinguish them.

Changing a light's geometry color has no effect on the color of the light itself. You can set the color the light emits on its *Intensity/Color/Distribution rollout (page 2–1352)* (for photometric lights), or its *Intensity/Color/Attenuation rollout (page 2–1345)* (for standard lights).

Procedure

To change the color of a light's geometry:

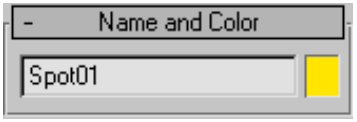
1. Create or select a light in your scene.
2. In the Name and Color rollout, click the color swatch to open a *Color Selector (page 1–161)*.
3. Choose a new color and click OK.

To change the name of a light:

1. Create or select a light in your scene.
2. In the Name and Color rollout, click the name field and enter the new name, then press the Enter key.

   The light's name has changed.

Interface



Name—The name of the selected light.

Note: When you rename a target-type light, the target object will be renamed to match the light.

Color—The color of the light's geometry.

(This has no effect on the color the light emits.)

Using Lights

In general, these are the reasons to use *light objects (page 2–1272)*:

- To improve the illumination of a scene.

  The default illumination in viewports might not be bright enough, or it might not illuminate all faces of a complicated object.

- To enhance a scene's realism through realistic lighting effects.

  Guidelines for Lighting (page 2–1280) has suggestions about making lighting appear realistic.

- To enhance a scene's realism by having lights cast shadows.

  All kinds of lights can cast shadows. Also, you can selectively control whether an object casts or receives shadows. See *Shadow Parameters (page 2–1337)*.

- To cast projections in a scene.

  All kinds of lights can project still or animated maps. See the Projector Map group in the *Advanced Effects rollout (page 2–1341)*.

- To help model a source of illumination in the scene, such as a flashlight.

  Light objects don't render, so to model a source of illumination, you also need to create geometry that corresponds to the light source. Use a *self-illuminating (page 3–1009)* material to make the geometry appear as if it's emitting light.

- To create lighting scenes using manufacturers' *IES (page 2–1328)*, *CIBSE (page 3–921)*, or *LTLI (page 3–964)* files.

  You can visualize commercially available lighting in your model by creating *photometric lights (page 2–1301)* based on manufacturer's photometric data files. By experimenting with different fixtures, and varying the light intensity and color temperature, you can design a lighting system that produces the results you want. See *Photometric Lights: Web Distribution (page 2–1325)*.

See also

Working with Lights (page 2–1274)

Properties of Light (page 2–1276)

Lighting in 3ds Max (page 2–1279)

Guidelines for Lighting (page 2–1280)

Positioning Light Objects (page 2–1282)

Animating Lights (page 2–1282)

Light Include/Exclude Tool (page 2–1283)

Light Lister (page 2–1285)

Working with Lights

The procedures in this topic apply to both standard and photometric lights.

Here are some general tips about working with lights:

- One simple way to light a scene is to convert the default lighting into light objects by using the command *Add Default Lights To Scene (page 1–49)*.

  Note: Add Default Lights To Scene works only if you have used the *Viewport Configuration dialog (page 3–853)* to have the scene use two default lights.

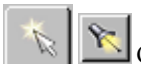
- You can turn the display of light objects on and off with an option in the *Display panel (page 3–775)*. See the "Procedures" section, below.

- You can use the *Place Highlight (page 1–467)* button to change a light's position. See the "Procedures" section, below.

- A *Light viewport (page 2–1338)* can be a useful way to adjust spotlights in your scene.

Procedures

To create a light:

1.  On the Create panel, click Lights.
2. Choose Standard or Photometric from the drop-down list. Standard is the default.
3. In the Object Type rollout, click the type of light you want to create.
4. Click a viewport to create the light. This step varies slightly depending on the type of light. For example, if the light has a target, you drag and click to set the target's location.

   Light objects replace the default lighting. As soon as you create a light, the default lighting is turned off. If you delete all lights in the scene, the default lighting is restored.
5. Set the creation parameters.

   Like all objects, lights have a name, a color, and a General Parameters rollout.

To create shadows, do one of the following:

- In the General Parameters rollout, make sure On is checked in the Shadows group. Adjust shadow parameters in the Shadow Parameters rollout and the additional (*Shadow Map (page 2–1363)*, *Advanced Ray-traced (page 2–1356)*, *Area Shadows, (page 2–1357)* or *Ray Traced Shadows (page 2–1363)*) shadow rollouts.
- Right-click the light, and turn on Cast Shadows in the Tools 1 (upper-left) quadrant of the quad menu.

  Turning on Cast Shadows also turns on the On toggle in the Shadows group of the General Parameters rollout, and vice versa.

  Shadows are visible only when rendered, either in a full *rendering (page 3–1)* or by *ActiveShade (page 3–17)*.

  Tip: To turn shadows on or off for multiple lights, select the lights and then use the *Light Lister (page 2–1285)*.

  You can set an object to not cast or not receive shadows. By default, objects do both. See *Object Properties (page 1–117)*.

To control the display of light objects:

-  On the Display panel, on the Hide By Category rollout, turn on Lights.

  All light objects in the scene disappear, but the lighting itself is unchanged.

  Light objects can cast light whether or not their display is turned off. The Zoom Extents commands are affected by whether light icons are displayed or not. When lights are displayed, *Zoom Extents (page 3–740)* and *Zoom Extents All (page 3–740)* includes the lights in the zoom.

  Tip: To control whether a light casts light in the scene, you can use its On toggle, or you can toggle its Renderable property on the light's *Object Properties dialog (page 1–117)*.

To change a light's parameters:

1. Select the light.

   Tip: Lights can be hard to select by clicking. Use the keyboard shortcut **H** to select the light by name.
2.  Open the Modify panel.
3. Change the light's parameters in the General Parameters rollout and other rollouts available for that light.

To position a light so it highlights a face:

1. Make sure the viewport you plan to render is active, and that the object you want to highlight is visible in it.

   The result of Place Highlight depends on what is visible in the viewport.

2. Select a light object.

3.  On the Main toolbar, choose Place Highlight from the *Align flyout (page 1–462)*.

   You can also choose Tools menu > Place Highlight.

4. Drag over the object to place the highlight.

   When you place an *omni light (page 2–1295)*, *free point (page 2–1304)*, *free linear (page 2–1307)*, *free area (page 2–1309)* or *free direct light (page 2–1293)*, the program displays a face normal for the face the mouse indicates.

   When you place a *target point (page 2–1303)*, *target area (page 2–1307)*, *target linear (page 2–1305)* or *target direct light (page 2–1292)*, the program displays the light's target and the base of its cone.

5. Release the mouse when the normal or target display indicates the face you want to highlight.

   The light now has a new position and orientation. You can see the highlight illumination in shaded viewports that show the face you chose, and when you render those views.

   Place Highlight (page 1–467) works with any kind of selected object. You can also use Place Highlight with a selection set of multiple objects. All objects maintain their initial distance from the face.

   Note: For materials, highlight rendering depends on the material's specular properties and the type of rendering you use.

Properties of Light

This topic describes light in the real world. When you light a scene, it can be helpful to know how light naturally behaves.

When light rays strike a surface, the surface reflects them, or at least some of them, enabling us to see the surface. The appearance of a surface depends on the light that strikes it combined with the properties of the surface material, such as color, smoothness, and opacity.

Materials (page 2–1395) let you specify the visual properties of surfaces.

Intensity

The intensity of light at its point of origin affects how brightly the light illuminates an object. A dim light cast on a brightly colored object shows only dim colors.



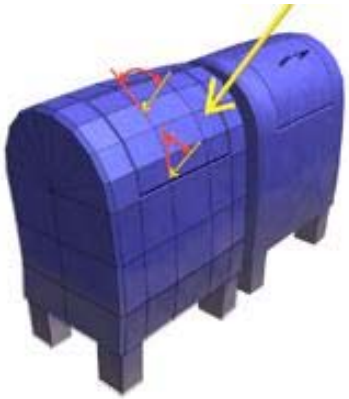
Left: A room lit by candles, which are a low-intensity source. Right: The same room lit by a higher-intensity light bulb.

Angle of Incidence

The more a surface inclines away from a light source, the less light it receives and the darker it appears. The angle of the surface normal relative to the light source is known as the *angle of incidence*.

When the angle of incidence is 0 degrees (that is, the light source strikes the surface perpendicularly), the surface is illuminated with the full intensity of the light source. As the angle of

incidence increases, the intensity of illumination decreases.

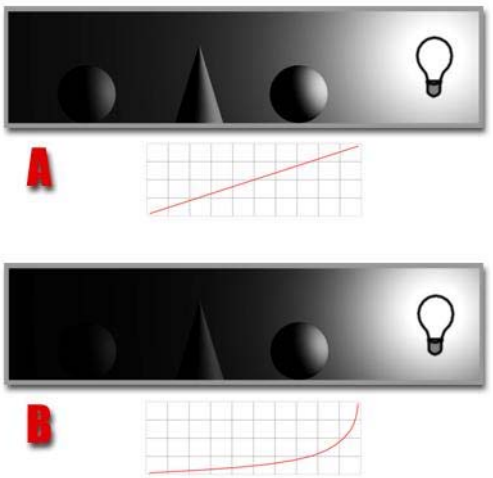


Angle of incidence affects intensity.

Attenuation

In the real world, light diminishes over distance. Objects far from the light source appear darker; objects near the source appear brighter. This effect is known as *attenuation*.

In nature, light attenuates at an inverse square rate. That is, its intensity diminishes in proportion to the square of the distance from the light source. It is common for attenuation to be even greater when light is dispersed by the atmosphere, especially when there are dust particles in the atmosphere, or fog or clouds.



A. Inverse decay

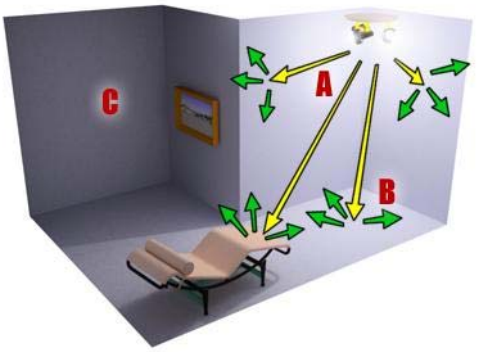
B. Inverse square decay

The graphs show the decay curves.

Reflected Light and Ambient Light

The light an object reflects can illuminate other objects. The more light a surface reflects, the more light it contributes to illuminating other objects in its environment.

Reflected light creates *ambient light*. Ambient light has a uniform intensity and is uniformly diffuse. It has no discernible source and no discernible direction.



A. Direct light

B. Reflected light

C. Resulting ambient light

Color and Light

The color of light depends partly on the process that generates the light. For example, a tungsten lamp casts orange-yellow light, a mercury vapor lamp casts cold blue-white light, and sunlight is yellow-white. Light color also depends on the medium the light passes through. For example, clouds in the atmosphere tint daylight blue, and stained glass can tint light a highly saturated color.

Light colors are *additive colors;* the primary light colors are red, green, and blue (RGB). As multiple colored lights mix together, the total light in the scene gets lighter and eventually turns white.



Additive mixing of colored lights

Color Temperature

Color *temperature* describes a color in terms of degrees Kelvin (K). This is useful for describing the color of light sources and other color values that are close to white. The following table shows color temperatures for some common types of light, with the equivalent hue number (from the HSV color description).

If you use these hue numbers for lights in a scene, set the value to full (255) and then adjust the saturation to meet the needs of your scene. Mentally we tend to correct light color so that objects appear to be lit by white light; usually the effect of color temperature in a scene should be subtle.

Light source	Color Temperature	Hue
Overcast daylight	6000 K	130
Noontime sunlight	5000 K	58
White fluorescent	4000 K	27
Tungsten/halogen lamp	3300 K	20
Incandescent lamp (100 to 200 W)	2900 K	16
Incandescent lamp (25 W)	2500 K	12

Light source	Color Temperature	Hue
Sunlight at sunset or sunrise	2000 K	7
Candle flame	1750 K	5

Lighting in 3ds Max

Lighting in 3ds Max simulates natural lighting. However, standard lights are simpler than natural lighting. Using *photometric lights (page 2–1301)* with a *radiosity solution (page 3–51)* with your lights provides a better model of the real world.

Intensity

The intensity of a standard light is its HSV Value. At full value (255), the light is at its brightest; at 0, the light is completely dark.

Note: See *Designing Materials (page 2–1395)* for more information about material color and how it interacts with light intensity.

The intensity of a photometric light is set by a real-world intensity value, measured in either lumens, candelas, or lux. See *Intensity/Color/Distribution Rollout (page 2–1352)*.

Angle of Incidence

3ds Max uses a vector from the light object to the face, along with the face normal, to calculate the angle of incidence.

A surface is fully illuminated when the angle of incidence is 0 degrees (that is, the light source strikes the surface perpendicularly). If the angle of incidence increases, attenuation is in effect, or if the light has a color, the surface intensity can be reduced.

In other words, the position and orientation of the light, relative to the object, are what control the angle of incidence in a scene. The *Place Highlight command (page 1–467)* is one way to fine-tune the location of a light.

Attenuation

For standard lights, *attenuation (page 3–912)* is turned off by default. To shade or render a scene with attenuation, you turn it on for one or more lights. All types of standard lights support attenuation. You can set explicitly where attenuation begins and where it ends. This is partly so you don't have to worry about setting up strictly realistic distances between light objects and the objects they illuminate. More importantly, this feature lets you fine-tune the effect of attenuation.

In outdoor scenes, attenuation can enhance the effect of distance. (Another way to model environmental effects is to use the atmospheric settings when you render. See *Environment and Atmosphere Effects (page 3–271)*.) In an indoor setting, attenuation is useful for low-intensity light sources such as candles.

Photometric lights always attenuate, using an inverse-square falloff, as in nature. (In the case of the IES Sun Light, its great intensity makes its attenuation hardly apparent.)

Reflected Light and Ambient Light

Rendering with the default renderer and standard lights does not calculate the effect of lights reflected from objects in the scene. Because of this, lighting a scene with standard lights often requires you to add more light objects than would be needed in real life. You can, however, use *radiosity (page 3–51)* to show the results of reflected light.

When you do not use a radiosity solution, you can use the *Environment panel (page 3–272)* to adjust the color and intensity of ambient light. Ambient light affects contrast. The higher the intensity of ambient light, the lower the contrast in the scene. The color of ambient light tints the scene. Sometimes ambient light is bounced light

that gets its color from other objects in the scene. Most of the time, however, the color of ambient light should be the complement of the color of the principal light source for the scene.

Tip: To better simulate reflected light and variations in it due to the varying reflectivity of objects in the scene, you can add more lights to a scene and set them to exclude the objects you don't want them to affect. You can also set up lights to affect only the ambient component of surfaces. See *General Lighting Parameters (page 2–1331)*.

Color

You can set the color of 3ds Max lights. You can use the RGB values for color temperatures as a guide for the principal lighting of a scene; see *Properties of Light (page 2–1276)*. Be aware, however, that we tend to perceive scenes as always being lit by white light (this is a perceptual phenomenon known as *color constancy),* so accurately reproducing the color of a light source can make the rendered scene appear to be tinted oddly. Use the light source values as a general guideline only.

Guidelines for Lighting

The guidelines for lighting used by photographers, filmmakers, and stage designers can also help you set up the lighting for scenes in 3ds Max.

Your choice of lighting depends on whether your scene simulates natural or artificial illumination. Naturally lit scenes, such as daylight or moonlight, get their most important illumination from a single light source. Artificially lit scenes, on the other hand, often have multiple light sources of similar intensity.

Note: If you use standard instead of photometric lights, both kinds of scenes require multiple secondary light sources for effective illumination.

Whether a scene is indoors or outdoors can also affect your choice of material colors. See *Designing Materials (page 2–1395)*.

Natural Light



Outdoor scene with natural sunlight

At ground level, for practical purposes, sunlight has parallel rays coming from a single direction. The direction and angle vary depending on the time of day, the latitude, and the season.

In clear weather, the color of sunlight is a pale yellow: for example, RGB values of 250, 255, 175 (HSV 45, 80, 255). Cloudy weather can tint sunlight blue, shading into dark gray for stormy weather. Particles in the air can give sunlight an orange or brownish tint. At sunrise and sunset, the color can be more orange or red than yellow.

3ds Max specifically provides a *daylight system* to simulate the sun. See *Sunlight and Daylight Systems (page 1–418)*. A single daylight system is appropriate as the main light source for sunlit scenes.

Shadows are more distinct the clearer the day is, and can be essential for bringing out the three-dimensionality of a naturally lit scene.

A directional light can also simulate moonlight, which is white but dim compared to the sun.

Artificial Light



Outdoor scene with natural twilight and one streetlight

Artificial light, whether used indoors or outdoors at night, uses multiple lights. The following guidelines are for creating normally lit, easily legible scenes. You don't have to follow the guidelines, of course, but then you call attention to the lighting itself, rather than to the subject of the scene.

The subject of a scene should be lit by a single bright light, known as the *key light*. Position the key light in front of the subject and slightly above.

In addition to the key light, position one or more other lights to illuminate the background and the side of the subject. These are known as *fill lights*. Fill lights are less bright than the key light.

When you use only one fill light, the angle at ground level between it, the subject, and the key light should be approximately 90 degrees.

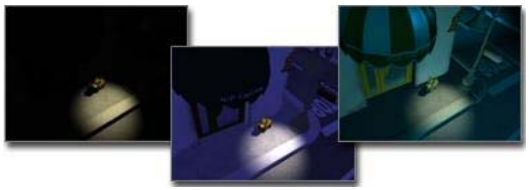
Key-and-fill lighting emphasizes the subject of a scene. It also emphasizes the three-dimensionality of the scene.

In 3ds Max, a spotlight is usually best for the key light, and either spotlights or omni lights are good for creating the fill lighting. See *Target Spot (page 2–1289)*, *Free Spot (page 2–1290)*, and *Omni (page 2–1295)*. Ambient light can be another element of your fill lighting.

You can also add lights to emphasize secondary subjects in a scene. In stage terminology, these lights are known as *specials*. Special lights are usually brighter than the fill light but less bright than the main key light.

To design using physically based energy values, distributions, and color temperature, you can create *photometric lights (page 2–1301)*.

Ambient Light



Left: No ambient light

Middle: Default ambient light

Right: User-adjusted ambient light

Ambient light in 3ds Max simulates the general illumination from light reflecting off diffuse surfaces. Ambient settings determine the illumination level of surfaces in shadow, or those not receiving direct illumination from light sources. The Ambient level on the Environment dialog establishes the scene's basic illumination level before any light sources are taken into account, and is the dimmest any portion of the scene can ever become.

Ambient light is most often used for exterior scenes, when the sky's broad lighting produces an even distribution of reflected light to surfaces not in direct sun. A common technique for deepening the shadows is to tint the ambient light color to be the complement of the scene's key light.

Unlike the outside, interior scenes typically have numerous lights, and a general ambient light level is not ideal for simulating the diffuse reflection of local light sources. For interiors, it's common to set the scene's environment ambient level to black,

and use lights that effect ambient only to simulate the regional areas of diffuse reflection.

You set the scene's ambient light using the Environment And Effects dialog > *Environment panel (page 3–272)*. You set a light to affect only ambient illumination with its *Advanced Effects rollout (page 2–1341)* > Ambient Only check box.

Positioning Light Objects

Once you have placed lights in your scene, you can use transforms to change a light's position or orientation.

Transforming Lights

Use transforms on light objects as follows:

 **Move:** Use *Move (page 1–439)* to change the position of lights. You can also use it to change the position of light targets.

 **Rotate:** Use *Rotate (page 1–439)* to change the orientation of lights.

Note: You can't rotate a target light about its local X or Y axes. Instead, use Move to move the light or its target. Rotating the light about its local Z axis can be useful if the light uses a rectangular beam or projects a bitmap.

Rotating a plain omni light or a photometric isotropic light has no effect as these lights cast light uniformly in all directions. However, rotating an omni or isotropic projector light causes the projected image to rotate.

Rotating a photometric isotropic light has no effect as these lights cast light uniformly in all directions. However, rotating an isotropic projector light causes the projected image to rotate.

 **Scale:** Scaling Point, Linear, or Area lights has no effect. Using *Scale (page 1–441)* with spotlights and directional lights changes the size of their light beam and attenuation ranges. Scaling omni lights changes only the attenuation ranges. Scaling photometric lights changes their attenuation rate.

Light viewports (page 3–750) are another convenient way to transform and change parameters of spotlights and directional lights.

Tip: When you adjust lights, it can be useful to turn on *Degradation Override (page 1–34)*. If Override is off and shaded viewports begin to display in wireframe, you can't see the result of the changes you make to lights.

Placing Highlights

You use *Place Highlight (page 1–467)* to position a light to create a specular highlight at a designated point on an object. Place Highlight is one of the buttons on the Align flyout. Place Highlight moves or rotates the selected light object to aim it at a face on an object you pick. The light maintains its original distance from the face. Place Highlight works with any kind of selected object. You can also use Place Highlight with a selection set that contains more than one object. All objects maintain their initial distance from the face.

Animating Lights

You animate lights by using transforms or changing creation parameters on different keyframes while the *Auto Key button (page 3–717)* is on. During animation, light transforms and parameter values are interpolated between keyframes.

Moving and Rotating Light Objects

You can use the following methods to move and rotate light objects.

- Move an *omni light (page 2–1295)* when it's a "practical" light within a scene (a light that appears in the scene itself). Combine the light with a self-illuminating geometric object. If you want to move a target type of light, select both the light and its target to animate them together.
- Use a *free spotlight (page 2–1290)* when a spotlight is to move within the scene. Free spotlights are especially intended to be animated along a path, using a *Path constraint (page 2–398)*. Unlike target spotlights, free spotlights can bank as they travel. Use *target spotlights (page 2–1289)* when the light position is fixed.
- If you do need to move a target spotlight, link both the light and its target to a *dummy object (page 2–16)*, then assign the path constraint to the dummy object.
- Use a *LookAt Constraint (page 2–406)* to have a spotlight track a moving object.

  If the spotlight is a target spotlight, its previous target is ignored.

  If the spotlight is a free spotlight, it effectively becomes a target spotlight, with the looked-at object the target.
- Change the parameters of a *daylight system (page 1–418)* or sunlight system to simulate different times of day or year.

Animating Light Creation Parameters

The following techniques can be used to animate Modify panel parameters for lights.

- To dim or brighten a light over time, animate its Multiplier parameter.
- To change the color of a light over time, animate its color parameters. Use a *smooth tangent (page 2–305)* for color change keys unless you want the color to change abruptly.
- To make a standard light flash on and off, set its Multiplier to 0 in repeated keyframes, and assign a *step tangent (page 2–305)* to this parameter.

Light Include/Exclude Tool

By default, the Light Include/Exclude tool is unavailable in the 3ds Max user interface. You can add it as a custom keyboard shortcut, quad or menu item, or toolbar button by using the *Customize User Interface dialog (page 3–792)*.

The Light Include/Exclude tool is a modeless dialog that lets you include or exclude objects on a light-by-light basis. When excluded, an object is not illuminated by the selected light and receives no shadows.

This dialog requires at least one light object in your scene. For a selected light, this dialog is a shortcut to the same functionality available on the *Exclude/Include dialog (page 2–1335)*.

Although light exclusion does not occur in nature, this feature is useful when you need exact control over the lighting in your scene. Sometimes, for example, you'll want to add lights specifically to illuminate a single object but not its surroundings, or you'll want a light to cast shadows from one object but not from another.

Note: By default, no objects are excluded for a new light. It's necessary only to include objects that have previously been excluded.

See also

Lights (page 2–1272)

Light Lister (page 2–1285)

Procedure

To exclude an object from receiving light:

1. In the scene, select the object or objects to exclude.
2. Open the Light Incl/Exl dialog.

See the path annotation, above. To use the Light Include/Exclude tool, you have to create a custom keyboard shortcut, quad or menu item, or toolbar button.

3. In the Geometry group, choose Exclude.
4. Click the Assign To Light button.
5. In the scene, click a light. If the light is hard to locate, use the *Select Objects dialog (page 1–78)* (keyboard shortcut H) to pick the light.

   The object is now excluded from the light. The Objects list at the bottom of the dialog lists the object.
6. Render to see the effect.

Interface



Geometry group

Include, Exclude—Toggles the state of a selected object to receive light from a particular light object. Default=include.

Assign to Light—Activates selection so you can choose the light you want to use. You can only choose one light at a time.

List Light Properties group

These controls let you view and edit the include/exclude status of objects on a light-by-light basis.

Current displayed light—Names the currently selected light.

Choose Light—Activates selection so you can choose another light.

Clear Light—Empties the Objects list, removing any included or excluded objects assigned to the currently displayed light.

Include, Exclude—Reverses the state of included or excluded objects to receive light from the currently displayed light.

Objects—Lists objects selected for inclusion or exclusion by the currently displayed light.

Help—Provides a quick reminder of procedures.

Light Lister

Tools menu > Light Lister

Light Lister is a modeless dialog that lets you control a number of features for each light. You can also make global settings that affect every light in your scene.

To display information, this dialog requires at least one light object in your scene. For a selected light, this dialog is a shortcut to the same functionality available on the Modify panel > Parameters rollout. Global settings are duplicated on the *Environment panel (page 3–272)*.

Note: The Light Lister cannot control more than 150 unique light objects at a time (instances of a light don't count). If there are more than 150 unique lights in your scene, the Lister displays controls for the first 150 it finds, and a warning that you should select fewer lights. Select fewer lights and then use the Selected Lights configuration.

See also

Light Include/Exclude Tool (page 2–1283)

Procedures

To use global settings:

1. On the Light Lister dialog, on the Configuration rollout, choose General Settings.

   The General Settings rollout appears.

2. Make changes to the settings. See Interface, below.

To set individual lights:

1. On the Configuration rollout, choose All Lights. The Lights rollout displays settings for all the scene lights (subject to the limit of 150 lights described above). Alternatively, you can select the lights to adjust, and then on the Configuration rollout choose Selected Lights.

   Tip: If you change the light selection, the Light Lister does not update interactively. Click Refresh to update the list.

2. On the Lights rollout, change the settings for any light on the list. See Interface, below.

   Some changes show up immediately in the viewport.

Interface

Configuration rollout



Chooses which lights and controls to affect, and lets you update light settings.

All Lights—The Lights rollout shows all lights in the scene (subject to the 150-light restriction, described above).

Selected Lights—The Lights rollout shows only selected lights.

General Settings—Displays the General Settings rollout.

Refresh—Updates the list of lights to use the current light selection (if Selected Lights is active) and current scene settings.

Note: When the Light Lister needs to be updated, the Refresh button is highlighted in yellow.

General Settings rollout



These controls are for general lighting settings.

- **Selected Lights—**When chosen, general settings affect only selected lights.
- **All Lights—**When chosen, general settings affect all lights in the scene (subject to the 150-light restriction, described above).

On—When on, affected lights are active in the scene. When off, affected lights go dark in the viewports and renderings.

Multiplier—Increases or decreases the light intensity of standard lights. See *General Lighting Parameters (page 2–1331)* for standard lights.

Color—By default, the color of all standard lights is white. Click to display the *Color Selector (page 1–161)* and change the light color for affected lights.

Warning: For photometric lights, this changes the filter color, *not* the color temperature.

Shadows—When on, affected lights cast shadows. When off, they don't cast shadows.

Shadow type drop-down list—Selects the shadow type for affected lights.

Map Size—Sets the size (in pixels squared) of the shadow map used by affected lights.

Bias—Setting depends on which shadow type is selected. In general, bias moves the shadow toward or away from the shadow-casting object. For specific effects, see *Shadow Map Parameters (page 2–1363)* and *Ray-Traced Shadow Parameters (page 2–1363)*.

Sm. Range (Sample Range)—For shadow-mapped shadows, sets the Sample Range value. See *Shadow Map Parameters (page 2–1363)*. This value has no effect for ray-traced or area shadows.

Transp. (Transparency)— When on, turns on transparency for advanced ray-traced and area shadows. Has no effect on shadow-mapped or standard ray-traced shadows. Default=off.

Int. (Integrity)—For advanced ray-traced shadows or area shadows, sets the Shadow Integrity. See *Advanced Ray-Traced Parameters Rollout (page 2–1356)* or *Area Shadows Rollout (page 2–1357)*.

Qual. (Quality)—For advanced ray-traced shadows or area shadows, sets the Shadow Quality. See *Advanced Ray-Traced Parameters Rollout (page 2–1356)* or *Area Shadows Rollout (page 2–1357)*.

Decay—(For standard lights.) Sets the type of decay: None, Inverse, or Inverse Square. See *Intensity/Color/Attenuation Parameters (page 2–1345)*.

Start—(For standard lights.) Sets the start range for decay. See *Intensity/Color/Attenuation Parameters (page 2–1345)*.

Length—(For photometric lights.) Sets the Length value for Linear and Area lights. See *Linear Light Parameters Rollout (page 2–1354)* or *Area Light Parameters Rollout (page 2–1354)*.

Width—(For photometric lights.) Sets the Width value for Area lights. See *Area Light Parameters Rollout (page 2–1354)*.

Global Tint—Adds a color tint to all lights in the scene except ambient light. Click to use the *Color Selector (page 1–161)*. The tint is in addition to the

global light color, or the color of individual lights. The default setting, white, has no tinting effect.

Global Level—Increases or decreases the overall lighting level for Standard lights. Default=1.0.

Note: This setting, designed for standard lights, reduces the level of photometric lights to near darkness. If you have photometric lights in your scene, leave this setting at the default.

Ambient Color—Changes the *ambient color (page 3–908)*, the color seen in shadows. Click to use the *Color Selector (page 1–161)*.

Lights rollout



This rollout is visible while All Lights or Selected Lights is active on the Configuration rollout. Its controls are for individual light objects. This rollout displays two lists: one for standard lights, and another for photometric lights.

Blank gray button (Select)—Click to select the named light. This button turns white for selected lights. Selecting a light opens the Modify panel for that light. This button has a gray box in the middle for the light that appears in the Modify panel. The gray box appears whether the light is selected or not.

Clicking Select deselects all other lights. If the entry is for an instanced light, clicking Select selects all instances (see the description of the Name field that follows). If the light is part of a group, the entire group is selected. If you are in Selected Lights mode (on the Configuration rollout), the list of selected lights is *not* refreshed automatically.

Note: You can set the controls for a light without selecting it.

Check box—Turns the light on or off. Default=on.

Name—Shows the name of the light object.

If there are multiple instances of a light, only one entry appears in the Light Lister, and the Name field becomes a drop-down list. The list lets you see the names of all instances, but it has no other effect: parameter changes made in the Light Lister affect all instances of the light.

Multiplier—(Standard lights only.) Increases or decreases the intensity of the light. Default=1.0.

Note: This control is unavailable if the light has a procedural controller.

Intensity—(Photometric lights only.) Sets the light intensity, in candelas. Default=1500.0.

Note: This control is unavailable if the light has a procedural controller.

Color—Click to display the *Color Selector (page 1–161)* and change the light color. Default=white.

Warning: For photometric lights, this changes the filter color, *not* the color temperature.

Shadows—Turns shadow casting on or off. Default=on.

Shadow type drop-down list—Selects the shadow type for the light.

Map Size—Available only when shadow map shadows is on. Sets the size (in pixels squared) of the shadow map used by all lights.

Bias—Setting depends on which shadow type is selected. In general, bias moves the shadow toward or away from the shadow-casting object. See *Shadow Map Parameters (page 2–1363)* and *Ray-Traced Shadow Parameters (page 2–1363)* for specific effects.

Sm. Range—For shadow-mapped shadows, sets the Sample Range value. See *Shadow Map Parameters (page 2–1363)*. This value has no effect for ray-traced shadows. Default=4.0.

Transp. (Transparency)—When on, turns on transparency for advanced ray-traced and area shadows. Has no effect on shadow-mapped or standard ray-traced shadows. Default=off.

The shadow transparency control is on the *Optimizations rollout (page 2–1361)*.

Int. (Integrity)—For advanced ray-traced shadows or area shadows, sets the Shadow Integrity. See *Advanced Ray-Traced Parameters Rollout (page 2–1356)* or *Area Shadows Rollout (page 2–1357)*. Default=1.

Qual. (Quality)—For advanced ray-traced shadows or area shadows, sets the Shadow Quality. See *Advanced Ray-Traced Parameters Rollout (page 2–1356)* or *Area Shadows Rollout (page 2–1357)*. Default=2.

Decay—(Standard Lights only.) Sets the type of decay: None, Inverse, or Inverse Square. See *Attenuation Parameters (page 2–1345)*. Default=None.

Start—(Standard lights only.) Sets the start range for decay. See *Attenuation Parameters (page 2–1345)*. Default=0.0.

Length—(Photometric lights only.) Sets the Length value for Linear and Area photometric lights. See *Linear Light Parameters Rollout (page 2–1354)* or *Area Light Parameters Rollout (page 2–1354)*.

Width—(Photometric lights only.) Sets the Width value for Area photometric lights. See *Area Light Parameters Rollout (page 2–1354)*.

Standard Lights

Standard lights are computer-based objects that simulate lights such as household or office lamps, the light instruments used in stage and film work, and the sun itself. Different kinds of light objects cast light in different ways, simulating different kinds of real-world light sources. Unlike *photometric lights (page 2–1301)*, standard lights do not have physically-based intensity values.

There are eight types of standard light objects:

Target Spotlight (page 2–1289)

Free Spotlight (page 2–1290)

Target Direct Light (page 2–1292)

Free Direct Light (page 2–1293)

Omni Light (page 2–1295)

Skylight (page 2–1296)

mr Area Omni Light (page 2–1298)

mr Area Spotlight (page 2–1299)

Most of the parameters for standard lights are common to all five types. These are described in the following topics:

General Lighting Parameters (page 2–1331)

Intensity/Color/Attenuation Parameters (page 2–1345)

Advanced Effects Rollout (page 2–1341)

Shadow Parameters (page 2–1337)

mental ray Indirect Illumination Rollout (for Lights) (page 2–1343)

Spotlights and Directional lights have rollouts specific to them:

Spotlight Parameters (page 2–1338)

Directional Parameters (page 2–1348)

For parameters specific to a particular kind of light, see that light's description.

Shadow Parameters

Different rollouts can appear, depending on which type of shadows you have chosen for a light:

Advanced Ray-Traced Parameters Rollout (page 2–1356)

Area Shadows Rollout (page 2–1357)

Optimizations Rollout (page 2–1361)

mental ray Shadow Map Rollout (page 2–1360)

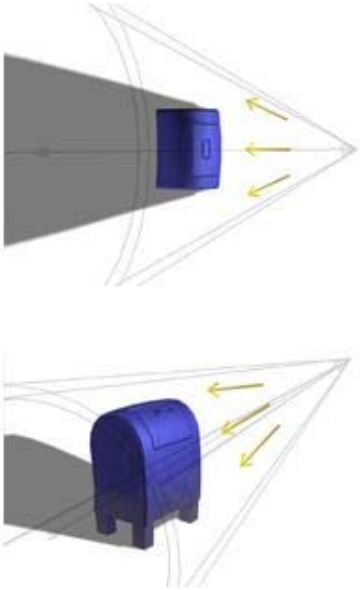
Ray-Traced Shadow Parameters Rollout (page 2–1363)

Shadow Map Parameters Rollout (page 2–1363)

Target Spotlight

Create panel > Lights > Standard > Target Spot button

Create menu > Standard Lights > Target Spotlight



Top: Top view of a target spotlight

Bottom: Perspective view of the same light

A spotlight casts a focused beam of light like a flashlight, a follow spot in a theater, or a headlight. A target spotlight uses a target object to aim the camera.

Note: When you add a target spotlight, the software automatically assigns a *Look At controller (page 2–344)* to it, with the light's target object assigned as the Look At target. You can use the controller settings on the Motion panel to assign any other object in the scene as the Look At target.

See also

Lights (page 2–1272)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

Intensity/Color/Attenuation Parameters (page 2–1345)

Advanced Effects Rollout (page 2–1341)

Shadow Parameters (page 2–1337)

Spotlight Parameters (page 2–1338)

Procedures

To create a target spotlight:

1.  On the Create panel, click Lights.

   Standard is the default choice of light type.
2. In the Object Type rollout, click Target Spot.
3. Drag in a viewport. The initial point of the drag is the location of the spotlight, and the point where you release the mouse is the location of the target.

   The light is now part of the scene.
4. Set the creation parameters.

To adjust a target spotlight:

1. Select the light.
2.  Use Move on the main toolbar to adjust the light. As an alternative, right-click the light and choose Move.

Because the spotlight is always aimed at its target, you can't rotate it about its local X or Y axis. However, you can select and move the target object as well as the light itself. When you move either the light or the target, the light's orientation changes so it always points at the target.

Note: The target's distance from the light does not affect the attenuation or brightness of the light.

To select the target:

The target, displayed as a small square, is often in the same area as objects that you want to illuminate. It can be difficult to select it by clicking.

1. Select the spotlight itself.
2. Right-click the light and choose Select Target from the upper-left (Tools 1) quadrant of the quad menu.

   Clicking the line that connects the light and its target selects both objects. However, region selection doesn't recognize the link line.

   Another way to adjust a spotlight is to use a *Light viewport (page 2–1338)*.

To change a viewport to a Light view:

1. Right-click a viewport label.

   The viewport right-click menu is displayed.
2. Choose Views.

   The name of each spotlight or directional light is displayed in the Views list.
3. Choose the name of the light you want.

   The viewport now shows the light's point of view. You can use the *Light Viewport Controls (page 3–750)* to adjust the light.

   The default keyboard shortcut for switching to a Light viewport is **$**.

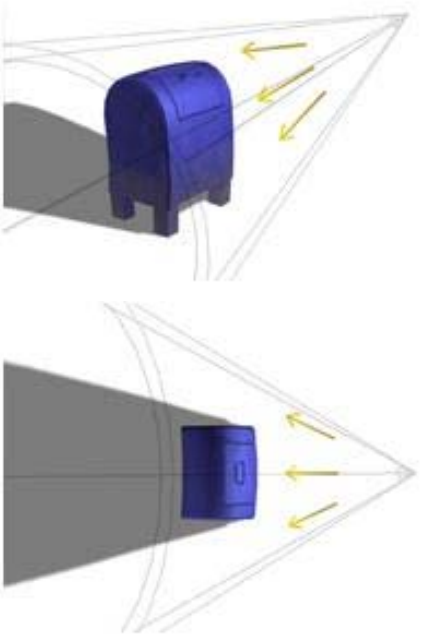
Interface

When you rename a target spotlight, the target is automatically renamed to match. For example, renaming *Light01* to *Klieg* causes *Light01.Target* to become *Klieg.Target*. The target's name must have the extension *.Target*. Renaming the target object does not rename the light object.

Free Spotlight

Create panel > Lights > Standard > Free Spot button

Create menu > Standard Lights > Free Spotlight



Top: Perspective view of a free spotlight

Bottom: Top view of the same light

A spotlight casts a focused beam of light like a flashlight, a follow spot in a theater, or a headlight. Unlike a targeted spotlight, a Free Spot has no target object. You can move and rotate the free spot to aim it in any direction.

See also

Lights (page 2–1272)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

Intensity/Color/Attenuation Parameters (page 2–1345)

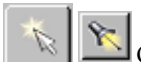
Advanced Effects Rollout (page 2–1341)

Shadow Parameters (page 2–1337)

Spotlight Parameters (page 2–1338)

Procedures

To create a free spotlight:

1.  On the Create panel, click Lights.

   Standard is the default choice of light type.
2. In the Object Type rollout, click Free Spot.
3. Click the viewport location where you want the light to be.

   The light is now part of the scene. It points away from you in the viewport you clicked.

   You can adjust the light's direction with Move and Rotate or by using a Light viewport.
4. Set the creation parameters.

To change a viewport to a Light view:

1. Right-click a viewport label.

   The viewport right-click menu is displayed.
2. Choose Views.

   The name of each spotlight or directional light is displayed in the Views list.
3. Choose the name of the light you want.

   The viewport now shows the light's point of view. You can use the *Light Viewport Controls (page 3–750)* to adjust the light.

   The default keyboard shortcut for switching to a Light viewport is **$**.

Interface

You aim a free spotlight by adjusting its orientation in a scene using Move and Rotate.

The free spotlight is useful when you want a spotlight to follow a path and either don't want to bother with linking a spotlight and target to a *dummy object (page 3–930)*, or you need banking along the path.

General Parameters rollout

When you create a Free Spot light, the Targeted parameter is adjustable on the *General Parameters rollout (page 2–1331)*. This is a fixed value for target lights.

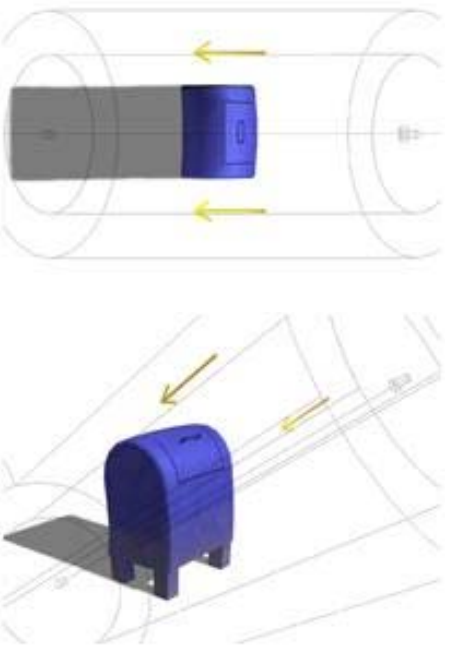
Targeted—When on, 3ds Max sets a point to use as an invisible target about which the Free Spot can orbit. The spinner adjusts the distance to the target.

The target distance is animatable.

Target Direct Light

Create panel > Lights > Standard > Target Direct button

Create menu > Standard Lights > Target Directional Light



Top: Top view of a target directional light

Bottom: Perspective view of the same light

Directional lights cast parallel light rays in a single direction, as the sun does (for all practical purposes) at the surface of the earth. Directional lights are primarily used to simulate sunlight. You can adjust the color of the light and position and rotate the light in 3D space.

A target directional light uses a target object to aim the light.

Because directional rays are parallel, directional lights have a beam in the shape of a circular or rectangular prism instead of a "cone."

Note: When you add a target directional light, the software automatically assigns a *Look At controller (page 2–344)* to it, with the light's target object assigned as the Look At target. You can use the controller settings on the Motion panel to assign any other object in the scene as the Look At target.

Note: Direct lights are supported in a *radiosity solution (page 3–51)* only if they are pointed downwards, outside the boundary box of the scene geometry.

Note: The mental ray renderer assumes that all directional lights come from infinity, so objects that are behind the direct light object in the 3ds Max scene will also be illuminated. In addition, with the mental ray renderer, directional lights cannot generate area shadows, and do not work with the Beam shader (in the lume library).

See also

Lights (page 2–1272)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

Intensity/Color/Attenuation Parameters (page 2–1345)

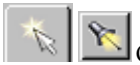
Advanced Effects Rollout (page 2–1341)

Shadow Parameters (page 2–1337)

Directional Parameters (page 2–1348)

Procedures

To create a target direct light:

1.  On the Create panel, click Lights. Standard is the default choice of light type.
2. In the Object Type rollout, click Target Direct.
3. Drag in a viewport. The initial point of the drag is the location of the light, and the point where you release the mouse is the location of the target.

   The light is now part of the scene.

4. Set the creation parameters.

   To adjust the light's direction, move the target object.

To change a viewport to a Light view:

1. Right-click the viewport label.

   The viewport right-click menu is displayed.

2. Choose Views.

   The name of each spotlight or directional light is displayed in the Views list.

3. Choose the name of the light you want.

   The viewport now shows the light's point of view. You can use the *Light Viewport Controls (page 3–750)* to adjust the light.

   The default keyboard shortcut for switching to a Light viewport is **$**.

Interface

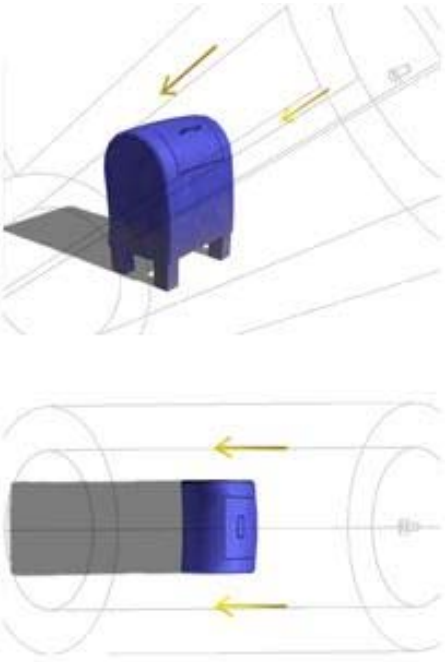
Clicking the line that connects the light and its target selects both objects. However, region selection doesn't recognize the link line.

When you rename a target directional light, the target is automatically renamed to match. For example, renaming *Light01* to *Sol* causes *Light01.Target* to become *Sol.Target*. The target's name must have the extension *.Target*. Renaming the target object does not rename the light object.

Free Direct Light

Create panel > Lights > Standard > Free Direct button

Create menu > Standard Lights > Directional Light



Top: Perspective view of a free directional light
Bottom: Top view of the same light

Directional lights cast parallel light rays in a single direction, as the sun does (for all practical purposes) at the surface of the earth. Directional lights are primarily used to simulate sunlight. You can adjust the color of the light and position and rotate the light in 3D space.

Unlike a targeted directional light, a Free Direct light has no target object. You can move and rotate the light object to aim it in any direction.

A Free Direct light is used when you select a Standard sun in your *Daylight system (page 1–418)*.

Because directional rays are parallel, directional lights have a beam in the shape of a circular or rectangular prism instead of a "cone."

Note: Direct lights are supported in a *radiosity solution (page 3–51)* only if they are pointed downwards, outside the boundary box of the scene geometry.

Note: The mental ray renderer assumes that all directional lights come from infinity, so objects that are behind the direct light object in the 3ds Max scene will also be illuminated. In addition, with the mental ray renderer, directional lights cannot generate area shadows, and do not work with the Beam shader (in the lume library).

See also

Lights (page 2–1272)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

Intensity/Color/Attenuation Parameters (page 2–1345)

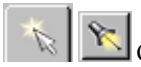
Advanced Effects Rollout (page 2–1341)

Shadow Parameters (page 2–1337)

Directional Parameters (page 2–1348)

Procedures

To create a free direct light:

1.  On the Create panel, click Lights.

   Standard is the default choice of light type.
2. In the Object Type rollout, click Free Direct.
3. Click a viewport.

   The light is now part of the scene. It points away from you in the viewport you clicked.
4. Set the creation parameters.

   To adjust the light's direction, you can rotate it as you would any object.

To change a viewport to a Light view:

1. Right-click the viewport label.

   The viewport right-click menu is displayed.
2. Choose Views.

   The name of each spotlight or directional light appears in the Views list.
3. Choose the name of the light you want.

   The viewport now shows the light's point of view. You can use the *Light Viewport Controls (page 3–750)* to adjust the light.

   The default keyboard shortcut for switching to a Light viewport is **$**.

Interface

Directional Parameters rollout

When you create a Free Direct light, the Targeted parameter is adjustable on the *General Parameters rollout (page 2–1331)*. This is a fixed value for target lights.

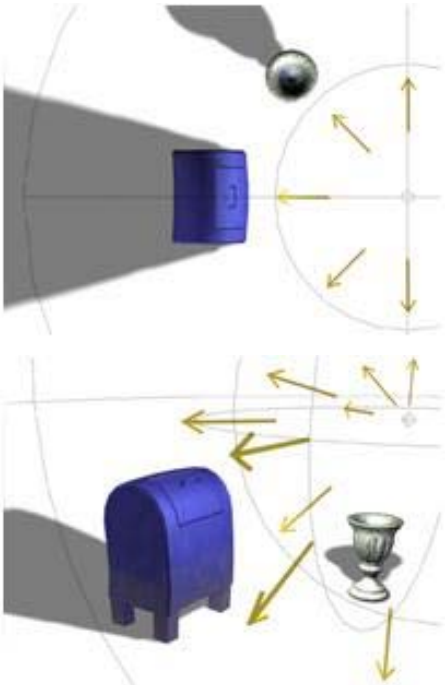
Targeted—When on, 3ds Max sets a point to use as an invisible target about which the Free Direct light can orbit. The spinner adjusts the distance to the target.

This parameter also affects the length of the light's cone display.

Omni Light

Create panel > Lights > Standard > Omni button

Create menu > Standard Lights > Omni Light



Top: Top view of an omni light

Bottom: Perspective view of the same light

An Omni light casts rays in all directions from a single source. Omni lights are useful for adding "fill lighting" to your scene, or simulating point source lights.

Omni lights can cast shadows and projections. A single shadow-casting omni light is the equivalent of six shadow-casting spotlights, pointing outward from the center.

When you set a map projected by an Omni light to be projected using the Spherical, Cylindrical, or Shrink Wrap Environment coordinates, the map is projected in the same way as it would be mapped to the environment. When you use the Screen Environment coordinates or Explicit Map Channel Texture coordinates, six copies of the map are projected radially.

Tip: Omni lights can generate up to six *quadtrees (page 3–999)*, so they generate ray-traced shadows more slowly than spotlights. Avoid using ray-traced shadows with omni lights unless your scene requires this.

See also

Lights (page 2–1272)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

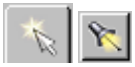
Intensity/Color/Attenuation Parameters (page 2–1345)

Advanced Effects Rollout (page 2–1341)

Shadow Parameters (page 2–1337)

Procedure

To create an omni light:

1.  On the Create panel, click Lights.

   Standard is the default choice of light type.
2. In the Object Type rollout, click Omni.
3. Click the viewport location where you want the light to be. If you drag the mouse, you can move the light around before releasing the mouse to fix its position.

   The light is now part of the scene.
4. Set the creation parameters.

   To adjust the light's effect, you can move it as you would any object.

Skylight

Create panel > Lights > Standard > Skylight button

Create menu > Lights > Skylight

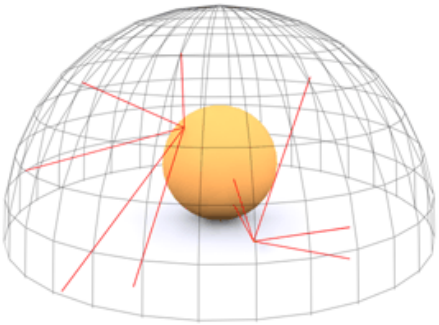
The Skylight light models daylight. It is meant for use with the *Light Tracer (page 3–44)*. You can set the color of the sky or assign it a map. The sky is modeled as a dome above the scene.



Model rendered with a single skylight, and light tracing

When you render with the *default scanline renderer (page 3–38)*, Skylight works best with advanced lighting: either the Light Tracer, or *radiosity (page 3–51)*.

Warning: When you render with the *mental ray renderer (page 3–78)*, objects illuminated by a Skylight appear dark unless you turn on *Final Gathering (page 3–940)*. The toggle for Final Gathering is on the *Final Gather rollout (page 3–111)* of the Render Scene dialog.



A skylight is modeled as a dome above the scene.

Tip: There are several ways to model daylight in 3ds Max, but if you use the Light Tracer, a Skylight often gives the best results.

Tip: If you encounter visual anomalies when rendering a bump-mapped material with a Skylight, convert the material to an *Advanced Lighting Override material (page 2–1601)* and then reduce the Indirect Light Bump Scale value.

Using a Map with the Skylight

If you use a map with a Skylight, the following guidelines can improve its effect:

- Make sure that the mapping coordinates are spherical or cylindrical.
- For light tracing, make sure you use sufficient samples. A good rule of thumb is to use at least 1,000 samples: set Initial Sample Spacing to 8x8 or 4x4, and increase the value of Filter Size to **2.0**.
- Use an image-processing application to blur the map before you use it. With a blurred map, you can use fewer samples to obtain good results. When used with Skylight, a blurred map will still render well.

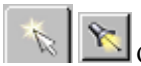
Be aware that using sufficient samples with a mapped Skylight will take longer to render than if the light were not mapped.

Using Render Elements with a Skylight

If you use *Render Elements (page 3–130)* to output the *lighting element (page 3–141)* of a skylight in a scene using either radiosity or the light tracer, you cannot separate the direct, indirect, and shadow channels of the light. All three elements of the skylight lighting are output to the Indirect Light channel.

Procedures

To create a Skylight:

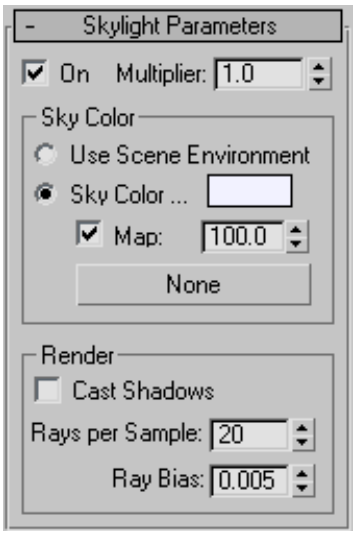
1.  On the Create panel, click Lights.

   Standard is the default choice of light type.
2. In the Object Type rollout, click Skylight.
3. Click a viewport.

   The light is now part of the scene.

   Note: The position of the Skylight, and its distance from objects, has no effect. The Skylight object is simply a helper. Skylight always comes from "overhead."
4. Set the creation parameters.

Interface



On—Turns the light on and off. When On is on, shading and rendering use the light to illuminate the scene. When off, the light is not used in shading or rendering. Default=on.

Multiplier—Amplifies the power of the light by a positive or negative amount. For example, if you set the *multiplier (page 3–977)* to 2, the light will be twice as bright. Default=1.0.

Using this parameter to increase intensity can cause colors to appear "burned out." It can also generate colors not usable in videos. In general, leave Multiplier set to its default of 1.0 except for special effects and special cases.

Sky Color group

Use Scene Environment—Colors the light using the environment set up on the *Environment panel (page 3–272)*.

This setting has no effect unless light tracing is active.

Sky Color—Click the color swatch to display a *Color Selector (page 1–161)* and choose a tint for the Skylight.



Map controls—These let you use a map to affect Skylight color. The button assigns a map, the toggle sets whether the map is active, and the spinner sets the percentage of the map to use (when the value is less than 100%, map colors are mixed with the Sky Color).

Tip: For best results, use an *HDR File (page 3–613)* for illumination.

The map has no effect unless light tracing is active.

Render group

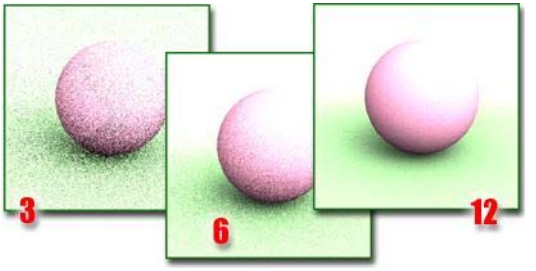
Note: If the renderer is not set to Default Scanline, or if the Light Tracer is active, these controls are disabled.

Cast Shadows—Causes the skylight to cast shadows. Default=off.

Note: The Cast Shadows toggle has no effect when using radiosity or the light tracer.

Note: Skylight objects will *not* cast shadows in an *ActiveShade rendering (page 3–17)*.

Rays per Sample—The number of rays used to calculate skylight falling on a given point in the scene. For animation, you should set this to a high value to eliminate flickering. A value of around 30 should eliminate flickering.



Increasing the number of rays increases the quality of your image. However, it also increases rendering time.

Ray Bias—The closest distance at which objects can cast shadows on a given point in the scene. Setting this value to 0 can cause the point to cast shadows upon itself, and setting it to a large value can prevent objects close to a point from casting shadows on the point.

mr Area Omni Light

Create panel > Lights > mr Area Omni button > Area Light Parameters rollout

When you render a scene using the *mental ray renderer (page 3–78)*, an area omni light emits light from a spherical or cylindrical volume, rather than from a point source. With the default scanline renderer, the area omni light behaves like any other standard omni light.

Note: In 3ds Max, area omni lights are created and supported by a MAXScript script. Only the mental ray renderer uses the parameters on the Area Light Parameters rollout. See *Enhancements to Standard Features (page 3–84)* for more details.

Tip: Area lights take longer to render than point lights. To create a quick test (or draft) rendering, you can use the Area/Linear Lights as Point Lights toggle in the *Common Parameters rollout (page 3–27)* of the Render Scene dialog to speed up your rendering.

See also

mr Area Spotlight (page 2–1299)

Procedures

To create an area omni light:

1.  On the Create panel, click Lights.
2. On the Object Type rollout, click mr Area Omni.
3. Click in a viewport.
4. Set the shape and size of the area light in the Area Light Parameters rollout.

   While you use the spinners to adjust the size of the area light, a gizmo (yellow by default) appears in viewports to show the adjusted size. This gizmo disappears once you finish adjusting the value.

   Tip: You can use Rotate to adjust the orientation of a cylindrical area omni light. However, no gizmo appears while you rotate the light.

To convert a standard 3ds Max light to an area light:

1. Select one or more lights.
2. Go to the Utilities panel.

3. On the Utilities rollout, click MAXScript.

   The MAXScript rollout is displayed.

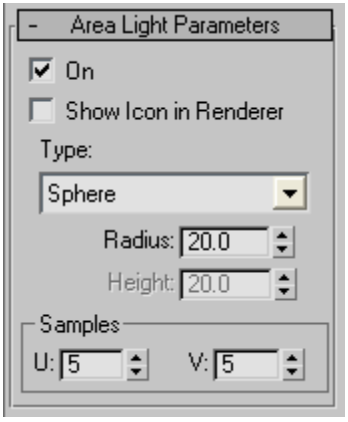
4. On the MAXScript rollout, choose "Convert To mr Area Lights" from the Utilities drop-down list.

   The "Convert To mr Area Lights" rollout is displayed.

5. On the "Convert To mr Area Lights" rollout, click Convert Selected Lights. A MAXScript alert is displayed, that says "Delete Old Lights?" Click Yes to delete the original light and replace it with the area light. Click No to leave the original light in place. If you click No there are now two lights in the scene: the original light, and the area light based on it.

6. Click Close to dismiss the "Convert To mr Area Lights" and MAXScript rollouts.

Interface



On—Turns the area light on and off. When On is on, the mental ray renderer uses the light to illuminate the scene. When On is off, the mental ray renderer doesn't use the light. Default=on.

Show Icon in Renderer—When on, the mental ray renderer renders a dark shape at the light's location. When off, the area light doesn't render. Default=off.

Type—Changes the shape of the area light. The choices are Sphere, for a spherical volume, and Cylinder, for a cylindrical volume. Default=Sphere.

Tip: You can use Rotate to adjust the orientation of a cylindrical area omni light. However, no gizmo appears while you rotate the light.

Radius—Sets the radius of the sphere or cylinder, in 3ds Max units. Default=20.0.

Height—Available only when Cylinder is the active type of area light. Sets the height of the cylinder, in 3ds Max units. Default=20.0.

As you use these spinners to adjust the size of the area light, a gizmo appears in viewports to show the adjusted size. This gizmo disappears once you finish adjusting the value.

Samples group

U and V—Adjust the quality of shadows cast by the area light. These values specify how many samples to take within the light's area. Higher values can improve rendering quality at a cost of rendering time. For a spherical light, U specifies the number of subdivisions along the radius, and V specifies the number of angular subdivisions. For a cylindrical light, U specifies the number of sampled subdivisions along the height, and V specifies the number of angular subdivisions. Default=5 for both U and V.

mr Area Spotlight

Create panel > Lights > mr Area Spot button > Area Light Parameters rollout

When you render a scene using the *mental ray renderer (page 3–78)*, an area spotlight emits light from a rectangular or disc-shaped area, rather than from a point source. With the default scanline renderer, the area spotlight behaves like any other standard spotlight.

Note: In 3ds Max, area spotlights are created and supported by a MAXScript script. Only the mental ray renderer uses the parameters in the Area Light Parameters rollout. See *Enhancements to Standard Features (page 3–84)* for more details.

Tip: Area lights take longer to render than point lights. To create a quick test (or draft) rendering, you can use the Area/Linear Lights as Point Lights toggle on the *Common Parameters rollout (page 3–27)* of the Render Scene dialog to speed up your rendering.

See also

mr Area Omni Light (page 2–1298)

Procedures

To create an area spotlight:

1.  Go to the Create panel and click Lights.
2. On the Object Type rollout, click mr Area Spot.
3. Drag in a viewport. The initial point of the drag is the location of the light, and the point where you release the mouse is the location of the target.

   The mental ray renderer will ignore the spotlight cone, but the location of the spotlight target determines the orientation of the plane of the area light, and the direction in which it's projected.
4. Set the shape and size of the area light in the Area Light Parameters rollout.

   While you use the spinners to adjust the size of the area light, a gizmo (yellow by default) appears in viewports to show the adjusted size. This gizmo disappears once you finish adjusting the value.

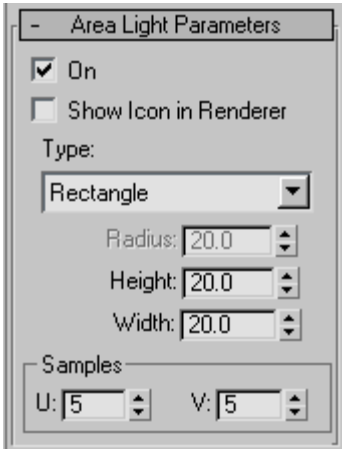
To convert a standard 3ds Max light to an area light:

1. Select one or more lights.
2. Go to the Utilities panel.
3. On the Utilities rollout, click MAXScript.

   The MAXScript rollout is displayed.
4. On the MAXScript rollout, choose Convert To Area Lights from the Utilities drop-down list.

   The Convert To Area Lights rollout is displayed.
5. On the Convert To Area Lights rollout, click Convert Selected Lights. A MAXScript alert is displayed, that says "Delete Converted Lights?" Click Yes to delete the original light and replace it with the area light. Click No to leave the original light in place. If you click No there are now two lights in the scene: the original light, and the area light based on it.
6. Click Close to dismiss the Convert To Area Lights and MAXScript rollouts.

Interface



On—Turns the area light on and off. When On is on, the mental ray renderer uses the light to illuminate the scene. When On is off, the mental ray renderer doesn't use the light. Default=on.

Show Icon in Renderer—When on, the mental ray renderer renders a dark shape where the area light is. When off, the area light is invisible. Default=off.

Type—Changes the shape of the area light. Can be either Rectangle, for a rectangular area, or Disc, for a circular area. Default=Rectangle.

Radius—Available only when Disc is the active type of area light. Sets the radius of the circular light area, in 3ds Max units. Default=20.0.

Height and Width—Available only when Rectangle is the active type of area light. Set the height and width of the rectangular light area, in 3ds Max units. Default=20.0 for both Height and Width.

While you use these spinners to adjust the size of the area light, a gizmo (yellow by default) appears in viewports to show the adjusted size. This gizmo disappears once you finish adjusting the value.

Samples group

U and V—Adjust the quality of shadows cast by the area light. These values specify how many samples to take within the light's area. Higher values can improve rendering quality at a cost of rendering time. For a rectangular light, U specifies the number of sampled subdivisions in one local dimension, and V the number of subdivisions in the other local dimension. For a circular (disc) light, U specifies the number of subdivisions along the radius, and V specifies the number of angular subdivisions. Default=5 for both U and V.

Photometric Lights

Photometric lights use *photometric (page 3–993)* (light energy) values that enable you to more accurately define lights as they would be in the real world. You can create lights with various distribution and color characteristics, or import specific photometric files available from lighting manufacturers.

Note: Photometric lights *always* attenuate using an inverse-square falloff, and rely on your scene using realistic units.

3ds Max includes the following types of photometric light objects:

Target Point Light (Photometric) (page 2–1303)

Free Point Light (Photometric) (page 2–1304)

Target Linear Light (Photometric) (page 2–1305)

Free Linear Light (Photometric) (page 2–1307)

Target Area Light (Photometric) (page 2–1307)

Free Area Light (Photometric) (page 2–1309)

IES Sun Light (Photometric) (page 2–1309)

IES Sky Light (Photometric) (page 2–1312)

mr Sky (page 2–1318)

mr Sun (page 2–1319)

Note: The remainder of this topic discusses photometric lights with the exception of mr Sky and mr Sun. For information about these light types, see *mental ray Sun & Sky (page 2–1313)*.

Distribution Capabilities of Photometric Lights

Each type of photometric light supports two or three different light distribution options.

The point light sources (target and free) support these distributions:

- *Isotropic (page 2–1323)*
- *Spotlight (page 2–1324)*
- *Web (page 2–1325)*

The linear and area light sources (target and free) support these distributions:

- *Diffuse (page 2–1323)*
- *Web (page 2–1325)*

Parameters for Photometric Lights

Most of the photometric light parameters are common to all photometric light types, and are described in the following topics:

Intensity/Color/Distribution Rollout (page 2–1352)

Photometric Linear and Area Lights have rollouts specific to them:

Linear Light Parameters Rollout (page 2–1354)

Area Light Parameters Rollout (page 2–1354)

Area Light Sampling Rollout (page 2–1354)

The remaining photometric light parameters are shared with standard lights, and are described in the following topics:

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

Shadow Parameters (page 2–1337)

Advanced Effects Rollout (page 2–1341)

Other photometric-specific topics include:

IES Standard File Format (page 2–1328)

Example of Photometric Data File (page 2–1329)

Common Lamp Values for Photometric Lights (page 2–1329)

Notes

- A scene's lighting can also be affected by the Ambient Light setting on the *Environment panel (page 3–272)*.
- You can use the *Place Highlight (page 1–467)* command to change a light's position.

Photometric Lights: Preset Lights

Create menu > Lights > Photometric Lights > Presets

Preset Lights are photometric lights that have a preset configuration in 3ds Max. Preset Lights load up correct values based on their equivalent commercial product. For example, a 60W light setting loads up the correct intensity, lighting, and distribution settings corresponding to a real 60W bulb. Using Preset Lights, you're able to quickly create lights that are correct simulations of real-world objects.

The Preset Lights available are:

- Generic 60W Bulb
- Generic 75W Bulb
- Generic 100W Bulb
- Generic Halogen Spotlight
- Generic Recessed 75W Lamp (web)
- Generic Recessed 75W Wallwash (web)
- Generic Recessed 250W Wallwash (web)
- Generic 4ft Pendant Fluorescent (web)
- Generic 4ft Cove Fluorescent (web)
- Generic Street 400W Lamp (web)
- Generic Stadium 100W Lamp (web)

See also

Photometric Lights (page 2–1301)

Target Point Light (Photometric) (page 2–1303)

Target Linear Light (Photometric) (page 2–1305)

Target Area Light (Photometric) (page 2–1307)

Free Point Light (Photometric) (page 2–1304)

Free Linear Light (Photometric) (page 2–1307)

Free Area Light (Photometric) (page 2–1309)

Procedure

The procedures for adding any of the 11 Preset Lights are identical. Follow the same procedure for any of the other Preset Lights as shown for the following:

To add a photometric preset generic 60W bulb to a scene:

1. From the Create menu, choose Photometric Lights > Preset Lights > Generic 60W Bulb.
2. Position your cursor where you want and click to place the light.

Interface

For each of the Preset Lights 3ds Max displays the appropriate parameters in the Command Panel. See the following interface descriptions for each Photometric Light object type:

Target Point Light (Photometric) (page 2–1303)

Target Linear Light (Photometric) (page 2–1305)

Target Area Light (Photometric) (page 2–1307)

Free Point Light (Photometric) (page 2–1304)

Free Linear Light (Photometric) (page 2–1307)

Free Area Light (Photometric) (page 2–1309)

Target Point Light (Photometric)

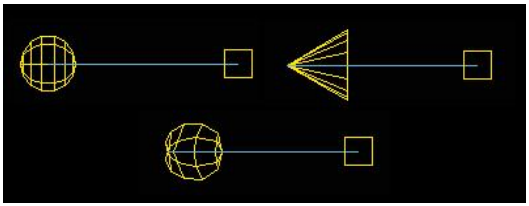
Create panel > Lights > Photometric Lights > Target Point button

Create menu > Lights > Target Point Light

A Target Point light, like a standard omni light, emits light from a geometric point. You can set the light distribution; this light has three types of distribution, with corresponding icons. You use a target object to aim the light.

Note: When you add a Target Point light, 3ds Max automatically assigns a *Look At controller (page 2–344)* to it, with the light's target object assigned as the Look At target. You can use the controller settings on the Motion panel to assign any other object in the scene as the Look At target.

Note: When you rename a Target Point light, the target is automatically renamed to match. For example, renaming *Point01* to *Klieg* causes *Point01.Target* to become *Klieg.Target*. The target's name must have the extension *.Target*. Renaming the target object does not rename the light object.



Icons for Target Point light with isotropic, spotlight and web distribution

See also

Lights (page 2–1272)

Isotropic Light Distribution (Photometric Lights) (page 2–1323)

Spotlight Distribution (Photometric Lights) (page 2–1324)

Web Distribution (Photometric Lights) (page 2–1325)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

Intensity/Color/Distribution Rollout (page 2–1352)

Shadow Parameters (page 2–1337)

Advanced Effects Rollout (page 2–1341)

Procedures

To create a Target Point light:

1.  On the Create panel, click Lights.

2. Choose Photometric from the drop-down list. (The default is Standard.)
3. In the Object Type rollout, click Target Point.
4. Drag in a viewport. The initial point of the drag is the location of the light, and the point where you release the mouse is the location of the target.

   The light is now part of the scene.
5. Set the creation parameters.

   You can use the Move transform to adjust the light.

To select the target:

The target, displayed as a small square, is often in the same area as objects that you want to illuminate. It can be difficult to select it by clicking.

1. First select the target point light.
2. Right-click the light and choose Select Target from the *quad menu (page 3–694)*.

   You can also choose Lights from the Selection Filters list on the toolbar, and then click the target. Clicking the line that connects the light and its target selects both objects.

To adjust the light and target:

1. Select the light or target or both.
2.  On the Main toolbar, click Move. Drag the selection to adjust the light.

   Because the light is always aimed at its target, you can't rotate it about its local X or Y axes. However, you can select and move the target object as well as the light itself. When you move either the light or the target, the light's orientation changes so it always points at the target.

   You can use the *Place Highlight (page 1–467)* command to change a light's position.

   For target point lights with spotlight distributions, you can also adjust the light using a *Light viewport (page 3–750)*.

To change a viewport to a light view:

Note: The viewport can only be set to a light view when the target point light's distribution is spotlight.

1. Right-click the viewport label.

   The viewport right-click menu is displayed.
2. Choose Views.

   The name of each light is displayed in the Views list. By default, Target Point lights are named Point01, Point02, and so on.
3. Choose the name of the light you want.

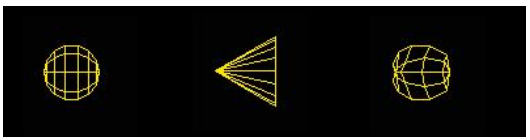
   The viewport now shows the light's point of view. You can use the *Light viewport (page 3–750)* controls to adjust the light.

   Tip: The default keyboard shortcut for Light viewports is **$**.

Free Point Light (Photometric)

Create panel > Lights > Photometric Lights > Free Point button
Create menu > Lights > Free Point Light

A Free Point light, like a standard omni light, emits light from a geometric point. You can set the light distribution; this light has three types of distribution, with corresponding icons. A free point light has no target object. You use transforms to aim the light.



Icons for Free Point light with isotropic, spotlight and web distribution

See also

Lights (page 2–1272)

Isotropic Light Distribution (Photometric Lights) (page 2–1323)

Spotlight Distribution (Photometric Lights) (page 2–1324)

Web Distribution (Photometric Lights) (page 2–1325)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

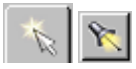
Intensity/Color/Distribution Rollout (page 2–1352)

Shadow Parameters (page 2–1337)

Advanced Effects Rollout (page 2–1341)

Procedures

To create a Free Point light:

1.  On the Create panel, click Lights.
2. Choose Photometric from the drop-down list. (Standard is the default.)
3. In the Object Type rollout, click Free Point.
4. Click the viewport location where you want the light to be.

   The light is now part of the scene. It points away from you in the viewport you clicked.
5. Set the creation parameters.

   You can *position the light (page 2–1282)* and adjust its direction with the transform tools or by using a Light viewport. You can also adjust the light's position with the *Place Highlights (page 1–467)* command.

To adjust the light:

1. Select the light.
2.  On the Main toolbar, click Move. Drag the selection to adjust the light.
3.  On the Main toolbar, click Rotate. Drag the selection to adjust the light.

   Tip: You can also adjust the light's position with the *Place Highlight (page 1–467)* command.

To change a viewport to a light view:

Note: This is only available for lights with Spot distribution.

1. Right-click the viewport label.

   The viewport right-click menu is displayed.
2. Choose Views.

   The name of each light is displayed in the Views list. By default, Free Point lights are named FPoint01, FPoint02, and so on.
3. Choose the name of the light you want.

   The viewport now shows the light's point of view. You can use the *Light Viewport Controls (page 3–750)* to adjust the light.

   The default keyboard shortcut for Light viewports is **$**.

Target Linear Light (Photometric)

Create panel > Lights > Photometric Lights > Target Linear button

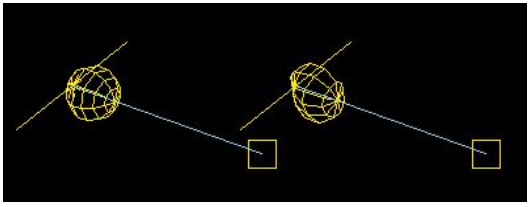
Create menu > Photometric Lights > Target Linear Light

A Target Linear light emits light from a line, like a fluorescent tube. You can set the light distribution; this light has two types of distribution, with corresponding icons. A target linear light uses a target object to aim the light.

Note: When you add a Target Linear light, the software automatically assigns a *Look At controller (page 2–344)* to it, with the light's target object

assigned as the Look At target. You can use the controller settings on the Motion panel to assign any other object in the scene as the Look At target.

Note: When you rename a Target Linear light, the target is automatically renamed to match. For example, renaming *Linear01* to *Klieg* causes *Linear01.Target* to become *Klieg.Target*. The target's name must have the extension *.Target*. Renaming the target object does not rename the light object.



Icons for Target Linear light with diffuse and web distribution

See also

Lights (page 2–1272)

Diffuse Distribution (Photometric Lights) (page 2–1323)

Web Distribution (Photometric Lights) (page 2–1325)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

Intensity/Color/Distribution Rollout (page 2–1352)

Shadow Parameters (page 2–1337)

Advanced Effects Rollout (page 2–1341)

Procedures

To create a Target Linear light:

1.  On the Create panel, click Lights.
2. Choose Photometric from the drop-down list. (The default is Standard.)
3. In the Object Type rollout, click Target Linear.
4. Drag in a viewport. The initial point of the drag is the location of the light, and the point where you release the mouse is the location of the target.

   The light is now part of the scene.
5. Set the creation parameters.

   You can use the Move transform to adjust the light.

To select the target:

The target, displayed as a small square, is often in the same area as objects that you want to illuminate. It can be difficult to select it by clicking.

1. First select the light itself.
2. Right-click the light to open the *quad menu (page 3–694)*. From the tools quadrant, choose Select Target.

   You can also choose Lights from the Selection Filters list on the toolbar, and then click the target. Clicking the line that connects the light and its target selects both objects.

To adjust the light and target:

1. Select the light or target or both.
2.  On the Main toolbar, click Move. Drag the selection to adjust the light.

   Because the light is always aimed at its target, you can't rotate it about its local X or Y axes. However, you can select and move the target object as well as the light itself. When you move either the light or the target, the light's orientation changes so it always points at the target.

   You can use the *Place Highlight (page 1–467)* command to change a light's position.

Free Linear Light (Photometric)

Create panel > Lights > Photometric Lights > Free Linear button

Create menu > Photometric Lights > Free Linear Light

A Free Linear light emits light from a line, like a fluorescent tube. You can set the light distribution; this light has two types of distribution, with corresponding icons. A free linear light has no target object. You use transforms to aim the light.



Icons for Free Linear light with diffuse and web distribution

See also

Lights (page 2–1272)

Diffuse Distribution (Photometric Lights) (page 2–1323)

Web Distribution (Photometric Lights) (page 2–1325)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

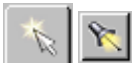
Intensity/Color/Distribution Rollout (page 2–1352)

Shadow Parameters (page 2–1337)

Advanced Effects Rollout (page 2–1341)

Procedures

To create a Free Linear light:

1.  On the Create panel, click Lights.
2. Choose Photometric from the drop-down list. (Standard is the default.)
3. In the Object Type rollout, click Free Linear.
4. Click the viewport location where you want the light to be.

   The light is now part of the scene. It points away from you in the viewport you clicked.
5. Set the creation parameters.

   You can *position the light (page 2–1282)* and adjust its direction with the transform tools. You can also adjust the light's position with the *Place Highlight (page 1–467)* command.

To adjust the light:

1. Select the light.
2.  On the Main toolbar, click Move. Drag the selection to position the light.
3.  On the Main toolbar, click Rotate. Drag the selection to adjust the light.

   Tip: You can also adjust the light's position with the *Place Highlight (page 1–467)* command.

Target Area Light (Photometric)

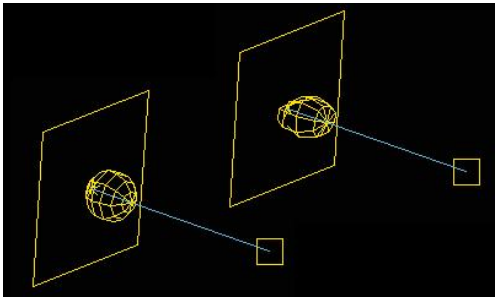
Create panel > Lights > Photometric Lights > Target Area button

Create menu > Photometric Lights > Target Area Light

A Target Area light emits light from a rectangular area, like a skylight. You can set the light distribution; this light has two types of distribution, with corresponding icons. A target area light uses a target object to aim the light.

Note: When you add a Target Area light, the software automatically assigns a *Look At controller (page 2–344)* to it, with the light's target object assigned as the Look At target. You can use the controller settings on the Motion panel to assign any other object in the scene as the Look At target.

Note: When you rename a Target Area light, the target is automatically renamed to match. For example, renaming *Area01* to *Klieg* causes *Area01.Target* to become *Klieg.Target*. The target's name must have the extension *.Target*. Renaming the target object does not rename the light object.



Icons for Target Area light with diffuse and web distribution

See also

Lights (page 2–1272)

Diffuse Distribution (Photometric Lights) (page 2–1323)

Web Distribution (Photometric Lights) (page 2–1325)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

Intensity/Color/Distribution Rollout (page 2–1352)

Shadow Parameters (page 2–1337)

Advanced Effects Rollout (page 2–1341)

Procedures

To create a Target Area light:

1.  On the Create panel, click Lights.
2. Choose Photometric from the drop-down list. (The default is Standard.)
3. In the Object Type rollout, click Target Area.
4. Drag in a viewport. The initial point of the drag is the location of the light, and the point where you release the mouse is the location of the target.

   The light is now part of the scene.
5. Set the creation parameters.

   You can use the Move transform to adjust the light.

To select the target:

The target, displayed as a small square, is often in the same area as objects that you want to illuminate. It can be difficult to select it by clicking.

1. First select the light itself.
2. Right-click the light to open the *quad menu (page 3–694)*. In the Tools quadrant, choose Select Target from the pop-up menu.

   You can also choose Lights from the Selection Filters list on the toolbar, and then click the target. Clicking the line that connects the light and its target selects both objects.

To adjust the light and target:

1. Select the light or target or both.
2.  On the Main toolbar, click Move. Drag the selection to adjust the light.

   Because the light is always aimed at its target, you can't rotate it about its local X or Y axes. However, you can select and move the target object as well as the light itself. When you move either the light or the target, the light's orientation changes so it always points at the target.

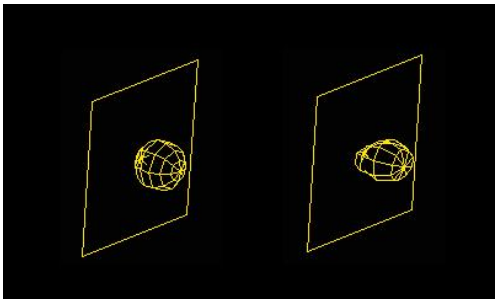
   You can use the *Place Highlight (page 1–467)* command to change a light's position.

Free Area Light (Photometric)

Create panel > Lights > Photometric Lights > Free Area button

Create menu > Photometric Lights > Free Area Light

A Free Area light emits light from a rectangular area, like a skylight. You can set the light distribution; this light has two types of distribution, with corresponding icons. A free area light has no target object. You use transforms to aim the light.



Icons for Free Area light with diffuse and web distribution

See also

Lights (page 2–1272)

Diffuse Distribution (Photometric Lights) (page 2–1323)

Web Distribution (Photometric Lights) (page 2–1325)

Name and Color Rollout (Lights) (page 2–1273)

General Lighting Parameters (page 2–1331)

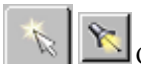
Intensity/Color/Distribution Rollout (page 2–1352)

Shadow Parameters (page 2–1337)

Advanced Effects Rollout (page 2–1341)

Procedures

To create a Free Area light:

1.  On the Create panel, click Lights.
2. Choose Photometric from the drop-down list. (Standard is the default.)
3. In the Object Type rollout, click Free Area.
4. Click the viewport location where you want the light to be.

   The light is now part of the scene. It points away from you in the viewport you clicked.
5. Set the creation parameters.

   You can *position the light (page 2–1282)* and adjust its direction with the transform tools. You can also adjust the light's position with the *Place Highlight (page 1–467)* command.

To adjust the light:

1. Select the light.
2.  On the Main toolbar, click Rotate. Drag the selection to adjust the light.

   Tip: You can also adjust the light's position with the *Place Highlight (page 1–467)* command.

IES Sun Light (Photometric)

Create panel > Systems > Create a Daylight system in your scene. > Modify panel > Daylight Parameters rollout > Sunlight > Choose IES Sun.

Create panel > Lights > Photometric > IES Sun

Create menu > Lights > Photometric > IES Sun

IES Sun is a physically-based light object that simulates *sunlight (page 3–1018)*. When used in conjunction with a *daylight system (page 1–418)*, its values are set automatically based on geographic location, time, and date. (IES stands

for Illuminating Engineering Society; see *IES Standard File Format (page 2–1328)*.)



Outdoor scene illuminated by the IES Sun light

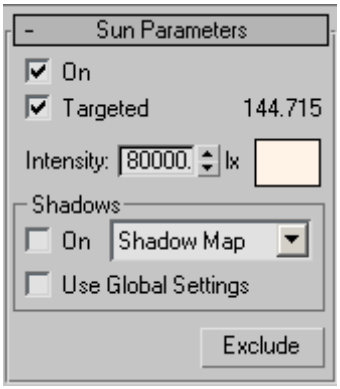
The mental ray renderer gives physically accurate results for IES Sun, and renderings that use it will appear similar to renderings done with the default scanline renderer. You do not need to turn on Final Gather for light from IES Sun to render.

Note: While you can create daylight lights manually, you should use them in conjunction with the Daylight system to get the best results. The Daylight system combines the two daylight components of sun and sky in a unified interface.

Tip: If you use the IES sun or sky with the *Logarithmic Exposure Control (page 3–297)*, turn on both the Daylight and Exterior options. This will provide greater control for properly mapping the higher energy levels to RGB colors. In addition, it is important to set the Physical Scale to the brightest light source in your scene. If the IES Sun is used, set the Physical Scale to 90000 cd.

Note: When you add an IES Sun light, 3ds Max automatically assigns a *Look At controller (page 2–344)* to it, with the light's target object assigned as the Look At target. You can use the controller settings on the Motion panel to assign any other object in the scene as the Look At target.

Interface



On—Turns sunlight on and off in the viewport.

Targeted—Applicable only when you add the IES Sun light directly, rather than as part of a *Daylight system (page 1–418)*. When on, the light is targeted and you can change the target distance by moving the target. The distance between the light and its target is displayed to the right of the check box. When off, you can set this value directly.

Cast Shadows—Sets whether the sunlight casts shadows or not.

Intensity—The intensity of the sunlight. The color swatch to the right of the spinner opens the *Color Selector (page 1–161)* to set the color of the light. Typical intensities in a clear sky are around 90,000 lux.

Note: Intensity is set automatically and cannot be set manually if the sun is under the control of a daylight system.

Shadows group

On—Determines whether the sunlight casts shadows or not.

Shadow Method drop-down list—Determines whether the renderer uses *shadow maps (page*

3–1010), *ray-traced shadows (page 3–1000)*, *advanced ray-traced shadows (page 3–1000)*, *mental ray shadow maps (page 2–1360)*, or *area shadows (page 3–911)* to generate shadows for this light.

Each shadow type has a rollout with its associated controls.

Tip: When you use the mental ray renderer and mental ray shadow maps, you can set up soft-edged shadows.

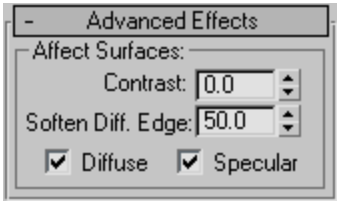
Use Global Settings—Turn on to use global settings for shadows cast by this light. Turn off to enable individual control of the shadows. If you choose not to use the global settings, you must choose which method the renderer will use to generate shadows for this particular light.

When Use Global Settings is on, the shadow parameters switch over to show you what the current global setting is. This data is shared by every other light of this class. When Use Global Settings is off, the shadow parameters are specific to that particular light.

Exclude—Excludes selected objects from the effects of the light. Click this button to display the *Exclude/Include dialog (page 2–1335)*.

Excluded objects still appear lit in shaded viewports. Exclusion takes effect only when you render the scene.

Advanced Effects Rollout



Contrast—Adjusts the contrast between the diffuse and ambient areas of the surface. Leave this set to 0 for normal contrast. Increase the value to increase the contrast for special effects: for example, the harsh light of outer space. Default=0.0.

Soften Diffuse Edge—Increasing the value of Soften Diffuse Edge softens the edge between the diffuse and ambient portions of a surface. This helps eliminate edges that can appear on a surface under certain circumstances. Default=50.

Note: Soften Diffuse Edge slightly reduces the intensity of the light. You can counter this, to some extent, by increasing the Multiplier value.

Diffuse—When on, the light affects the diffuse properties of an object's surface. When off, the light has no effect on the diffuse surface. Default=on.

Specular—When on, the light affects the specular properties of an object's surface. When off, the light has no effect on the specular properties. Default=on.

For example, by using the Diffuse and Specular check boxes you can have one light color the specular highlights of an object, while not coloring its diffuse component, and then have a second light color the diffuse component of the surface while not creating specular highlights.

Optimizations rollout

This rollout is identical to the *Optimizations rollout (page 2–1361)* for advanced ray-traced and area shadows.

IES Sky Light (Photometric)

Create panel > Systems > Create a Daylight system in your scene. > Modify panel > Daylight Parameters rollout > Skylight > IES Sky

Create panel > Lights > Photometric > IES Sky

Create menu > Lights > Photometric > IES Sky

IES Sky is a physically-based light object that simulates atmospheric effects on *skylight (page 3–1012)*.



Outdoor scene lit by the IES sky light

Note: (IES stands for Illuminating Engineering Society; see *IES Standard File Format (page 2–1328)*.)

You can create daylight lights manually, but to get the best results, you should use them in conjunction with the *Daylight system (page 1–418)*. The Daylight system combines the two daylight components of sun and sky in a unified interface. It allows you to set date and time positions with the light type you want to use.

IES Sky works correctly only when the sky object is pointing down from the Z axis, meaning that it points down when looking from the Top view.

When using the default scanline renderer, IES Sky produces the best results when used in conjunction with one of the *advanced lighting options (page 3–44)*: radiosity or light tracing.

Warning: When you render with the *mental ray renderer (page 3–78)*, objects illuminated by IES Sky appear dark unless you turn on *Final Gathering (page 3–940)*. The toggle for Final Gathering is on the *Final Gather rollout (page 3–111)* of the Render Scene dialog.

When you add an IES Sky light, 3ds Max automatically assigns a *Look At controller (page 2–344)* to it, with the light's target object assigned as the Look At target. You can use the controller settings on the Motion panel to assign any other object in the scene as the Look At target.

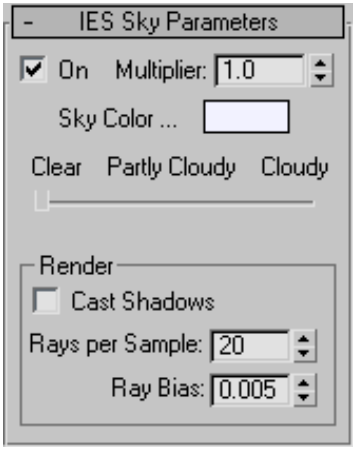
Using Render Elements with an IES Sky Light

If you use *Render Elements (page 3–130)* to output the *lighting element (page 3–141)* of an IES Sky light in a scene using either radiosity or the light tracer, you cannot separate the direct, indirect, and shadow channels of the light. All three elements of the IES Sky lighting are output to the Indirect Light channel.

See also

IES Sun Light (Photometric) (page 2–1309)

Interface



On—Turns the sky light on and off in the viewport.

Multiplier—Adjusts the intensity of the skylight.

When this is set to 1.0, the intensity will be physically accurate based on angle. You can, however, override this by changing the value. This is useful for doing night shots with artificial lighting.

Sky Color—The color swatch opens the *Color Selector (page 1–161)*, which lets you set the color of the sky.

Coverage group

Clear, Partly Cloudy, Cloudy—This choice determines the extent to which light is scattered through the sky. Use the slider to set the control to one of the named settings or anywhere in between.

Render group

Note: These controls are unavailable when the mental ray renderer is active.

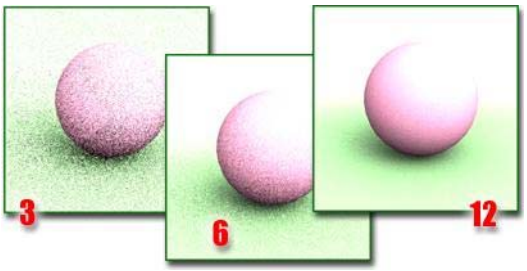
Cast Shadows—Causes the sky light to cast shadows.

A great deal of processing is dedicated to the calculation of the subtle shadows cast by *skylight (page 3–1012)*. If shadows are not important in your model, you can disable them and save substantial amounts of processing time. However, the results will not be as realistic.

Note: The Cast Shadows toggle has no effect when using radiosity or the light tracer.

Note: IES Sky objects will *not* cast shadows in an *ActiveShade rendering (page 3–17)*.

Rays per Sample—The number of rays used to calculate skylight falling on a given point in the scene. For animation you should set this to a high value to eliminate flickering. A value of around 30 should eliminate flickering.



Increasing the number of rays increases the quality of your image. However, it also increases rendering time.

Ray Bias—The closest distance at which objects can cast shadows on a given point in the scene. Setting this value to 0 can cause the point to cast shadows upon itself, and setting it to a large value can prevent objects close to a point from casting shadows on the point.

mental ray Sun & Sky

Daylight system > Modify panel > Daylight Parameters rollout > Sunlight/Skylight drop-down lists > mr Sun/mr Sky

The mental ray Sun & Sky solution is designed to enable physically plausible daylight simulations and accurate renderings of daylight scenarios. In 3ds Max, this is achieved through the use of two

special photometric lights and an environment shader that all work together:

- The mr Sun photometric light is responsible for the sunlight; the direct light from the sun.



The scene is lit by mr Sun only.

- The mr Sky photometric light is responsible for the skylight; it simulates the real-world phenomenon of indirect light created by the scattering of sunlight in the atmosphere.



The scene is lit by mr Sky only.

- The mr Physical Sky environment shader is responsible for the visible representation of the sun disk and the sky, both to the camera and in reflections and refraction, as well as for the *virtual ground plane* (gray in the following illustration).



mr Physical Sky is visible, but no lighting is present in the scene.

While you can use these separately, their intended functionality is unleashed only when used together. mr Sun and mr Sky are meant primarily for use within the 3ds Max *Daylight system (page 1–418)*. When used in combination, this solution is called *Sun & Sky*.

Note: Sun & Sky is a true high-dynamic range photometric lighting system that requires the Logarithmic Exposure Control to be enabled. In some of the mental ray *rendering presets (page 3–23)*, such as those starting with "mental.ray.daylight," this is done automatically. Also, when you add a Daylight System from the Create menu, you have the option to activate Logarithmic Exposure Control automatically.

Using Sun & Sky with SSS Materials

To use a mental ray *fast SSS material (page 2–1583)* together with the high-dynamic range Sun & Sky solution, be sure to turn on Scatter Indirect Illumination on the material's Advanced Options rollout so that the material can scatter the skylight, which is considered indirect.

Also, turn off Screen (Soft) Compositing Of Layers, also on the Advanced Options rollout; otherwise the output of the SSS shaders is clamped to a low dynamic range and will appear to render black.

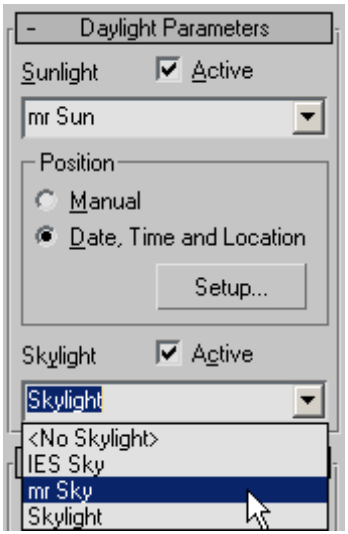
Procedure

To use mental ray Sun & Sky:

1. Make sure *mental ray (page 3–78)* is the active renderer.
2. From the Create menu, choose Lights > Daylight System.

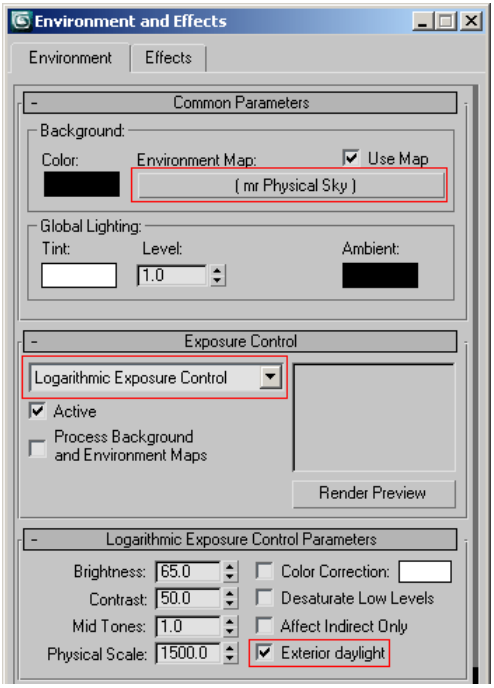
   You can also find Daylight System on the Create menu under Systems.

   Note: The Daylight System is also available from Create panel > Systems, but if you go that route you're not prompted to turn on Logarithmic Exposure control (see next step), so when using Sun & Sky it's best to make a habit of choosing Daylight System from the Create menu.
3. You're prompted to activate Logarithmic Exposure Control with the Exterior Daylight Flag, if both are not already active. Click Yes to do so automatically.
4. Create a Daylight system in the Perspective (or a Camera) viewport: Drag and release to create the compass rose, move the mouse to position the light, and then click to place the light and finish creating the system.
5.  Go to the Modify panel, and on the Daylight Parameters rollout, set Sunlight to mr Sun and Skylight to mr Sky.



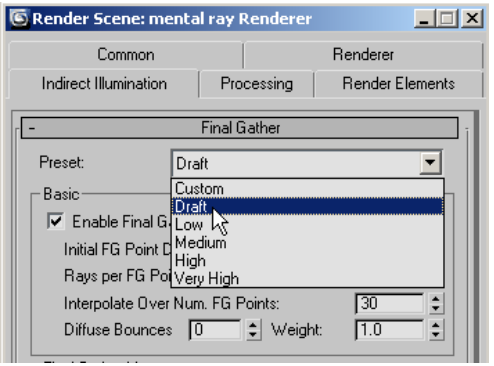
6. When you choose mr Sky, you're prompted to add a mr Physical Sky environment map. Click Yes to do so.

   The parameters rollouts for mr Sun and mr Sky now appear on the Modify panel.
7. Open the *Environment dialog (page 3–272)* (press 8) and make sure Common Parameters > Environment Map is set to mr Physical Sky, *Exposure Control (page 3–293)* is set to Logarithmic Exposure Control, and Logarithmic Exposure Control Parameters has Exterior Daylight on.



Tip: Should you wish to edit the mr Physical Sky parameters, an easy way to access them is to open the Material Editor (press M) and then drag the Environment Map button from the Environment panel > Common Parameters rollout to a material slot (sample sphere) in the Material Editor. When prompted for the copy method, choose Instance.

8. Open the *Render Scene dialog (page 3–2)* (F10) to the Indirect Illumination panel and turn on Enable Final Gather. Choose the Draft preset as a starting point.



Important: Don't forget to turn on Enable Final Gather when using Sun & Sky, since the skylight is a form of indirect light, which can be rendered only with the help of Final Gather. If this is not done, Sun shadows will be unnaturally dark or black rather than bluish, as they should be.

9. Render with Sun & Sky. To specify different times, dates, and locations in the Daylight system for different effects, use the Motion panel > Control Parameters settings or turn on Manual Override to place the sun by hand.

Tip: Adjust the Brightness and Contrast of the Logarithmic Exposure Control for a pleasing image.

Common Parameters

Certain parameters do the same things in mr Sun, mr Sky, and mr Physical Sky. For physical correctness, it is necessary to keep these parameters in sync with each other in all three elements. For example, a sun with a different Haze value than the sky cannot be guaranteed to be physically plausible.

For this reason the concept of parameter inheritance is included. Both the mr Sun light and the mr Physical Sky shader have check boxes labeled "Inherit from mr Sky" that are on by default. When these are on, the common

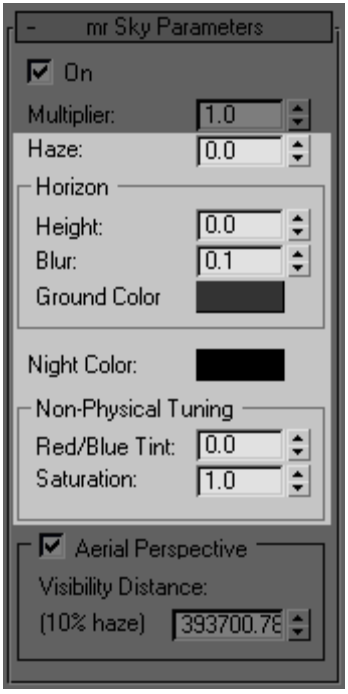
parameters are all guided by the mr Sky light, which is the central point of control.

With parameter inheritance, you can concentrate on tweaking the parameters in one spot (mr Sky) and, as long as Inherit From mr Sky is on for both mr Sun and mr Physical Sky, you're sure to obtain consistent results.

Interface (common parameters)

The most important common parameters are those that drive the entire shading and colorization model.

Note: mr Sun and mr Physical Sky can inherit parameters from mr Sky even if the latter is off.



Haze—The amount of water vapor or other particulate matter in the air. Possible values range from 0.0 (a completely clear day) to 15.0 (extremely overcast, or a sandstorm in the Sahara). The Haze value influences the intensity and color of the sky and horizon, the intensity and color of sunlight, the softness of the sun's shadows, the softness of the glow around the sun, and the strength of the aerial perspective.

Horizon Height—The vertical position of the horizon.

The default value, 0.0, places the horizon at a standard height. However, because the horizon is infinitely far away, this can cause trouble joining up with any finite geometry that is supposed to represent the ground. It can also cause issues rendering locations that are supposed to be at a high altitude, like mountain tops or the top of New York skyscrapers where the horizon really is visibly "below" the viewer.

This parameter allows tuning the position of the horizon. Note that this horizon doesn't actually exist at a specific height in 3D space; it is a shading effect for rays that go below a certain angle. This parameter tweaks that angle. The total range available range is somewhat extreme, reaching from -10.0 (the horizon is "straight down") to 10.0 (the horizon is at the zenith). In practice, only much smaller values are actually useful. For example, to push the horizon down just below the edge of a finite visible ground plane, use –0.2.

Note: The Horizon height affects not only the visual representation of the horizon in the mr Sky light, but also the color of the mr Sun itself. In other words, the point where the sun "sets" will change for a Horizon height settings other than 0.0.

Horizon Blur—The "blurriness" with which the horizon is rendered.

At 0.0 the horizon is completely sharp. Generally only values lower than 0.5 are useful, but the full range is up to 10.0 for a horizon that consists of blur only, with no actual horizon at all.

Ground Color—The color of the virtual ground plane. Note that this is a diffuse reflectance value (that is, albedo). The ground appears as a Lambertian reflector with this diffuse color, lit by

the sun and sky only, and does not receive any shadows.

Tip: Some sky models neglect the influence of bounce light from the ground, assuming only the sky is illuminating the scene. To compare the output of mr Sky with, for example, the IES Sky light, set Ground Color to black.

Night Color—The minimum color of the sky: The sky will never become darker than this value. It can be useful for adding things like moon, stars, high-altitude cirrus clouds that remain lit long after sunset, etc. As the sun sets and the sky darkens, the contribution from Night Color is unaffected and remains as the base light level.

Red/Blue Tint—Gives artistic control over the redness of the light. The default value of 0.0 is the physically correct value (calculated for a 6500K whitepoint), but can be changed with this parameter, which ranges from -1.0 (extremely blue) to 1.0 (extremely red).

Saturation—Also an artistic control, where 1.0 is the physically calculated saturation level. The parameter ranges from 0.0 (black and white) to 2.0 (extremely boosted saturation).

mr Sky

Daylight system > Modify panel > Daylight Parameters rollout > Skylight drop-down list > mr Sky
Create panel > Lights > Photometric > mr Sky
Create menu > Lights > Photometric Lights > mr Sky

The mr Sky light is intended primarily for use in the mental ray Sun & Sky combination. This topic mainly provides information on parameters unique to this component. A number of mr Sky parameters are common to all three Sun & Sky components. For some of those parameters, this topic provides a brief explanation plus a link to the main topic with additional details.

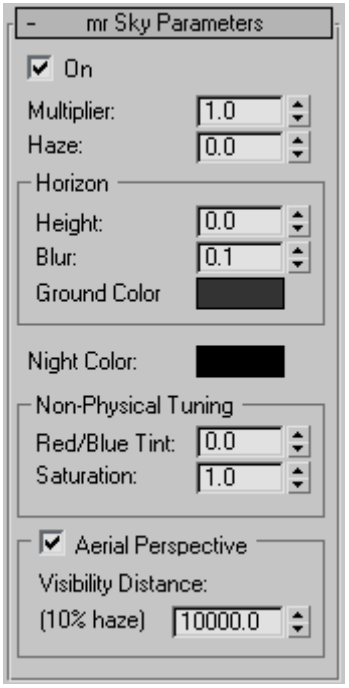
See also

mental ray Sun & Sky (page 2–1313)

mr Sun (page 2–1319)

mr Physical Sky (page 2–1321)

Interface



On—Turns the light on and off.

Multiplier—A scalar multiplier for the light output. Default=1.0.

Haze—The amount of particulate matter in the air. Possible values range from 0.0 (a completely clear day) to 15.0 (extremely overcast, or a sandstorm in the Sahara). Default=0.0.

For details, see *Haze*.

Horizon group

Height—The vertical position of the horizon. Default=0.0.

For details, see *Horizon Height*.

Blur—The "blurriness" with which the horizon is rendered. Default=0.1.

At 0.0 the horizon is completely sharp. Generally only values lower than 0.5 are useful, but the full range is up to 10.0 for a horizon that consists of blur only, with no actual horizon at all.

Ground Color—The color of the virtual ground plane. Note that this is a diffuse reflectance value (that is, albedo). The ground appears as a Lambertian reflector with this diffuse color, lit by the sun and sky only, and does not receive any shadows.

Tip: Some sky models neglect the influence of bounce light from the ground, assuming only the sky is illuminating the scene. To compare the output of mr Sky with, for example, the IES Sky model, one must therefore set Ground Color to black.

Night Color—The minimum color of the sky: The sky will never become darker than this value. It can be useful for adding things like moon, stars, high-altitude cirrus clouds that remain lit long after sunset, etc. As the sun sets and the sky darkens, the contribution from Night Color is unaffected and remains as the base light level.

Non-Physical Tuning group

Red/Blue Tint—Provides artistic control over the redness of the sky light. The default value of 0.0 is the physically correct value (calculated for a 6500K whitepoint), but can be changed with this parameter, which ranges from -1.0 (extremely blue) to 1.0 (extremely red).

Saturation—Provides artistic control over saturation of the sky light. The default value of 1.0 is the physically calculated saturation level. Possible values range from 0.0 (black and white) to 2.0 (extremely high saturation).

Aerial Perspective group

Visibility Distance—Aerial Perspective is a term used by painters to convey how distant objects are perceived as hazier and tinted towards the blue end of the spectrum. mr Sky emulates this with the Visibilty Distance parameter. When nonzero, it defines the "10% distance", that is, the distance at which approximately 10% of haze is visible at a Haze level of 0.0.

mr Sun

Daylight system > Modify panel > Daylight Parameters rollout > Sunlight drop-down list > mr Sun
Create panel > Lights > Photometric > mr Sun
Create menu > Lights > Photometric Lights > mr Sun

The mr Sun light is intended primarily for use in the mental ray Sun & Sky combination. This topic mainly provides information on parameters unique to this component. A number of mr Sky parameters are common to all three Sun & Sky components. For some of those parameters, this topic provides a brief explanation plus a link to the main topic with additional details.

See also

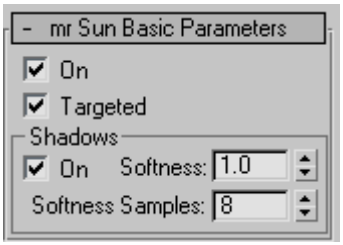
mental ray Sun & Sky (page 2–1313)

mr Sky (page 2–1318)

mr Physical Sky (page 2–1321)

Interface

mr Sun Basic Parameters rollout



On—Turns the mr Sun light on and off. Default=on.

Targeted—Applicable only when you add an mr Sun light directly to the scene via Create panel > Lights > Photometric, rather than as part of a *Daylight system (page 1–418)*. When on, the light is targeted and you can change the target distance by moving the target. The distance between the light and its target is displayed to the right of the check box. When off, you can set this value directly. Default=on.

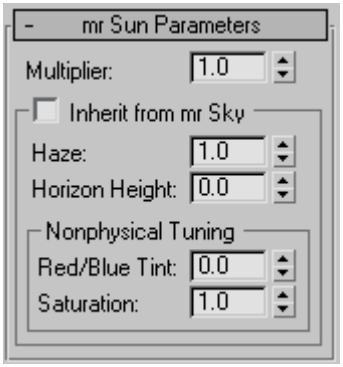
Shadows group

On—Toggles shadowing for the light. Default=on.

Softness—The softness of shadow edges. The default value of 1.0 accurately matches the softness of real solar shadows. Lower values make the shadows sharper and higher values make them softer.

Softness—The number of shadow samples for the soft shadows. If it is set to 0, no soft shadows are generated. Default=8.

mr Sun Parameters rollout



Multiplier—A scalar multiplier for the light output. Default=1.0.

Inherit from mr Sky—Uses the equivalent settings from the mr Sky Parameters rollout for the remaining mr Sun Parameters rollout settings. Default=on.

For further information, see *Common Parameters (page 2–1316)*

Haze—The amount of particulate matter in the air. Possible values range from 0.0 (a completely clear day) to 15.0 (extremely overcast, or a sandstorm in the Sahara). Default=0.0.

For details, see *Haze*.

Height—The vertical position of the horizon. Default=0.0.

For details, see *Horizon Height*.

Non-Physical Tuning group

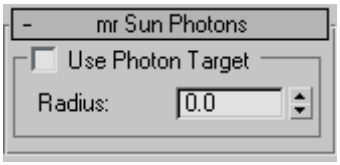
Red/Blue Tint—Provides artistic control over the redness of the sky light. The default value of 0.0 is the physically correct value (calculated for a 6500K whitepoint), but can be changed with this parameter, which ranges from -1.0 (extremely blue) to 1.0 (extremely red).

Saturation—Provides artistic control over saturation of the sky light. The default value of

1.0 is the physically calculated saturation level. Possible values range from 0.0 (black and white) to 2.0 (extremely high saturation).

mr Sun Photons rollout

Use these settings to focus global-illumination photons on an area of interest. For example, if you've modeled a huge city as a backdrop, but are rendering only a room interior, you probably don't want mental ray to shoot photons over the entire city, with the result that only a few will find their way into the room.



Use Photon Target—When on, uses the Radius setting with respect to the light target.

Radius—Sets the radial distance from the target in which the mr Sun casts GI photons.

mr Physical Sky

Daylight system > Modify panel > Daylight Parameters rollout > Skylight drop-down list > mr Sky > Confirm prompt.

The mr Physical Sky shader is intended primarily for use in the mental ray Sun & Sky combination. This topic mainly provides information on parameters unique to this component. A number of mr Physical Sky parameters are common to all three Sun & Sky components. For some of those parameters, this topic provides a brief explanation plus a link to the main topic with additional details.

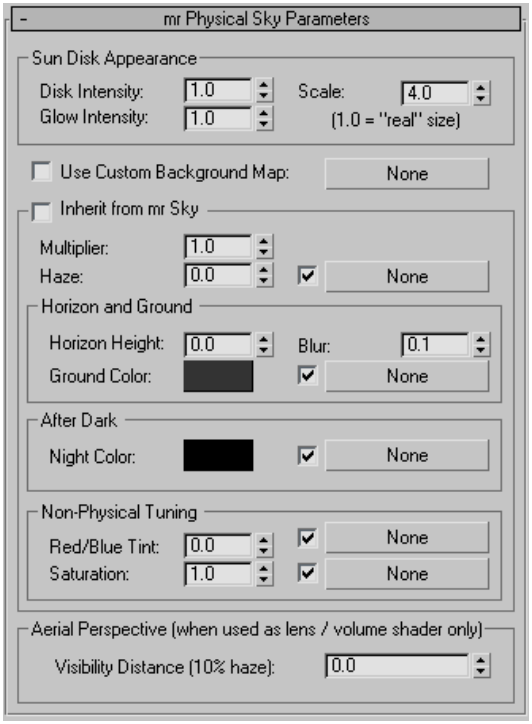
See also

mental ray Sun & Sky (page 2–1313)

mr Sky (page 2–1318)

mr Sun (page 2–1319)

Interface



Sun Disk Appearance group

Use these settings to adjust the visible appearance of the sun in the sky.

Disk Intensity—The brightness of the sun.

Glow Intensity—The brightness of the glow surrounding the sun.

Scale—The size of the sun disk.

Use Custom Background Map—When on but no background map is specified, the background of the rendering is transparent black, suitable for external compositing. If you supply a background shader by clicking the button and then specifying a map or shader, the background of the rendering will come from that shader (for example, a texture map that uses a background photograph). In

either case the mr Physical Sky will still be visible in reflections and refraction.

To apply a map or shader to this parameter, click the None button.

Note: The background area in the rendered image's alpha channel is always fully transparent, as is the case in general in 3ds Max.

Inherit from mr Sky—Uses the equivalent settings from the mr Sky Parameters rollout for the remaining mr Physical Sky Parameters rollout settings, except for Aerial Perspective. Default=on.

For further information, see *Common Parameters (page 2–1316)*

Haze—The amount of particulate matter in the air. Possible values range from 0.0 (a completely clear day) to 15.0 (extremely overcast, or a sandstorm in the Sahara). Default=0.0.

For details, see *Haze*.

To apply a map or shader to this parameter, click the None button.

Horizon and Ground group

Horizon Height—The vertical position of the horizon. Default=0.0.

For details, see *Horizon Height*.

Blur—The "blurriness" with which the horizon is rendered. Default=0.1.

At 0.0 the horizon is completely sharp. Generally only values lower than 0.5 are useful, but the full range is up to 10.0 for a horizon that consists of blur only, with no actual horizon at all.

Ground Color—The color of the virtual ground plane. Note that this is a diffuse reflectance value (that is, albedo). The ground appears as a Lambertian reflector with this diffuse color, lit by the sun and sky only, and does not receive any shadows.

To apply a map or shader to this parameter, click the None button.

Tip: Some sky models neglect the influence of bounce light from the ground, assuming only the sky is illuminating the scene. To compare the output of mr Sky with, for example, the IES Sky model, one must therefore set Ground Color to black.

After Dark group

Night Color—The minimum color of the sky: The sky will never become darker than this value. It can be useful for adding things like moon, stars, high-altitude cirrus clouds that remain lit long after sunset, etc. As the sun sets and the sky darkens, the contribution from Night Color is unaffected and remains as the base light level.

To apply a map or shader to this parameter, click the None button.

Non-Physical Tuning group

Red/Blue Tint—Provides artistic control over the redness of the sky light. The default value of 0.0 is the physically correct value (calculated for a 6500K whitepoint), but can be changed with this parameter, which ranges from -1.0 (extremely blue) to 1.0 (extremely red).

To apply a map or shader to this parameter, click the None button.

Saturation—Provides artistic control over saturation of the sky light. The default value of 1.0 is the physically calculated saturation level. Possible values range from 0.0 (black and white) to 2.0 (extremely high saturation).

To apply a map or shader to this parameter, click the None button.

Aerial Perspective group

Note: This parameter is for advanced users, and functions only when you use mr Physical Sky is

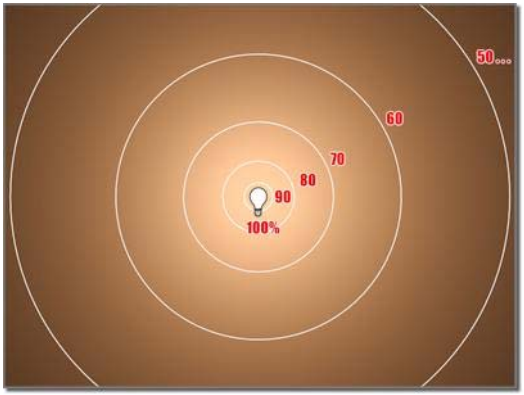
used as a lens shader or volume shader. When used as a lens shader, the aerial perspective "mist" applies to primary rays only; it doesn't appear in reflections. When used as a volume shader, it applies to the entire scene and every ray in it, including reflections and refraction.

Visibility Distance—Aerial Perspective is a term used by painters to convey how distant objects are perceived as hazier and tinted towards the blue end of the spectrum. mr Sky emulates this with the Visibilty Distance parameter. When nonzero, it defines the "10% distance", that is, the distance at which approximately 10% of haze is visible at a Haze level of 0.0.

Isotropic Light Distribution (Photometric Lights)

Create panel > Lights > Photometric Lights > Click Target Point or Free Point. > Intensity/Color/Distribution rollout > Choose Isotropic in Distribution list. > Click in viewport to create light.

An Isotropic light distributes the light equally in all directions.



Isotropic light distribution



Icon for Target Point light with isotropic distribution

See also

Lights (page 2–1272)

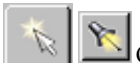
Target Point Light (Photometric) (page 2–1303)

Free Point Light (Photometric) (page 2–1304)

Intensity/Color/Distribution Rollout (page 2–1352)

Procedure

To create a light with an Isotropic distribution:

1.  On the Create panel, click Lights.
2. Choose Photometric from the drop-down list. (The default is Standard.)
3. In the Object Type rollout, select Free Point.
4. In the Intensity/Color/Distribution rollout, select Isotropic from the Distribution list.
5. Click in a viewport to position the light.

   The light is now part of the scene.
6. Set the creation parameters.

   You can use the Move transform to adjust the light.

   You can also use the *Place Highlight (page 1–467)* command to change a light's position.

Diffuse Distribution (Photometric Lights)

Create panel > Lights > Photometric Lights > Click Target Linear, Free Linear, Target Area, or Free Area. > Intensity/Color/Distribution rollout > Choose Diffuse in Distribution list. > Click in viewport to create light.

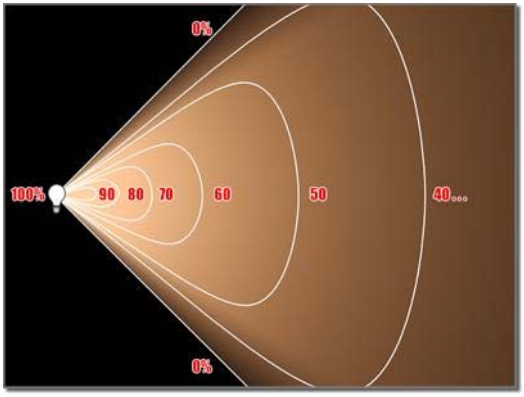
Diffuse distribution emits light from a surface. Light that leaves the surface at a right angle is

at the light's greatest intensity. At increasingly oblique angles, the intensity of the emitted light diminishes.

Spotlight Distribution (Photometric Lights)

Create panel > Lights > Photometric Lights > Click Target Point or Free Point. > Intensity/Color/Distribution rollout > Choose Spotlight in Distribution list. > Click in viewport to create light.

A spotlight distribution casts a focused beam of light like a flashlight, a follow spot in a theater or a headlight. At the light's *beam angle,* intensity has fallen to 50 per cent. At its *field angle,* intensity has fallen to zero. (The beam angle is similar to the hotspot angle for standard lights, but all of a hotspot is at 100 per cent intensity. The field angle is comparable to the falloff angle for standard lights.)



Spotlight distribution



Icon for Target Point light with spotlight distribution

See also

Lights (page 2–1272)

Target Point Light (Photometric) (page 2–1303)

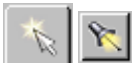
Free Point Light (Photometric) (page 2–1304)

Spotlight Parameters (page 2–1338)

Intensity/Color/Distribution Rollout (page 2–1352)

Procedures

To create a light with a spotlight distribution:

1.  On the Create panel, click Lights.
2. Choose Photometric from the drop-down list. (Standard is the default.)
3. In the Object Type rollout, click Target Point or Free Point.
4. In the Intensity/Color/Distribution rollout, select Spotlight from the Distribution list.
5. If you are creating a Target Point light, drag in a viewport. The initial point of the drag is the location of the light, and the point where you release the mouse is the location of the target. If you are creating a Free Point light, click in a Viewport to position the light.

   The light is now part of the scene.
6. Set the creation parameters.

   You can use the Move transform to adjust the light.

   You can also use the *Place Highlight (page 1–467)* command to change a light's position.

To adjust a target point spotlight:

1. Select the light.
2.  Use Move on the main toolbar to adjust the light.

   Because the spotlight is always aimed at its target, you can't rotate it about its local X or Y

axis. However, you can select and move the target object as well as the light itself. When you move either the light or the target, the light's orientation changes so it always points at the target.

You can use the *Place Highlight (page 1–467)* command to change a light's position. Another way to adjust a spotlight is to use a *Light viewport (page 3–750)*.

To select the target:

If the spotlight distribution is on a target point light, you can move the target to adjust the direction of the spotlight. The target, displayed as a small square, is often in the same area as objects that you want to illuminate. It can be difficult to select it by clicking.

1. First select the spotlight itself.
2. Right-click the light and choose Select Target from the pop-up menu. Clicking the line that connects the light and its target selects both objects.

To change a viewport to a light view:

1. Right-click the viewport label.

   The Viewport Right-Click menu is displayed.
2. Choose Views.

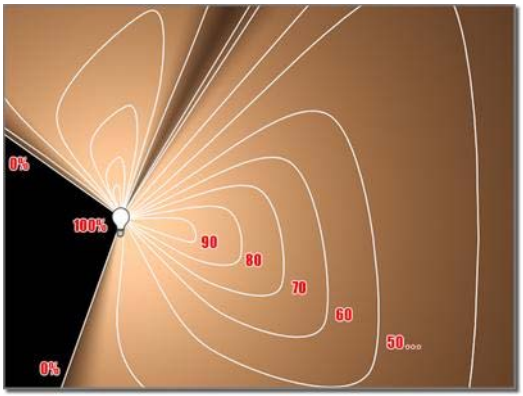
   The name of each light is displayed in the Views list. By default, photometric lights are named Point01, Point02, and so on.
3. Choose the name of the light you want.

   The viewport now shows the light's point of view. You can use the *Light Viewport Controls (page 3–750)* to adjust the light.

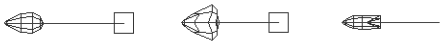
   The default keyboard shortcut for Light viewports is **$**.

Web Distribution (Photometric Lights)

A Web distribution uses a *photometric web definition (page 2–1326)* to distribute the light. A photometric web is a 3D representation of the light intensity distribution of a light source. Web definitions are stored in files. Many lighting manufacturers provide web files that model their products; these are often available on the Internet. A web file can be in the *IES (page 2–1328)*, *LTLI (page 3–964)*, or *CIBSE (page 3–921)* format. Controls for assigning the web file are on the Web Parameters rollout.



Example of web distribution



Various web distributions

See also

Lights (page 2–1272)

Target Point Light (Photometric) (page 2–1303)

Free Point Light (Photometric) (page 2–1304)

Target Linear Light (Photometric) (page 2–1305)

Free Linear Light (Photometric) (page 2–1307)

Target Area Light (Photometric) (page 2–1307)

Free Area Light (Photometric) (page 2–1309)

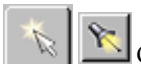
Intensity/Color/Distribution Rollout (page 2–1352)

Web Parameters Rollout (page 2–1355)

IES Standard File Format (page 2–1328)

Procedure

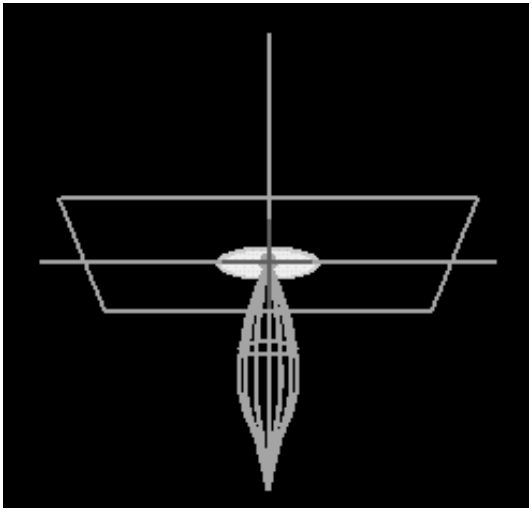
To create a light from a manufacturer's IES file:

1.  On the Create panel, click Lights.
2. Choose Photometric from the drop-down list. (The default is Standard.)
3. In the Object Type rollout, select any photometric light type.
4. In the Intensity/Color/Distribution rollout, select Web from the Distribution list.
5. If you are creating a Target Point, Linear, or Area light, drag in a viewport. The initial point of the drag is the location of the light, and the point where you release the mouse is the location of the target. If you are creating a Free Point, Linear or Area light, click in a viewport to position the light.

   A light is created using the default photometric web which is a diffuse distribution shining from one edge.
6. Go to the Web Parameters rollout, and select the IES file that you want to use. See *Web Parameters (page 2–1355)*.
7. Set the creation parameters.

   You can use the Move transform to *position the light (page 2–1282)*. You can also use the *Place Highlight (page 1–467)* command to change a light's position.

Photometric Webs



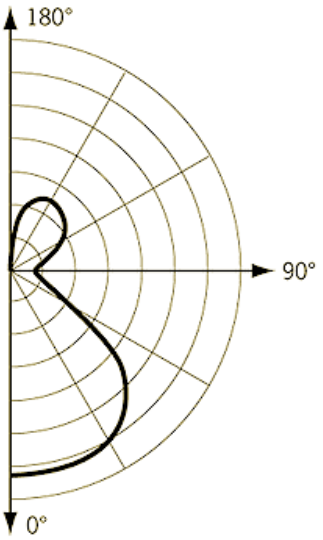
Light that uses a photometric web

A photometric web is a 3D representation of the light intensity distribution of a light source. This directional light distribution information is stored in a photometric data file in the IES format using the *IES LM-63-1991 standard file format (page 2–1328)*, or in the LTLI or CIBSE formats for photometric data. You can load photometric data files provided by various manufacturers as *web parameters (page 2–1355)*. The light icon represents the photometric web you select.

To describe the directional distribution of the light emitted by a source, 3ds Max approximates the source by a point light placed at its photometric center. With this approximation, the distribution is characterized as a function of the outgoing direction only. The luminous intensity of the source for a predetermined set of horizontal and vertical angles is provided, and the system can compute the luminous intensity along an arbitrary direction by interpolation.

Goniometric Diagrams

Photometric data is often depicted using a goniometric diagram.



Goniometric diagram of a web distribution

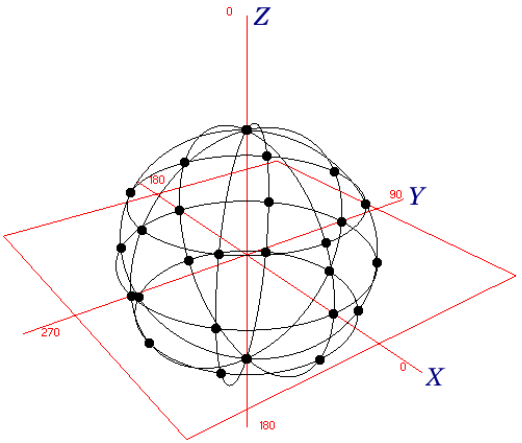
This type of diagram visually represents how the luminous intensity of a source varies with the vertical angle. However, the horizontal angle is fixed and, unless the distribution is axially symmetric, more than one goniometric diagram may be needed to describe the complete distribution.

Photometric Web

The photometric web is a three-dimensional representation of the light distribution. It extends the goniometric diagram to three dimensions, so that the dependencies of the luminous intensity on both the vertical and horizontal angles can be examined simultaneously. The center of the photometric web represents the center of the light object.

The luminous intensity in any given direction is proportional to the distance between this web and the photometric center, measured along a line leaving the center in the specified direction.

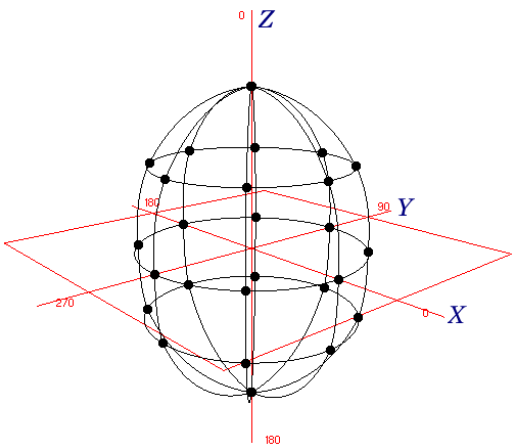
Example 1: Isotropic Distribution



Example of isotropic distribution

A sphere centered around the origin is a representation of an isotropic distribution. All the points in the diagram are equidistant from the center and therefore light is emitted equally in all directions.

Example 2: Ellipsoidal Distribution



Example of ellipsoidal distribution

In this example, the points in the negative Z direction are the same distance from the origin as the corresponding points in the positive Z direction, so the same amount of light shines upward and downward. No point has a very large X or Y component, either positive or negative, so less light is cast laterally from the light source.

IES Standard File Format

You can create a photometric data file in the IES format using the IES LM-63-1991 standard file format for photometric data. (IES stands for Illuminating Engineering Society.) However, only the information relevant to 3ds Max is described here. For a complete description of the IES standard file format, see IES Standard File Format for Electronic Transfer of Photometric Data and Related Information, prepared by the IES Computer Committee (http://www.iesna.org).

The luminous intensity distribution (LID) of a luminaire is measured at the nodes of a photometric web for a fixed set of horizontal and vertical angles. The poles of the web lie along the vertical axis, with the nadir corresponding to a vertical angle of zero degrees. The horizontal axis corresponds to a horizontal angle of zero degrees and is oriented parallel to the length of the luminaire. This type of photometric web is generated by a Type C goniometer and is the most popular in North America; other types of goniometry are supported by the IES standard file format but are not discussed here.

The photometric data is stored in an ASCII file. Each line in the file must be less than 132 characters long and must be terminated by a carriage return/line-feed character sequence. Longer lines can be continued by inserting a carriage return/line-feed character sequence. Each field in the file must begin on a new line and must appear exactly in the following sequence:

1. IESNA91
2. [TEST] the test report number of your data
3. [MANUFAC] the manufacturer of the luminaire
4. TILT=NONE
5. 1
6. The initial rated lumens for the lamp used in the test or -1 if absolute photometry is used and the intensity values do not depend on different lamp ratings.
7. A multiplying factor for all the candela values in the file. This makes it possible to easily scale all the candela values in the file when the measuring device operates in unusual units—for example, when you obtain the photometric values from a catalog using a ruler on a goniometric diagram. Normally the multiplying factor is 1.
8. The number of vertical angles in the photometric web.
9. The number of horizontal angles in the photometric web.
10. 1
11. The type of unit used to measure the dimensions of the luminous opening. Use 1 for feet or 2 for meters.
12. The width, length, and height of the luminous opening. Currently, Lightscape ignores these dimensions because you can associate a given luminous intensity distribution with any of the luminaire geometric entities supported by Lightscape. It is normally given as 0 0 0.
13. 1.0 1.0 0.0
14. The set of vertical angles, listed in increasing order. If the distribution lies completely in the bottom hemisphere, the first and last angles must be 0° and 90°, respectively. If the distribution lies completely in the top hemisphere, the first and last angles must be 90°

and 180°, respectively. Otherwise, they must be 0° and 180°, respectively.

15. The set of horizontal angles, listed in increasing order. The first angle must be 0°. The last angle determines the degree of lateral symmetry displayed by the intensity distribution. If it is 0°, the distribution is axially symmetric. If it is 90°, the distribution is symmetric in each quadrant. If it is 180°, the distribution is symmetric about a vertical plane. If it is greater than 180° and less than or equal to 360°, the distribution exhibits no lateral symmetries. All other values are invalid.
16. The set of candela values. First all the candela values corresponding to the first horizontal angle are listed, starting with the value corresponding to the smallest vertical angle and moving up the associated vertical plane. Then the candela values corresponding to the vertical plane through the second horizontal angle are listed, and so on until the last horizontal angle. Each vertical slice of values must start on a new line. Long lines may be broken between values as needed by following the instructions given earlier.

Here is an *example of a photometric data file (page 2–1329)*.

Example of Photometric Data File

The following is an example of a photometric data file:

```
IESNA91
[TEST] Simple demo intensity distribution
[MANUFAC] Lightscape Technologies, Inc.
TILT=NONE
1
-1
1
8
1
1
2
0.0 0.0 0.0
1.0 1.0 0.0
0.0 5.0 10.0 20.0 30.0 45.0 65.0 90.0
0.0
1000.0 1100.0 1300.0 1150.0 930.0 650.0 350.0 0.0
```

See also

IES Standard File Format (page 2–1328)

Common Lamp Values for Photometric Lights

The following table lists some commonly used lamps values that you can use as a guide for defining photometric lights. The information in the table is approximate; however, you can refer to manufacturer's documentation for more precise photometric data for these lamps.

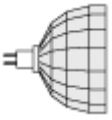
General Purpose Lamps

Note: In the tables below, "Class." stands for "Classification"; the values under "Intensity" are expressed in candelas.



Class.	Watts	Type	Intensity	Beam	Field
A-19/ Med	60	Point	70		
A-19/ Med	75	Point	95		
A-19/ Med	100	Point	139		

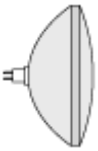
M16 Low Voltage Lamps



Class.	Watts	Type	Intensity	Beam	Field
Narrow Beam	20	Spot	3300	6	12
Narrow Beam	50	Spot	9150	12	25

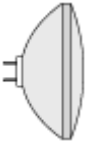
Class.	Watts	Type	Intensity	Beam	Field
Medium Beam	50	Spot	3000	25	50
Wide Beam	20	Spot	460	38	75
Wide Beam	50	Spot	1500	38	75

Par36 Low Voltage Lamps



Class.	Watts	Type	Intensity	Beam	Field
Narrow Beam	25	Spot	4200	9	15
Narrow Beam	50	Spot	8900	10	15
Medium Beam	50	Spot	1300	30	60
Wide Beam	25	Spot	250	36	75
Wide Beam	50	Spot	600	39	75

Par56 Line Voltage Lamps



Class.	Watts	Type	Intensity	Beam	Field
Narrow Beam	300	Spot	68000	9	15
Narrow Beam	500	Spot	95000	9	15
Medium Beam	300	Spot	24000	18	36
Medium Beam	500	Spot	47500	18	36
Wide Beam	300	Spot	10000	30	60
Wide Beam	500	Spot	18000	30	60

Par38 Line Voltage Lamps



Class.	Watts	Type	Intensity	Beam	Field
Narrow Beam	45	Spot	4700	14	28
Narrow Beam	75	Spot	5200	12	25
Narrow Beam	150	Spot	10500	14	28
Medium Beam	45	Spot	1700	28	60
Medium Beam	75	Spot	1860	30	60
Medium Beam	150	Spot	4000	30	60

R40 Line Voltage Lamps



Class.	Watts	Type	Intensity	Beam	Field
Narrow Beam	150	Spot	5400	22	50
Wide Beam	150	Spot	1040	76	130
Wide Beam	300	Spot	1950	76	130

Common Lighting Rollouts

General Lighting Parameters

Create panel > Lights > Create a light. > General Parameters rollout

Create menu > Photometric Lights > Create a light. > Modify panel > General Parameters rollout.

The General Parameters rollout is displayed for all types of lights. These controls turn a light on and off, and exclude or include objects in the scene.

On the Modify panel, the General Parameters rollout also lets you control the light's target object and change the light from one type to another.

Shadow Types

The General Parameters rollout also lets you turn shadow-casting on or off for the light, and choose which type of shadow the light uses.

The following table describes the advantages and disadvantages of each shadow type:

Shadow Type	Advantages	Disadvantages
Advanced Ray-Traced	Supports transparency and opacity mapping. Uses less RAM than standard ray-traced shadows. Recommended for complex scenes with many lights or faces.	Slower than shadow maps. Does not support soft shadows. Processes at every frame.
Area Shadows	Supports transparency and opacity mapping. Uses very little RAM. Recommended for complex scenes with many lights or faces. Supports different formats for area shadows.	Slower than shadow maps. Processes at every frame.
mental ray Shadow Maps	Can be quicker than ray-traced shadows with the mental ray renderer.	Not as accurate as ray-traced shadows.
Raytrace Shadows	Supports transparency and opacity mapping. Processes only once if there are no animated objects.	Can be slower than shadow maps. Does not support soft shadows.
Shadow Maps	Produces soft shadows. Processes only once if there are no animated objects. Fastest shadow type.	Uses a lot of RAM. Does not support objects with transparency or opacity maps.

Procedures

To turn a light on and off:

- Turn the On toggle on or off.

  Default=on.

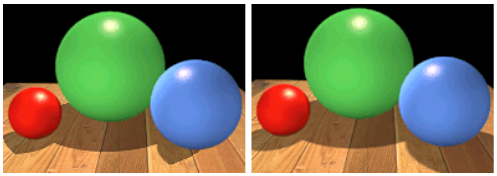
  Tip: To make a standard light flash on and off, set its Multiplier parameter to 0 in repeated keyframes, and assign a *step tangent (page 2–305)* to this parameter.

To have a light cast shadows:

- In the General Parameters rollout > Shadows group, make sure On is selected.

The light will now cast shadows when you render the scene.

Note: With Overshoot turned on, standard spotlights cast light in all directions but cast shadows only within the falloff cone; standard directional lights cast light throughout the scene, but cast shadows only within the falloff area.



Left: A spotlight's projection cone truncates shadows.
Right: An omni light casts complete shadows.

To have a light use the global settings for shadows:

- On the General Parameters rollout, turn on Use Global Settings.

  When Use Global Settings is on, the other shadow controls are set to the values used by all other shadow-casting lights in the scene that have Use Global Settings set.

  Changing the affected parameters for one light with Use Global Settings set, changes them for all lights with Use Global Settings set.

To set a light's shadow parameters individually:

- On the General Parameters rollout, turn off Use Global Settings.

  The settings revert to the individual settings for the light.

  By default, Shadow Map is the active shadow type. In the *Shadow Map Params rollout (page 2–1363)*, the default settings are: Map Bias=1; Size=512; Sample Range=4.0; Absolute Map Bias=Off.



Scene with shadow-mapped shadows
Shadows rendered using default parameter settings

Note: When you render a scene, you can turn rendering of shadows on or off.

To cast area shadows:

- On the General Parameters rollout, choose Area Shadows from the drop-down list.

  Use controls on the *Area Shadows rollout (page 2–1357)* to adjust the shadow properties.

To cast advanced ray-traced shadows:

Advanced ray-traced shadows are similar to ray-traced shadows, however they provide control over *antialiasing (page 3–907)*, letting you fine-tune how shadows are generated.

- On the General Parameters rollout, choose Advanced Ray-traced Shadows from the drop-down list.

  Use controls on the *Advanced Ray-traced Params rollout (page 2–1356)* to adjust the shadow properties.

To cast shadow-mapped shadows:

1. On the General Parameters rollout, choose Shadow Maps from the drop-down list.
2. Go to the *Shadow Map Params rollout (page 2–1363)*.

- Use the Size spinner to set the size of the shadow map.
- Use the Bias spinner to adjust the shadow offset, if necessary.
- Use the Sample Range spinner to create a soft-edged shadow.

To cast ray-traced shadows:

Ray-traced shadows (page 3–1000) are generated by tracing the path of rays sampled from a light source. Ray-traced shadows are more accurate than shadow-mapped shadows.

1. On the General Parameters rollout, choose Ray-Traced Shadows from the drop-down list.
2. Use controls on the *Ray Traced Shadow Params rollout (page 2–1363)* to adjust the shadow offset, if necessary.

To keep an object from casting shadows:

1. Select the object.
2. Right-click the object to display the quad menu and choose Properties from the Transform quadrant.

   The *Object Properties dialog (page 1–117)* is displayed.
3. Turn off Cast Shadows, and then click OK.

   Now when you render the scene, the object casts no shadows.

To make an object not receive shadows:

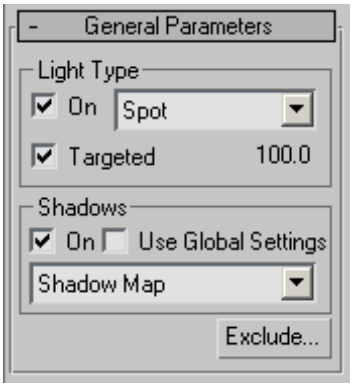
1. Select the object.
2. Right-click the object to display the quad menu and choose Properties from the lower-right (Transform) quadrant.

   The *Object Properties dialog (page 1–117)* is displayed.
3. Turn off Receive Shadows, and then click OK.

   Now when you render the scene, the object receives no shadows.

Note: You can also prevent objects from casting shadows by excluding them from a light.

Interface



Light Type group (Modify panel)

On—(Both Create panel and Modify panel) Turns the light on and off. When On is on, shading and rendering use the light to illuminate the scene. When On is off, the light is not used in shading or rendering. Default=on.

In viewports, the interactive renderer shows the effect of turning lights on or off.

Light Type List—Changes the type of the light. If you have selected a *standard light type (page 2–1288)*, the light can be changed to an omni, spot, or directional light. If you have selected a *photometric light type (page 2–1301)*, you can change the light to a point, linear, or area light.

This parameter is available only in the Modify panel.

Targeted—When on, the light is targeted. The distance between the light and its target is displayed to the right of the check box. For a free light, you can set this value. For target lights, you can change it by turning off the check box, or by moving the light or the light's target object.

Shadows group

On—Determines whether the current light casts shadows or not. Default=on.

Shadow Method drop-down list—Determines whether the renderer uses *shadow maps (page 3–1010)*, *ray-traced shadows (page 3–1000)*, *advanced ray-traced shadows (page 3–1000)*, or *area shadows (page 3–911)*, to generate shadows for this light.

The "mental ray Shadow Map" type is provided for use with the *mental ray renderer (page 3–78)*. When you choose this shadow type and enable shadow maps (on the *Shadows & Displacement rollout (page 3–114)* of the Render Scene dialog), shadows use the mental ray shadow-map algorithm. If this type is chosen but you render with the default scanline renderer, no shadows appear in the rendering.

Note: When shadow maps are enabled and the shadow map type is Shadow Map, the mental ray renderer attempts to translate the shadow map settings into comparable settings for mental ray shadow maps. (The results might not be what you expected.) In all other cases, the mental ray renderer generates ray-traced shadows.

Each shadow type has its particular controls:

- **Advanced ray-traced shadows—**See *Advanced Ray-Traced Parameters Rollout (page 2–1356)* and *Optimizations Rollout (page 2–1361)*.
- **Area shadows—**See *Area Shadows Rollout (page 2–1357)* and *Optimizations Rollout (page 2–1361)*.
- **mental ray shadow maps—**See *mental ray Shadow Map Rollout (page 2–1360)*.
- **Ray-traced shadows—**See *Ray-Traced Shadow Parameters Rollout (page 2–1363)*.
- **Shadow maps—**See *Shadow Map Parameters Rollout (page 2–1363)*.

Tip: Use ray-traced or advanced ray-traced shadows when you want shadows cast by opacity-mapped objects. Shadow-mapped shadows don't recognize the transparent portions of the mapping, and as a result they don't look convincing.

Use Global Settings—Turn on to use global settings for shadows cast by this light. Turn off to enable individual control of the shadows. If you choose not to use the global settings, you must choose which method the renderer will use to generate shadows for this particular light.

When Use Global Settings is on, the shadow parameters switch over to show you what the global settings are. This data is shared by every other light of this class. When Use Global Settings is off, the shadow parameters are specific to that particular light.

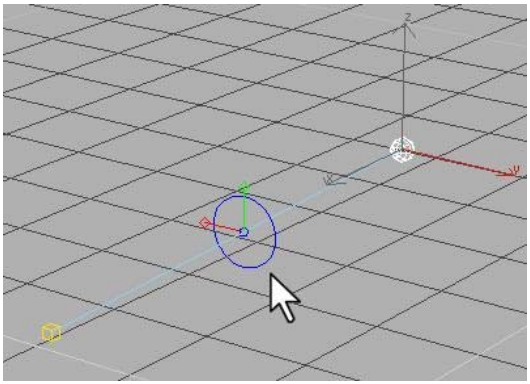
Exclude button

Excludes selected objects from the effects of the light. Click this button to display the *Exclude/Include dialog (page 2–1335)*.

Excluded objects still appear lit in shaded viewports. Exclusion takes effect only when you render the scene.

Roll Angle Manipulator

To change the roll angle of a target light, you can use a manipulator. This can be useful if the light does not cast a round beam, or if it is a projector light (see *Advanced Effects Rollout (page 2–1341)*). To display the manipulator, select the light, right-click it, and then click Roll Angle Manipulator Toggle in the Tools 1 (upper left) quadrant of the quad menu.



Roll angle manipulator on a target light

Drag the circular portion of the manipulator to rotate the light by an arbitrary number of degrees. Click one of the red or green "clock hands" of the manipulator to rotate the light in 90-degree increments.

Note: You can also access the roll angle manipulator by selecting the light object and then turning on *Select And Manipulate (page 2–15)* on the main toolbar.

Exclude/Include Dialog

Create a light. > General Parameters rollout > Exclude button

Rendering menu > Render > Render Scene dialog > Render Elements panel > Render Elements rollout > Add button > Render Elements dialog > Select an existing Matte element in the list. > Matte Texture Element rollout > Include button

The Exclude/Include dialog determines which objects are not illuminated by a selected light, or which objects are considered in a *Matte render element (page 3–130)*.

Although light exclusion does not occur in nature, this feature is useful when you need exact control over the lighting in your scene. Sometimes, for example, you'll want to add lights specifically to illuminate a single object but not its surroundings, or you'll want a light to cast shadows from one object but not from another.

Procedures

To exclude objects from a light:

1. In the General Parameters rollout, click Exclude.

   The Exclude/Include dialog is displayed.
2. Make sure Exclude is on.
3. Choose to exclude objects from Illumination, Shadow Casting, or Both.
4. In the list of object names on the left, highlight the objects you want to exclude.
5. Click the arrow pointing to the right to move the selected object names into the right-hand list.

   Objects in the right-hand list will be excluded. To remove an object from this list, highlight its name and click the arrow pointing to the left. To remove all objects from this list, click Clear.
6. Click OK.

   Now when you shade or render the scene, the light will have no effect on the objects you chose.

   The object still appears lit in shaded viewports. Exclusion takes effect only when you render the scene.

To include objects that were excluded:

1. Click Exclude.

   The Exclude/Include dialog is displayed.
2. Select the names of objects in the right-hand list, and then click the arrow pointing to the left.
3. Click OK.

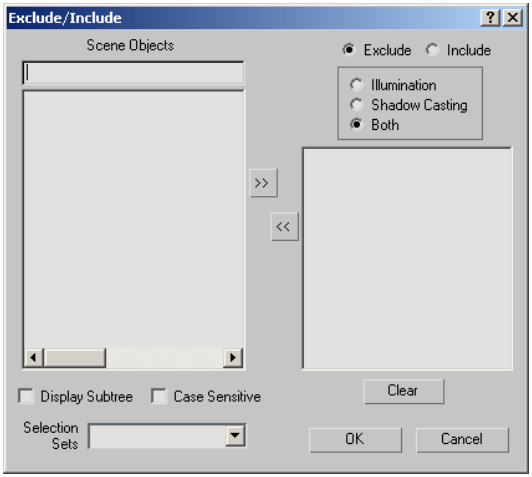
To include all excluded objects at once:

1. Click Exclude or Include.

   The Exclude/Include dialog is displayed.

2. Click Clear.
3. Click OK.

Interface



The Exclude/Include dialog contains the following controls:

Exclude/Include—Determines whether the light (or Matte render element) will exclude or include the objects named in the list on the right.

Illumination—Excludes or includes illumination of the object's surface.

Note: This control has no effect on the matte render element.

Shadow Casting—Excludes or includes creation of the object's shadow.

Note: This control is not available with the matte render element.

Both—Excludes or includes both of the above.

Note: This control is not available with the matte render element.

Scene Objects—Select objects from the Scene Objects list on the left, then use the arrow buttons to add them to the exclusion list on the right.

The Exclude/Include dialog treats a group as an object. You can exclude or include all objects in a group by selecting the group's name in the Scene Objects list. If a group is nested within another group, it isn't visible in the Scene Objects list. To exclude a nested group or individual objects within the group, you have to ungroup them before you use this dialog.

Search field—The edit box below the Scene Objects list searches for objects by name. You can enter names that use wildcards.

Display Subtree—Indents the Scene Objects list according to the object hierarchy.

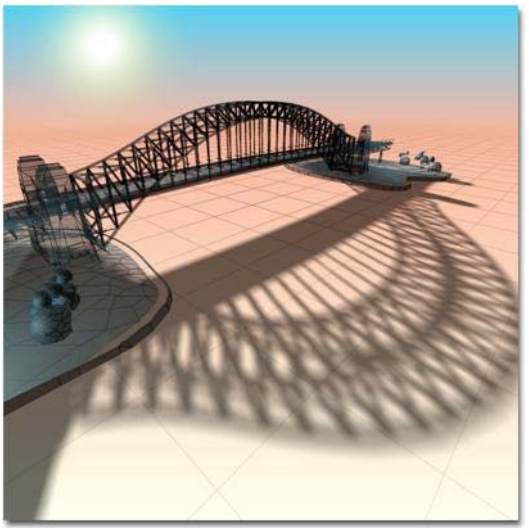
Case Sensitive—Uses case sensitivity when searching object names.

Selection Sets—Displays a list of *named selection sets (page 1–83)*. Choosing a selection set from this list selects those objects in the Scene Objects list.

Clear—Clears all entries from the Exclude/Include list on the right.

Shadow Parameters

Create panel > Lights > Create a light. > Shadow Parameters rollout



A bridge's shadow cast by sunlight

The Shadow Parameters rollout is displayed for all light types except Skylight and IES Sky, and for all shadow types. It lets you set shadow colors and other general shadow properties.

The controls also let atmospheric effects cast shadows.

See also

Advanced Ray-Traced Parameters Rollout (page 2–1356)

Area Shadows Rollout (page 2–1357)

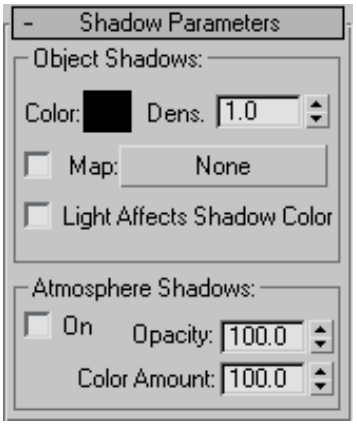
mental ray Shadow Map Rollout (page 2–1360)

Optimizations Rollout (page 2–1361)

Ray-Traced Shadow Parameters Rollout (page 2–1363)

Shadow Map Parameters Rollout (page 2–1363)

Interface



Color—Displays a *Color Selector (page 1–161)* to choose a color for the shadows cast by this light. Default=black.

You can animate the shadow color.

Dens. (density)—Adjusts the density of shadows.



The shadow density increases from right to left.

Increasing the Density value increases the density (darkness) of shadows. Decreasing Density makes shadows less dense. Default=1.0.

Density can have a negative value, which can help simulate the effect of reflected light. A white shadow color and negative Density render dark shadows, though the quality of these is not as good as a dark shadow color and positive Density.

You can animate the Density value.

Map check box—Turn on to use the map assigned with the Map button. Default=off.

Map—Assigns a map to the shadows. The map's colors are blended with the shadow color. Default=none.

Light Affects Shadow Color—When on, blends the light's color with the shadow color (or shadow colors, if the shadow is mapped). Default=off.



A checker map is used to alter the shadow cast by the piano.

Atmosphere Shadows group

These controls let atmospheric effects cast shadows.

On—When on, atmospheric effects cast shadows as the light passes through them. Default=off.

Note: This control is independent of the On toggle for normal Object Shadows. A light can cast atmospheric shadows but not normal shadows, or vice versa. It can cast both kinds of shadows, or neither.

Opacity—Adjusts the opacity of the shadows. This value is a percentage. Default=100.0.

Color Amount—Adjusts the amount that the atmosphere's color is blended with the shadow color. This value is a percentage. Default=100.0.



A cloud casts a colored shadow on the city.

Spotlight Parameters

Create a standard Target Spotlight or Free Spotlight, or a photometric light with Spotlight distribution. > Spotlight Parameters rollout

The Spotlight Parameters rollout appears when you create or select a *Target Spot (page 2–1289)* or *Free Spot (page 2–1290)*, or a photometric light with *spotlight distribution (page 2–1324)*.

Procedures

To see the spotlight cone in viewports:

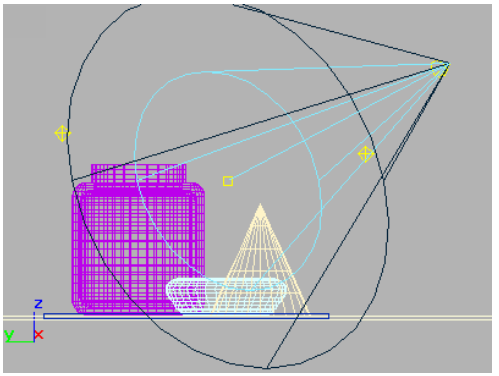
The cone is always visible while the light is selected. This setting keeps the cone visible when the light is unselected.

1. Select the spotlight.
2. In the Spotlight Parameters rollout > Light Cone group, turn on Show Cone.
3. Select another object, deselecting the light.

   A wireframe outline of the light's cone appears. The hotspot region is outlined in light blue, and the falloff region is outlined in dark gray.

   Note: The cone is always visible while the light is selected.

Tip: The hotspot and falloff borders are not visible in shaded viewports. Use controls in the Light Cone group to adjust hotspot and falloff in the viewport. Render the scene to see the full effect of the light's border.



Spotlight cone displayed in a viewport

To adjust the hotspot and falloff, do one of the following:

- Use the Hotspot and Falloff spinners to increase or decrease the size of the hotspot and falloff regions. For spotlights, Hotspot and Falloff are expressed in degrees.
- Use manipulators to adjust hotspot and falloff by dragging in a viewport, as described in the procedure that follows.

By default, falloff value is always constrained to be at least two degrees or two units greater than the hotspot value. The size of this constraint is a preference that you can change in the *Rendering panel (page 3–826)* of the Preferences dialog. The falloff value can range from 0 to 10,000 units.

When the hotspot and falloff values are close in value, the light casts a hard-edged beam.

To use manipulators to control hotspot and falloff:

1. Select the spotlight.

   Tip: When you select multiple spotlights, all their manipulators are accessible.

2.  On the default main toolbar, click to turn on Select And Manipulate.

   Now when you move the mouse over the hotspot or falloff circle, the circle turns red to show you can adjust it by dragging. Also, a tooltip displays the spotlight name, the parameter, and its value.

3. Drag the hotspot or falloff circle to adjust the value.

   The hotspot and falloff constrain each other, as their spinner controls do.

   Spotlight manipulators are described at the end of this topic.

To set the shape of the light beam:

1. Choose either Rectangle or Circle.

   The shape of the light's cone changes to reflect the shape you chose.

2. If you chose Circle, you are done. If you chose Rectangle, you can now adjust the *aspect ratio (page 3–911)* of the rectangular light in the Aspect field.

To change a rectangular light's aspect ratio:

1. Choose Rectangle, as described above.
2. Change the Aspect value to the aspect ratio that you want.

   The Bitmap Fit button is another way to set the aspect ratio. It is mainly used with projections.

Interface



Light Cone group

These parameters control *hotspots and falloff (page 3–954)* for spotlights.

Show Cone—Turns display of the cone on or off.

Note: The cone is always visible when a light is selected, so turning off this check box has no apparent effect until you deselect the light.

Overshoot—When *Overshoot (page 3–988)* is set, the light casts light in all directions. However, projections and shadows occur only within its falloff cone.

Hotspot/Beam—Adjusts the angle of a light's cone. The Hotspot value is measured in degrees. Default=43.0.

For photometric lights, the Beam angle is the angle at which the light's intensity has fallen to 50 per cent of its full intensity. (As opposed to the Hotspot, where intensity is still 100 per cent.)

Falloff/Field—Adjusts the angle of a light's falloff. The Falloff value is measured in degrees. Default=45.0.

For photometric lights, the Field angle is comparable to the Falloff angle. It is the angle at which the light's intensity has fallen to zero.

You can manipulate the hotspot and falloff by dragging manipulators in viewports, as described in the previous procedures, and the later section "Manipulators for Spotlights."

You can also adjust hotspot and falloff angles in a Light viewport (looking at the scene from the point of view of the spotlight).

Circle/Rectangle—Determine the shape of the falloff and hotspot areas. Set Circle when you want a standard, circular light. Set Rectangle when you want a rectangular beam of light, such as light cast through a window or doorway.

Aspect—Sets the *aspect ratio (page 3–911)* for the rectangular light beam. The Bitmap Fit button lets you make the aspect ratio match a specified bitmap. Default=1.0.

Bitmap Fit—If the light's projection aspect is rectangular, sets the aspect ratio to match a particular bitmap. This is useful when you are using the light as a *projector light (page 3–998)*.

Manipulators for Spotlights


Manipulators are visible and usable while the *Select And Manipulate button (page 2–15)* is turned on. This button is on the default *main toolbar (page 3–685)*. When you move the mouse over a manipulator, the manipulator turns red to show that dragging or clicking it will have an effect. Also, a tooltip appears, showing the name of the object, the parameter, and its value.

Tip: When you select multiple spotlights, all their manipulators are accessible.

For more information on using the spotlight manipulators, see the Procedures section at the top of this topic.

Hotspot manipulator: In a viewport, drag the hotspot circle to adjust the hotspot value.

Falloff manipulator: In a viewport, drag the falloff to adjust the falloff value.

The hotspot and falloff constrain each other, as their spinner controls do.

Advanced Effects Rollout

Create panel > Lights > Create a light. > Advanced Effects rollout

The Advanced Effects rollout provides controls that affect how lights affect surfaces, and also includes a number of fine adjustments and setting for projector lights.

You can make a *light object (page 2–1272)* into a projector by choosing a map for the light to project. A projected map can be a still image or an animation.



Casting a projection with a light

Note: If your scene includes animated bitmaps, including materials, projector lights, environments, and so on, the animation file is reloaded once per frame. If your scene uses multiple animations, or if the animations are large files, this can slow down rendering performance.

Procedures

To make a light a projector:

1.  Open the *Material Editor (page 2–1409)*. The Material Editor is where you adjust the map's parameters.
2. Use an unused sample slot to display a map.
3. Drag the map from the Material Editor to the light's Map button in the Advanced Effects rollout. A dialog asks if the projection map should be a copy or an instance. Choose Instance. If you choose Copy, adjusting the map in the Material Editor has no effect on the projected map.

   You can also drag from any other used map button, as in the Environment dialog.

   Assigning the map displays the map name in the button, and turns on the Projector toggle. After you have set up the map, you can turn off Projector to test rendering the scene without the projected image.

   An alternative is to click the Map button. This displays the *Material/Map Browser (page 2–1412)*, which lets you choose the map type. At this point, the light behaves as a projector. To assign a map or adjust its parameters, you need to use the Material Editor.

   Note: Lights project maps only within their cone, even if Overshoot is turned on.

To put the map in the Material Editor:

- Drag from the light's Map button to an unused sample slot in the Material Editor.

  Note: Sample slots with white triangles in the corners indicate materials that are used in the current scene.

  You can adjust the map in the Material Editor by changing the map's parameters.

To blur a projection map:

- Increase the value of Blur Offset in the map's Coordinates rollout in the Material Editor.

  You can animate Blur Offset to have a projected map go in or out of focus.

To make the shape of the light fit the projected bitmap:

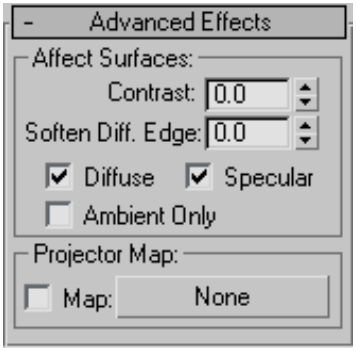
Note: This procedure only applies to standard spot and direct lights.

1. Choose a bitmap to project, as described in the preceding procedures.
2. Make sure the light's shape is set to Rectangular, and then click Bitmap Fit.

   A file selection dialog is displayed.
3. Choose the same bitmap you chose for the standalone map, and then click OK.

   You can also choose a bitmap other than the one the light projects.

Interface



Affect Surfaces group

Contrast—Adjusts the contrast between the diffuse and ambient areas of the surface. Leave this set to 0 for normal contrast. Increase the value to increase the contrast for special effects: for example, the harsh light of outer space. Default=0.0.

Soften Diffuse Edge—Increasing the value of Soften Diffuse Edge softens the edge between the diffuse and ambient portions of a surface. This helps eliminate edges that can appear on a surface under certain circumstances. Default=50.

Note: Soften Diffuse Edge reduces the amount of light, slightly. You can counter this, to some extent, by increasing the Multiplier value.

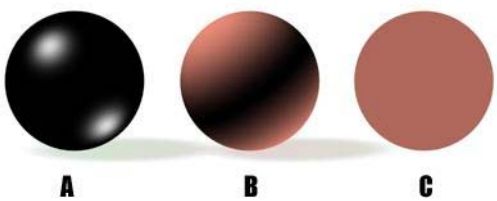
Diffuse—When on, the light affects the diffuse properties of an object's surface. When off, the light has no effect on the diffuse surface. Default=on.

Specular—When on, the light affects the specular properties of an object's surface. When off, the light has no effect on the specular properties. Default=on.

For example, by using the Diffuse and Specular check boxes you can have one light color the specular highlights of an object, while not coloring its diffuse area, and then have a second light color the diffuse portion of the surface while not creating specular highlights.

Ambient Only—When on, the light affects only the ambient component of the illumination. This gives you more detailed control over the ambient illumination in your scene. When Ambient Only is on, Contrast, Soften Diffuse Edge, Diffuse, and Specular are unavailable. Default=off.

The effect of Ambient Only isn't visible in viewports. It appears only when you render the scene.



A: Affect specular only

B: Affect diffuse only

C: Affect ambient only

Projector Map group

These controls make the photometric light a projector.

Check box—Turn on to project the map selected by the Map button. Turn off to turn off projection.

Map—Names the map used for the projection. You can drag from any map specified in the Material Editor, or any other map button (as on the Environment panel), and drop that map on the light's Map button. Clicking Map displays the Material/Map Browser. You can choose the map type using the Browser, then drag the button to the Material Editor, and use the Material Editor to select and adjust the map.

mental ray Indirect Illumination Rollout (for Lights)

Create panel > Lights > Create a light. > Modify panel > mental ray Indirect Illumination rollout

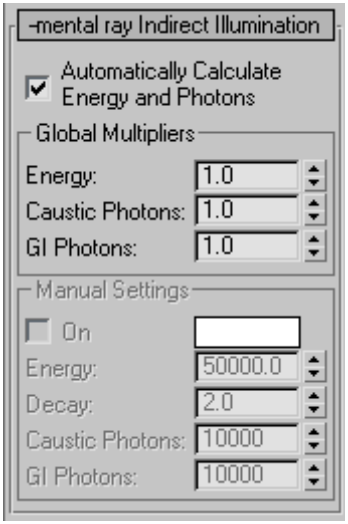
Note: This rollout does not appear on the Create panel.

The mental Indirect Illumination rollout provides controls for light behavior with the *mental ray renderer (page 3–78)*. The settings in this rollout have no effect on rendering with the default scanline renderer, or on advanced lighting (the light tracer or a radiosity solution). These settings control how the light behaves when it generates indirect illumination; that is, *caustics (page 3–92)* and *global illumination (page 3–93)*.

Tip: By default, each light uses the global settings that are found in the Global Light Properties group on Render Scene Dialog > Indirect Illumination panel > *Indirect Illumination rollout (page 2–1343)*. It is more convenient to adjust all lights in the scene at once. If you need to adjust a specific light, you can use the multiplier controls for energy and photons. In general, you should rarely if ever need to turn off Use Global Settings and specify local light settings for indirect illumination.

Important: In addition to the values specified here, the light must also be set to generate caustics, global illumination, or both. These controls are on the *mental ray panel (page 1–126)* of the Object Properties dialog. You also need to turn on Caustics, Global, or both, on the Render Scene dialog's *Indirect Illumination rollout (page 2–1343)*.

Interface



Automatically Calculate Energy and Photons—When on, the light uses the global light settings for indirect illumination, rather than local settings. Default=on.

When this toggle is on, only the controls in the Global Multipliers group are available.

Global Multipliers group

Energy—Multiplies the global Energy value to increase or decrease the energy of this particular light. Default=1.0.

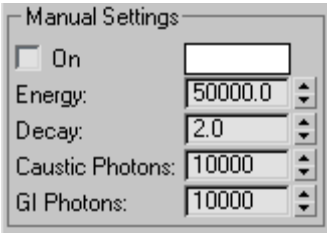
Caustic Photons—Multiplies the global Caustic Photons value to increase or decrease the count of

photons used to generate caustics by this particular light. Default=1.0.

GI Photons—Multiplies the global GI Photons value to increase or decrease the count of photons used to generate global illumination by this particular light. Default=1.0.

Manual Settings group

When Automatically Calculate is off, the Global Multipliers group becomes unavailable, and the manual settings for indirect illumination become available.



On—When on, the light can generate indirect illumination effects. Default=off.

Filter color—Click to display a *Color Selector (page 1–161)* and choose a color that filters the light energy. Default=white.

Energy—Sets the energy of the light. Energy, or "flux," is the amount of light used in indirect illumination. Each photon carries a fraction of the light's energy. This value is independent of the light intensity determined by the light's color and Multiplier, so you can use the Energy value to fine-tune indirect illumination effects without changing the light's other effects in a scene (such as providing diffuse illumination). Default=50000.0.

Decay—Specifies how photon energy decays as it moves away from the light source. Range=0.0 to 100.0. Default=2.0 (inverse square; physically correct falloff).

This value is applied as an exponent to the distance (*r*) between the light source and each indirectly illuminated point to determine the energy at that point, in the formula $1/r^{[Decay]}$), and is intended for tweaking the GI solution. The most commonly used values are:

- **0.0 (no decay)—**The energy doesn't decay, and photons provide equal indirect illumination throughout the scene.
- **1.0 (inverse)—**The energy decays proportionally to its distance from the light, with linear falloff. That is, a photon's energy is $1/r$, where r is the distance from the light source.
- **2.0 (inverse square)—**The energy decays at an inverse square rate. That is, a photon's energy is the inverse of the square of the distance (r) from the light source ($1/r^2$).

In the real world, light decays at the inverse square rate, but this gives strictly realistic results only if you provide a realistic value for the energy of the light. You can use other values to help adjust indirect illumination without worrying about physical accuracy.

Caustic Photons—Sets the number of photons emitted by the light for use in caustics. This is the number of photons in the *photon map (page 3–994)* used for caustics. Increasing this value increases the accuracy of caustics, but also increases the amount of memory used and the length of render time. Decreasing this value improves memory usage and render time, and can be useful for previewing caustic effects. Default=10000.

GI Photons—Sets the number of photons emitted by the light for use in global illumination. This is the number of photons in the photon map used for global illumination. Increasing this value increases the accuracy of global illumination, but also increases the amount of memory used and the length of render time. Decreasing this value improves memory usage and render time, and can be useful for previewing global-illumination effects. Default=10000.

mental ray Light Shader Rollout

Select a light. > Modify panel > mental ray Light Shader rollout

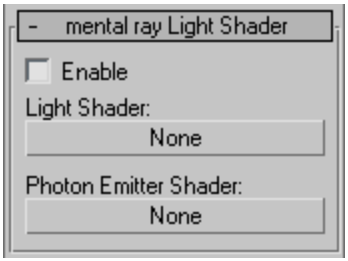
Note: This rollout appears only if you have enabled the mental ray extensions by using the mental ray Preferences panel.

Note: This rollout does not appear on the Create panel; only on the Modify panel.

The mental ray Light Shader rollout lets you add *mental ray shaders (page 2–1710)* to lights. When you render with the *mental ray renderer (page 3–78)*, light shaders can alter or adjust the light's effect.

To adjust the settings for a light shader, drag the shader's button to an unused Material Editor sample slot. If the software prompts you to choose Instance or Copy, be sure to choose Instance. (If you edit a copy of the shader, you will have to drag the sample slot back to the shader button on the Light Shader rollout before any changes take effect.)

Interface



Enable—When on, rendering uses the light shaders you have assigned to this light. When off, the shaders have no effect on rendering. Default=off.

Light Shader—Click the button to display a *Material/Map Browser (page 2–1412)* and choose a light shader. Once you have chosen a shader, its name appears on the button.

These are the light shaders provided with 3ds Max:

Shader	Library
Ambient/Reflective Occlusion	base
Light Infinite	base
Light Point	base
Light Spot	base

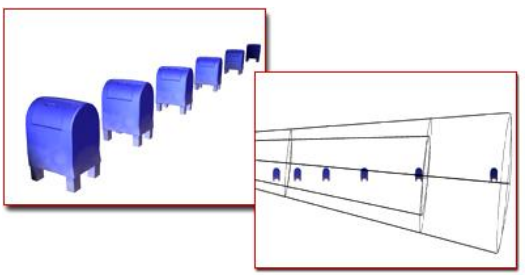
Photon Emitter Shader—Click the button to display a *Material/Map Browser (page 2–1412)* and choose a shader. Once you have chosen a shader, its name appears on the button.

Note: No photon emitter shaders are provided with 3ds Max. This option is for users who have access to light map shaders via other shader libraries or custom shader code.

Additional Rollouts for Standard Lights

Intensity/Color/Attenuation Parameters

Create panel > Lights > Create a Standard light. > Attenuation Parameters rollout



Attenuation settings cause distant objects to be dimmer.

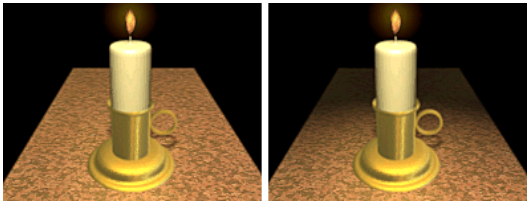
The Intensity/Color/Attenuation rollout allows you to set the color and intensity of your light. You can also define the attenuation of your light.

Attenuation (page 3–912) is the effect of light diminishing over distance. In 3ds Max, you can set attenuation values explicitly. The effects can vary from real-world lights, giving you more direct control over how lights fade in or fade out.

Note: With no attenuation, an object can paradoxically appear to grow brighter as it moves away from the light source. This is because the angle of incidence more closely approaches 0 degrees for more of the object's faces.

Two sets of values control attenuation for objects. The Far attenuation value sets the distance at which the light drops off to zero. The Near attenuation value sets the distance at which the light "fades in." Both these controls are turned on and off by a toggle called Use.

When Use is set for far attenuation, the light at its source uses the value specified by its color and multiplier controls. It remains at this value from the source to the distance specified by Start, then its value drops off to zero at the distance specified by End.



Adding attenuation to a scene

When Use is set for Near attenuation, the light value remains at zero up to the distance specified by Start. From Start to the distance specified by End, its value increases. Beyond End, the light remains at the value specified by the color and multiplier controls, unless far attenuation is also active.

You can't set the Near and Far attenuation distances so they overlap.

The Decay controls are an additional way to make a light fade out.

Tip: When lights are attenuated, the light might be too bright on near surfaces or too dim on far surfaces. If you see this in renderings, exposure control can help correct the problem. It adjusts the larger dynamic range of the (simulated) physical scene, into the smaller dynamic range of the display. See *Environment Panel (page 3–272)* for additional information on exposure control.

Procedures

To choose the color of a light using the color selector:

1. Click the color swatch next to the On check box.

   A modal *Color Selector (page 1–161)* is displayed.

2. Use the Color Selector to choose a new color, and then click OK.

   Shaded viewports update to show the new light color, which also appears in renderings.

To change a light's intensity:

- Increase the light's Multiplier value.

To use attenuation:

1. Set the Start and End values.
2. Turn on Use.

   Attenuation is now in effect for this light when the scene is shaded or rendered.

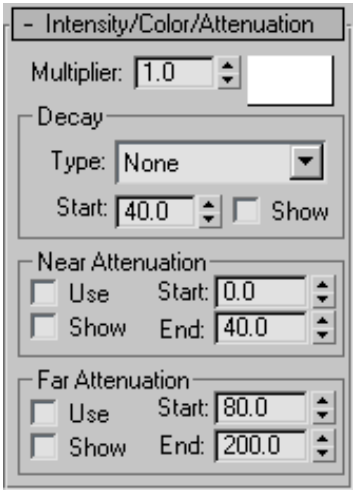
   Note: You can also change the decay type (and values) to reduce the light's intensity over distance.

To see the attenuation range in viewports:

- Set Show for far or near attenuation.

  You can preview the effect of attenuation in shaded viewports only if you turn this on as an option in *viewport preferences (page 3–821)*.

Interface



Multiplier—Amplifies the power of the light by a positive or negative amount. For example, if you set the *multiplier (page 3–977)* to 2, the light will be twice as bright. Multipliers are also useful for subtracting light and selectively placing dark areas in the scene. Default=1.0.

Using this parameter to increase intensity can cause colors to appear "burned out." It can also generate colors not usable in videos. In general, leave Multiplier set to its default of 1.0 except for special effects and special cases.

High Multiplier values wash out colors. For example, if you set a spotlight to be red but then increase its Multiplier to 10, the light is white in the hotspot and red only in the falloff area, where the Multiplier isn't applied.

Negative Multiplier values result in "dark light." That is, the light darkens objects instead of illuminating them.

Color Swatch—Shows the color of the light. Clicking the color swatch displays the *Color Selector (page 1–161)* so you can choose a color for the light.

Decay group

Decay is an additional way to make a light's intensity reduce over distance.

Type—Sets the type of decay to use. There are three types to choose from.

- **None—**(The default.) Applies no decay. The light maintains full strength from its source to infinity, unless you turn on far attenuation.
- **Inverse—**Applies inverse decay. The formula is luminance=R_0/R, where R_0 is the radial source of the light if no attenuation is used, or the Near End value of the light if attenuation is used. *R* is the radial distance of the illuminated surface from R_0.
- **Inverse Square—**Applies inverse-square decay. The formula for this is $(R_0/R)^2$. This is actually the "real-world" decay of light, but you might find it too dim in the world of computer graphics.

  This is the decay formula used by *photometric lights (page 2–1301)*.

  Tip: If Inverse Square decay makes the scene too dim, you can try using the *Environment Panel (page 3–272)* to increase the Global Lighting Level value.

The point at which decay begins depends on whether or not you use attenuation:

- With no attenuation, decay begins at the source of the light.
- With near attenuation, the decay begins at the Near End position.

Once the beginning point is established, the decay follows its formula to infinity, or until the light itself is cut off by the Far End distance. In other words, the distance between Near End and Far End does not scale, or otherwise affect, the apparent ramp of decaying light.

Tip: Because decay continues to calculate dimmer and dimmer values as the distance of the light throw increases, it's a good idea to set at least the Far End of attenuation to eliminate unnecessary calculations.

Near Attenuation group

Start—Sets the distance at which the light begins to fade in.

End—Sets the distance at which the light reaches its full value.

Use—Enables near attenuation for the light.

Show—Displays the near attenuation range settings in viewports. For spotlights, attenuation ranges appear as lens-shaped sections of the cone. For directional lights, the ranges appear as circular sections of the cone. For omni lights and spot or directional lights with Overshoot turned on, the ranges appear as spheres. By default, Near Start is dark blue and Near End is light blue.

Far Attenuation group

Start—Sets the distance at which the light begins to fade out.

End—Sets the distance at which the light has faded to zero.

Use—Enables far attenuation for the light.

Show—Displays the far attenuation range settings in viewports. For spotlights, attenuation ranges appear as lens-shaped sections of the cone. For directional lights, the ranges appear as circular sections of the cone. For omni lights and spot or directional lights with Overshoot turned on, the ranges appear as spheres. By default, Far Start is light brown and Far End is dark brown.

Directional Parameters

Create a standard Target Direct or Free Direct light. > Directional Parameters rollout

The Directional Parameters rollout appears when you create or select a *target direct (page 2–1292)* or *free direct (page 2–1293)* light. These parameters control *hotspots (page 3–954)* and *falloff (page 3–954)*.

Procedures

To see the directional cone in viewports:

The cone is always visible while the light is selected. This setting keeps the cone visible when the light is unselected.

1. Select the spotlight.
2. In the Directional Parameters rollout > Light Cone group, turn on Show Cone.
3. Select another object, deselecting the light.

   A wireframe outline of the light's cone appears. The hotspot region is outlined in light blue, and the falloff region is outlined in dark gray.

   The cone is always visible while the light is selected.

   Tip: The hotspot and falloff borders are not visible in shaded viewports. Use the cone values to adjust hotspot and falloff in the viewport. Render the scene to see the full effect of the light's border.

To adjust the hotspot and falloff:

- Use the Hotspot and Falloff spinners to increase or decrease the size of the hotspot and falloff regions. For directional lights, Hotspot and Falloff are expressed in units.

  By default, the falloff value is always constrained to be at least two degrees or two units greater than the hotspot value. The size of this constraint is a preference that you can change

in the *Rendering panel (page 3–826)* of the Preferences dialog. The falloff value can range from 0 to 10,000 units.

When the hotspot and falloff values are equal, the light casts a hard-edged beam.

To set the shape of the light beam:

1. Choose either Rectangle or Circle.

   The shape of the light's cone changes to reflect the shape you chose.

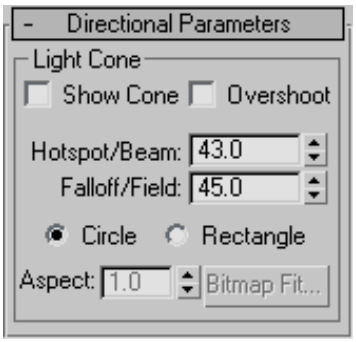
2. If you chose Circle, you are done. If you chose Rectangle, you can now adjust the *aspect ratio (page 3–911)* of the rectangular light using the spinner labeled Aspect.

To change a rectangular light's aspect ratio:

1. Choose Rectangle, as described above.
2. Change the Aspect value to the aspect ratio that you want.

   The Bitmap Fit button is another way to set the aspect ratio. It is mainly used with projections.

Interface



Light Cone group

These parameters control *hotspots (page 3–954)* and *falloff (page 3–954)* for spotlights.

Show Cone—Turns display of the cone on or off.

Note: The cone is always visible when a light is selected, so clearing this check box has no apparent effect until you deselect the light.

Overshoot—When *Overshoot (page 3–988)* is set, the light casts light in all directions. However, projections and shadows occur only within its falloff cone.

Hotspot/Beam—Adjusts the size of a light's cone. The Hotspot value is measured in 3ds Max units. Default=43.0.

Falloff/Field—Adjusts the size of a light's falloff. The Falloff value is measured in 3ds Max units. Default=45.0.

You can also adjust hotspot and falloff angles in a *Light viewport (page 2–1338)* (looking at the scene from the point of view of the spotlight).

Circle/Rectangle—Determine the shape of the falloff and hotspot areas. Set Circle when you want a standard, circular light. Set Rectangle when you want a rectangular beam of light, such as light cast through a window or doorway.

Aspect—Sets the *aspect ratio (page 3–911)* for the rectangular light beam. The Bitmap Fit button lets you make the aspect ratio match a specified bitmap. Default=1.0.

Bitmap Fit—If the light's projection aspect is rectangular, sets the aspect ratio to match a particular bitmap. This is useful when you are using the light as a *projector light (page 3–998)*.

Atmospheres and Effects for Lights

Modify panel > Select light object. > Atmospheres & Effects rollout

The Atmospheres & Effects rollout lets you assign, delete, and set up parameters for atmospheres and rendering effects associated with the light. This

rollout appears only in the Modify panel; it doesn't appear at creation time.

Adding an atmosphere or effect associates that atmosphere or effect with the light object. This rollout is a shortcut to either the *Environment panel (page 3–272)* or the *Effects panel (page 3–219)* on the Environment And Effects dialog.

Procedures

To add a new atmosphere or rendering effect:

1. In the Atmospheres & Effects rollout, click Add. This displays the *Add Atmosphere or Effect dialog (page 2–1351)*.
2. In the dialog, choose Atmosphere, Effect, or All. Choose New.
3. Choose an atmosphere or effect in the list, then click OK.

   This associates a new atmosphere or rendering effect with the light.

To add an existing atmosphere or rendering effect:

1. In the Atmospheres & Effects rollout, click Add. This displays the *Add Atmosphere or Effect dialog (page 2–1351)*.
2. In the dialog, choose Existing.
3. Choose an atmosphere or effect in the list, then click OK.

   This creates a duplicate atmosphere or rendering effect for the light. Its settings are initially identical to the atmosphere or effect you chose. You can adjust them using Setup.

To delete an atmosphere or rendering effect:

- In the Atmospheres & Effects rollout, choose the name of an atmosphere or effect in the list, then click Delete.

To set up the parameters for an atmosphere or rendering effect:

- In the Atmospheres & Effects rollout, choose the name of an atmosphere or effect in the list, then click Setup.

  If you chose an atmosphere, clicking Setup displays the *Environment panel (page 3–272)*. If you chose an effect, clicking Setup displays the *Effects panel (page 3–219)*.

Interface

These controls are displayed for all light types.



Add—Displays the *Add Atmosphere or Effect dialog (page 2–1351)*, which lets you add an atmosphere or a rendering effect to the light.

Delete—Deletes the atmosphere or effect you have selected in the list.

List of atmospheres and effects—Displays the names of all atmospheres or effects you have assigned to this light.

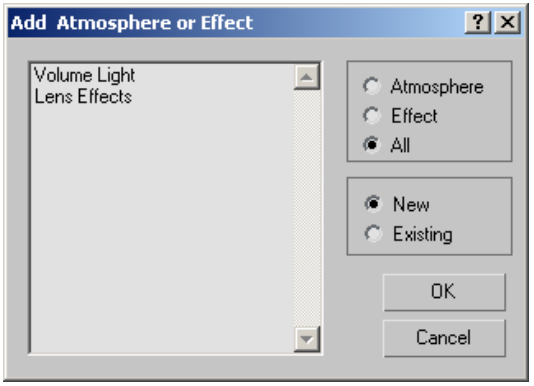
Setup—Lets you set up the atmosphere or rendering effect you have selected in the list. If the item is an atmosphere, clicking Setup displays the *Environment panel (page 3–272)*. If the item is an effect, clicking Setup displays the *Effects panel (page 3–219)*.

Add Atmosphere or Effect Dialog

Modify panel > Select light object. > Atmospheres & Effects rollout > Add button

The Add Atmosphere or Effect dialog lets you associate an atmosphere or a rendering effect with the light. The list shows either atmospheres, rendering effects, or both. It shows only atmospheres and effects that can be associated with light objects, or that use light objects as their apparatus.

Interface



List of atmospheres and effects—Displays the atmospheres or effects that you can associate with the light.

List filter group

These radio buttons choose what to show in the list.

Atmosphere—Lists only atmospheres.

Effect—Lists only rendering effects.

All—Lists both atmospheres and rendering effects.

New or existing group

These radio buttons choose between new or existing effects.

New—Lists only new atmospheres or effects.

Existing—Lists only atmospheres or effects that have been already assigned to the light.

Adding an existing atmosphere or effect creates a new atmosphere or effect whose settings are initially identical to the previous one.

Hair Light Attr(ibutes) Rollout

Add a supported light. > Select an object with the Hair And Fur modifier applied. > Render Settings > Effects panel > Select the light (in a viewport). > Hair And Fur rollout > Lighting group > Click Add Hair Properties. > Modify panel > Hair Light Attr rollout

Add a supported light and an object with the Hair And Fur modifier applied. > Rendering menu > Effects > Highlight the Hair And Fur render effect. > Hair And Fur rollout > Lighting group > Select the light (in a viewport). > Click Add Hair Properties. > Modify panel > Hair Light Attr rollout

The Hair Light Attr(ibutes) rollout lets you adjust properties for the hair shadow map generated by the light. This rollout appears for supported lights at render time. If you want to adjust a light's hair settings before render time, you can display it explicitly by following the steps shown in the path annotations at the beginning of this topic.

You specify which lights to use in the *Hair And Fur render effect (page 3–220)*. By default, *Use All Lights At Render Time* is on, so all supported lights in the scene can illuminate hair. Use this rollout to set hair-specific shadow-map values. Hair shadows ignore the light's ordinary shadow-map settings (for other 3ds Max objects).

Important: The settings on this rollout apply only to "buffer"-rendered hair (the default rendering method, set in the Hair And Fur render effect). The "geometry" and "mr prim" methods use the light's ordinary shadow settings.

Interface



Light hair—When on, the light can illuminate and cast shadows from the hair.

Note: The light must also be set to cast shadows (General Parameters rollout > Shadows > On).

Resolution—Sets the size (in pixels squared) of the shadow map that's computed for the light.

The shadow map size specifies the amount of subdivisions for the map. The higher the value, the more detailed the map will be, but the more memory is required.

Fuzz—Sets the softness of the shadow edges. See *Sample Range (page 3–1004)*.



Left: Fuzz=0.0

Center: Fuzz=5.0

Right: Fuzz=10.0

The arrow points to the shadow cast by the hair on the underlying surface.

Additional Rollouts for Photometric Lights

Intensity/Color/Distribution Rollout

Create panel > Lights > Create a Photometric light. > Intensity/Color/Distribution rollout

The Intensity/Color/Distribution rollout lets you set the type of distribution for *photometric lights (page 2–1301)*. You can also define the color and intensity of lights.

See also

Lights (page 2–1272)

Photometric Lights (page 2–1301)

Isotropic Light Distribution (Photometric Lights) (page 2–1323)

Spotlight Distribution (Photometric Lights) (page 2–1324)

Diffuse Distribution (Photometric Lights) (page 2–1323)

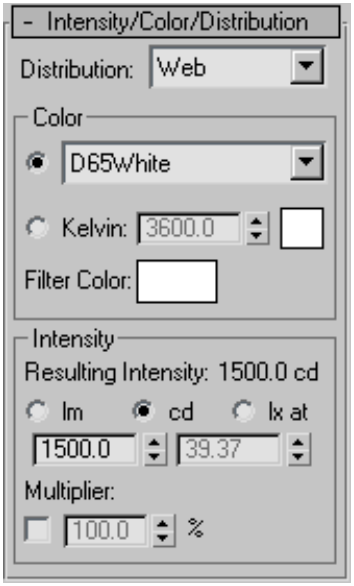
Web Distribution (Photometric Lights) (page 2–1325)

Procedure

To change the distribution type of a photometric light:

1. Select the light.
2. In the Modify panel > Intensity/Color/Distribution rollout, choose a type from the Distribution list.

Interface



Distribution—Describes the directional distribution of the light emitted by a source. There are several different distributions:

- *Isotropic (page 2–1323)*
- *Spotlight (page 2–1324)*
- *Web (page 2–1325)*
- *Diffuse (page 2–1323)*

Note: Target Point and Free Point lights can have either Isotropic, Spotlight, or Web distribution. All other photometric lights use either Web or Diffuse distribution.

Color group

Light—Pick a common lamp specification to approximate the spectral character of your light. The color swatch next to the Kelvin parameter is updated to reflect the light you select.

Kelvin—Set the color of the light by adjusting the color temperature spinners. The color temperature is displayed in degrees Kelvin. The corresponding color is visible in the color swatch next to the temperature spinners.

Filter—Use a color filter to simulate the effect of a color filter placed over the light source. For example, a red filter over a white light source casts red light. Set the filter color by clicking the color swatch to display the *Color Selector (page 1–161)*. Default=white (RGB=255,255,255; HSV=0,0,255).

Intensity group

These controls specify the strength or brightness of photometric lights in physically based quantities.

You set the intensity of a light source using one of the following units:

- **lm (lumen)—**Measures the overall output power of the light (*luminous flux (page 3–955)*). A 100-watt general purpose light bulb has a luminous flux of about 1750 lm.
- **cd (candela)—**Measures the maximum *luminous intensity (page 3–965)* of the light, usually along the direction of aim. a 100-watt general purpose light bulb has a luminous intensity of about 139 cd.
- **lx at (lux)—**Measures *illuminance (page 3–955)* caused by the light shining on a surface at a certain distance and facing in the direction of the source. The lux is the International Scene unit, which is equivalent to 1 lumen per square meter. The AS unit for illuminance is the footcandle (fc) which is equivalent to 1 lumen per square foot. To convert from footcandles to lux, multiply by 10.76. For example, to specify an illuminance of 35 fc, set illuminance to 376.6 lx.

  To specify the illuminance of a light, you must set the lx value, and enter the distance at which that illuminance is measured.

Note: You can obtain these values directly from lighting manufacturers. A table of some common

lamp types is provided in *Photometric Lights: Common Lamp Values (page 2–1329)*.

Linear Light Parameters Rollout

Create panel > Lights > Create a photometric linear light. > Linear Light Parameters rollout

Select a linear light. > Modify panel > Linear Light Parameters rollout

The Linear Light parameters rollout lets you set the length of your photometric linear light.

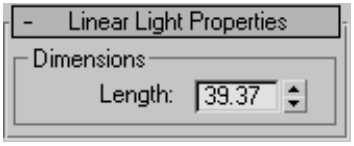
See also

Photometric Lights (page 2–1301)

Target Linear Light (Photometric) (page 2–1305)

Free Linear Light (Photometric) (page 2–1307)

Interface



Length—Sets the length of your linear light.

Area Light Parameters Rollout

Create panel > Lights > Create a photometric area light. > Area Light Parameters rollout

The Area Light parameters rollout lets you set the dimensions of your photometric area light.

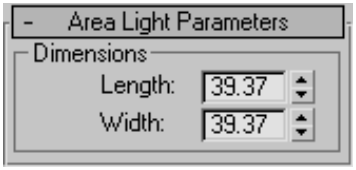
See also

Photometric Lights (page 2–1301)

Target Area Light (Photometric) (page 2–1307)

Free Area Light (Photometric) (page 2–1309)

Interface



Length—Sets the length of your area light.

Width—Sets the width of your area light.

Area Light Sampling Rollout

Select a linear or photometric area light. > Modify panel > Area Light Sampling rollout

The Area Light Sampling rollout lets you control the number of samples to use for the minimum subdivision size conversion. The values in this rollout affect rendering performed by the *default scanline renderer (page 3–38)* and the *mental ray renderer (page 3–78)*. It's also used for Radiosity in order to compute *Shoot Direct Lights (page 3–69)* for area and linear lights. The number of samples are divided into the area of an area light or into the length of a linear light.

Note: This rollout appears only on the Modify panel.

See also

Photometric Lights (page 2–1301)

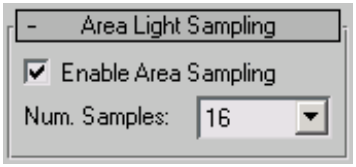
Target Linear Light (Photometric) (page 2–1305)

Free Linear Light (Photometric) (page 2–1307)

Target Area Light (Photometric) (page 2–1307)

Free Area Light (Photometric) (page 2–1309)

Interface



Enable Area Sampling—Toggles area sampling. Default=on.

Num. Samples—The number of samples to use for subdivision conversion. Default=16.

Web Parameters Rollout

Create panel > Lights > Photometric Lights > Create a photometric light. > Intensity/Color/Distribution Parameters rollout > Choose Web distribution type.

The Web Parameters rollout is displayed on the Modify panel when you create or select a photometric light with a *web distribution (page 2–1325)*. Use these parameters to select a photometric web file and to adjust the orientation of the web.

3ds Max can use the *IES (page 2–1328)*, *CIBSE (page 3–921)*, or *LTLI (page 3–964)* photometric web formatting.

See also

Target Point Light (Photometric) (page 2–1303)

Free Point Light (Photometric) (page 2–1304)

Target Linear Light (Photometric) (page 2–1305)

Free Linear Light (Photometric) (page 2–1307)

Target Area Light (Photometric) (page 2–1307)

Free Area Light (Photometric) (page 2–1309)

Procedures

To select a photometric web file:

1. On the Web Parameters rollout for a selected light, click the Web File button.

   The Open A Photometric Web dialog is displayed.

2. Navigate to the location of the file you want to import. Select the file and click Open.

   If you are modifying a light, the light's icon is replaced by a wireframe outline of the photometric web. If you are creating a light, the photometric web icon appears when you drag your cursor in the viewport to create the light.

   Note: You can obtain photometric web files from various lighting manufacturers. You can use the Content Browser to download IES files or use IES files available in the Lightscape Libraries provided with Lightscape software.

To change the orientation of the photometric web:

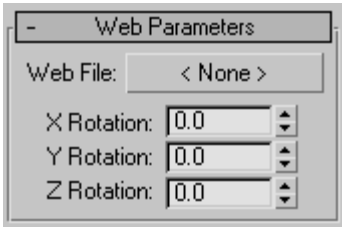
- Use the rotation spinners to change the X, Y, or Z rotation or the photometric web.

  Note: These parameters cannot be animated.

  To change the direction of a light, you can also use the Rotation transform or for target point lights, move the light's target point.

  Tip: To see the shape of the photometric web, you can make the light icon larger by increasing the Non Scaling Object size in Tools menu > Options dialog > Viewports panel > Viewport Parameters.

Interface



Web File—Selects an IES file to use as a photometric web. The default web is a diffuse distribution shining from one edge.

X Rotation—Rotates the photometric web about the X axis. The center of rotation is the photometric center of the web. Range=-180 degrees to 180 degrees.

Y Rotation—Rotates the photometric web about the Y axis. The center of rotation is the photometric center of the web. Range=-180 degrees to 180 degrees.

Z Rotation—Rotates the photometric web about the Z axis. The center of rotation is the photometric center of the web. Range=-180 degrees to 180 degrees.

Photometric web icon—Shows a three-dimensional representation of the light distribution. See *Web Distribution (page 2–1326)*.

Rollouts for Specific Shadow Types

Advanced Ray-Traced Parameters Rollout

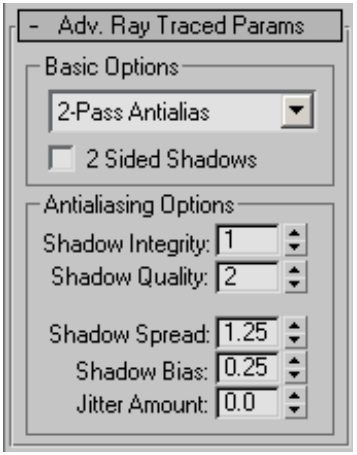
Create panel > Lights > Create or select a light. > General Parameters rollout > Shadow type > Advanced Ray-traced Shadows

Advanced Ray-traced shadows are similar to *ray-traced shadows (page 3–1000)*, however they give you more control over shadow behavior. Additional controls are available in the *Optimizations rollout (page 2–1361)*.



Advanced ray-traced shadows cast by an area light.

Interface



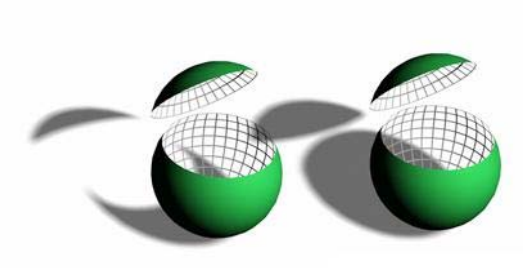
Basic Options group

Mode—Selects the type of raytracing for generating shadows:

- **Simple—**Casts a single ray of light toward the surface. No *antialiasing (page 3–907)* is performed.
- **Antialias 1-Pass—**Casts a bundle of rays. The same number of rays is cast from each illuminated surface. The number of rays is set using the Pass 1 Quality spinner.
- **Antialias 2-Pass—**Casts two bundles of rays. The first batch of rays determines if the point in question is fully illuminated, shadowed, or in the penumbra (soft area) of the shadow. If the point is in the penumbra, a secondary

batch of rays is cast to further refine the edge. The number of initial rays is specified using the Pass 1 Quality spinner. The number of secondary rays is specified using the Pass 2 Quality spinner.

2-Sided Shadows—When on, backfaces are not ignored when calculating shadows. Objects seen from the inside are not lit by lights from the outside. This costs a bit more render time. When off, backfaces are ignored. Rendering is quicker, but outside lights illuminate object interiors.



The faces inside the sliced sphere do not cast shadows if 2-Sided Shadows is not selected.

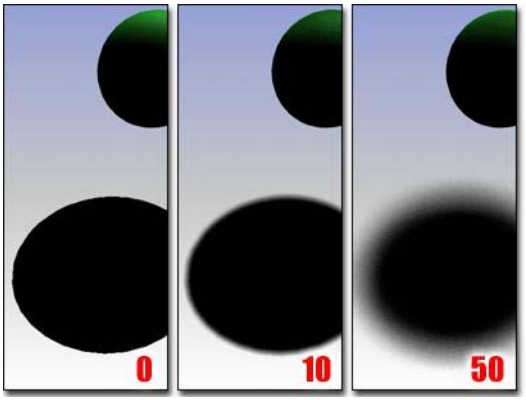
Antialiasing Options group

Shadow Integrity—The number of rays cast from an illuminated surface.

Shadow Quality—The number of secondary rays cast from an illuminated surface.

Shadow Spread—The radius, in pixels, to blur the antialiased edge.

Note: As this value increases, so does the quality of the blur. However, increasing this value also increases the likelihood of missing small objects. To avoid this problem, increase the value of Pass 1 Quality.



Increasing the Shadow Spread value softens shadow edges.

Shadow Bias—The minimum distance from the point being shaded that an object must be to cast a shadow. This prevents blurred shadows from affecting surfaces they shouldn't.

Note: As you increase the blur value, you should also increase the bias.

Jitter Amount—Adds randomness to the ray positions. The rays are initially in a very regular pattern, which can show up in the blurry part of the shadow as regular artifacts. Jittering will convert these artifacts to noise, which is generally less noticeable to the eye. Recommended values are 0.25 to 1.0. However, very blurry shadows will require more jitter.

Area Shadows Rollout

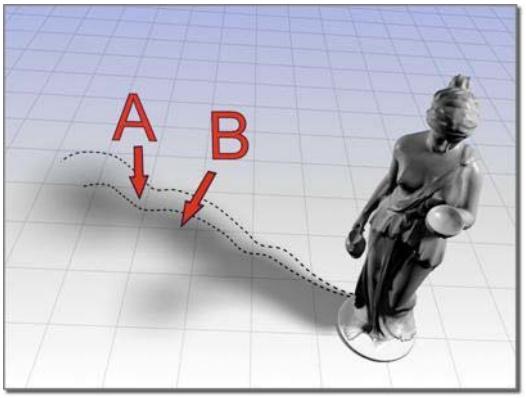
Create panel > Lights > Create or select a light. > General Parameters rollout > Shadow type > Area Shadows rollout

The Area Shadow generator can be applied to any light type to achieve the effect of an area shadow effect. In order to create an area shadow, the user needs to specify the dimensions of a virtual light created to "fake" an area shadow.

Note: When using area lights, the user should try to make the properties of the light match the

properties in the Area Light Dimensions group of the Area Shadows rollout.

Tip: Area shadows can take a fair amount of time to render. If you want to create a quick test (or draft) rendering, you can use the Area/Linear Lights as Point Lights toggle in the *Common Parameters rollout (page 3–27)* of the Render Scene dialog to speed up your rendering. When this toggle is on, the shadows are processed as if the light object is a point source.



Area shadows create a soft edge that becomes more noticeable as the distance between the object and the shadow increases.

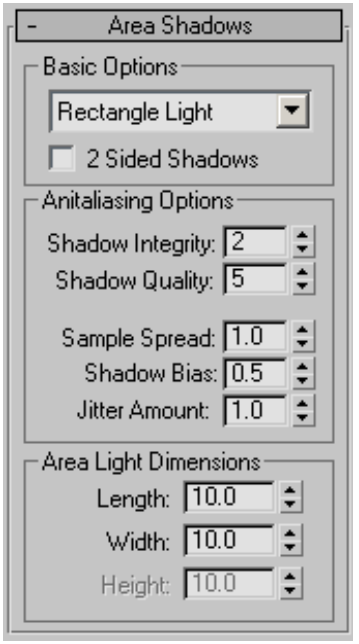
A: Penumbra (soft area)

B: Shadow



Area shadow cast by a point light

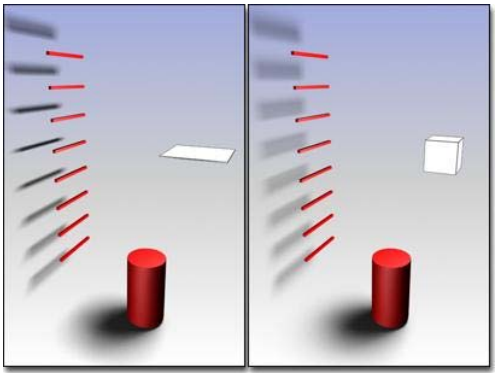
Interface



Basic Options group

Mode—Selects the way the area shadows are generated:

- **Simple**—Casts a single ray from the light toward the surface. No *antialiasing (page 3–907)* or area light calculation is performed.
- **Rectangle Light**—Casts rays from the light in a rectangular array.
- **Disc Light**—Casts rays from the light in a circular array.
- **Box Light**—Casts rays from the light as if it were a box.
- **Sphere Light**—Casts rays from the light as if it were a sphere.

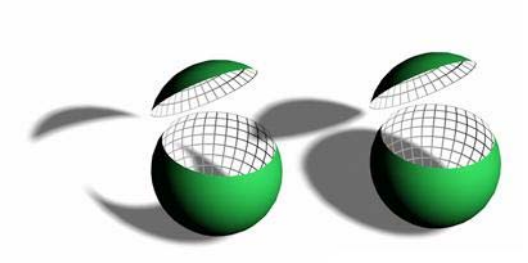


The shape of the area shadow array affects how shadows are cast.

Left: Rectangle

Right: Box

2-Sided Shadows—When on, backfaces are not ignored when calculating shadows. Objects seen from the inside are not lit by lights from the outside. This costs a bit more render time. When off, backfaces are ignored. Rendering is quicker, but outside lights illuminate object interiors.



The faces inside the sliced sphere do not cast shadows if 2-Sided Shadows is off.

Antialiasing Options group

Shadow Integrity—Sets the number of rays in the initial bundle of rays cast. These rays are projected from every surface that receives light from the light source.

The number of rays is as follows:

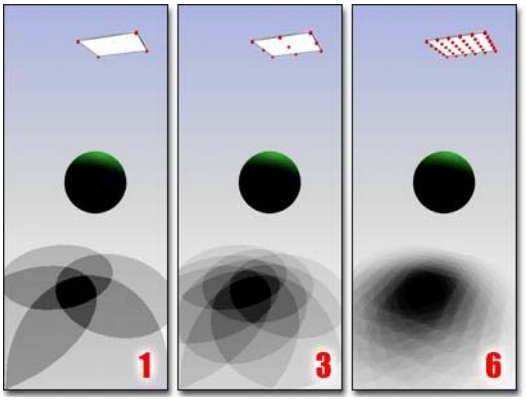
1=4 rays

2=5 rays

3 to N=NxN rays.

For example, setting Shadow Integrity to 5 generates 25 rays.

This is the primary control for "finding" small objects and thin spaces between objects. If the shadows are missing a small object in your scene, try increasing Shadow Integrity one step at a time. Also, if the penumbra (soft area) is blotchy, try increasing this setting.



Increasing the Shadow Integrity value creates more accurate shadow contour and detail.

Shadow Quality—Sets the total number of rays cast in the penumbra (soft area) region, including the rays fired in the first pass.

These rays are projected from every point that is in the penumbra, or antialiased edge of the shadow, to smooth it out.

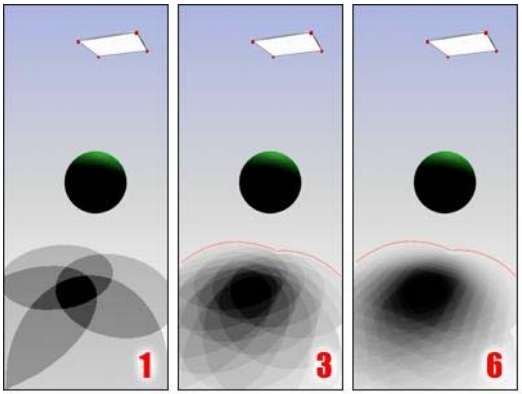
The number of rays is as follows:

- 2=5 rays
- 3 to N = NxN

For example, setting Shadow Quality to 5 generates 25 rays.

The Shadow Quality value should always be greater than the Shadow Integrity value. This is because the software overlays the secondary rays atop the first pass's rays using the same algorithm.

Increase Shadow Quality to fix banding in the penumbra, and eliminate the noise pattern from jittering.



Increasing the Shadow Quality value produces a more accurate penumbra (soft area) within the contour defined by the Shadow Integrity value.

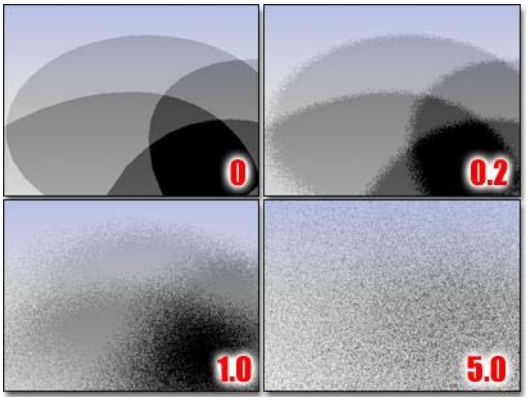
Sample Spread—The radius, in pixels, to blur the antialiased edge.

Note: As this value increases, so does the quality of the blur. However, increasing this value also increases the likelihood of missing small objects. To avoid this problem, increase the value of Shadow Integrity.

Shadow Bias—The minimum distance that an object must be from the point being shaded in order to cast a shadow. This prevents blurred shadows from affecting surfaces they shouldn't.

Note: As you increase the blur value, you should also increase the bias.

Jitter Amount—Adds randomness to the ray positions. The rays are initially in a very regular pattern, which can show up in the blurry part of the shadow as regular artifacts. Jittering will convert these artifacts to noise, which is generally less noticeable to the eye. Recommended values are 0.25 to 1.0. However, very blurry shadows require more jitter.



Increasing jitter blends the individual shadow samples.

Area Light Dimensions group

The software uses these dimensions to compute the area shadowing. They do not affect the actual light object.

Length—Sets the length of the area shadow.

Width—Sets the width of the area shadow.

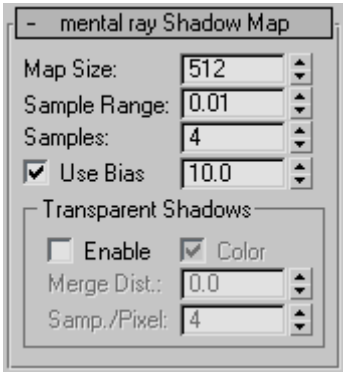
Height—Sets the height of the area shadow.

mental ray Shadow Map Rollout

Create panel > Lights > Create a light. > General Parameters rollout > Select "mental ray Shadow Map shadows." > mental ray Shadow Map rollout

Choosing "mental ray Shadow Map" as the shadow type tells the *mental ray renderer (page 3–78)* to generate shadows using the mental ray shadow-map algorithm. If you render with the *default scanline renderer (page 3–38)* instead, no shadows appear in the rendering.

Interface



Map Size—Sets the resolution of the shadow map. The size of the map is the square of this value. Greater resolutions require more time to process, but produce more accurate shadows. Default=512.

Sample Range—When greater than zero, generates soft-edged shadows. This value specifies the area of the map to soften, by removing portions of the map as specified by Samples. Default=0.0.

If you set Sample Range to be greater than zero, you must also set Samples to be greater than zero, to obtain a soft shadow effect.

Directional lights require Sample Range to have a greater value than spotlights require.

Samples—Sets the number of samples to remove from a shadow map when generating soft shadows. Default=1.

Use Bias—When on, changes the *map bias (page 3–966)*. Increasing the value moves the shadow farther away from the shadow casting object. Default=10.

Transparent Shadows group

Enable—When on, shadow maps are saved with multiple Z-layers, and can have transparency. Default=off.

Color—When on, surface color affects the color of the shadow. Default=on.

Turning off Color saves memory at rendering time.

Merge Dist.—The minimum distance between two surfaces for them to be considered "distinct." If two surfaces are closer than this value, the shadow map treats them as a single surface. When set to 0.0, the mental ray renderer automatically calculates a distance value to use. Default=0.0 (automatic).

Larger Merge Distance values reduce memory consumption, but can reduce shadow quality. Low Merge Distance values increase memory consumption and slow down rendering speed.

Samp./Pixel—The number of samples used to generate a pixel in the shadow map. Higher values increase the quality and detail of the shadow, at a cost of render time. Default=5.

If a mapped shadow appears to be *aliased (page 3–907)*, increase the value of Samples/Pixel. This setting is especially useful when shadows are cast by finely detailed geometry.

Optimizations Rollout

Create panel > Lights > Create a light. > General Parameters rollout > Select Advanced Ray-traced Shadows or Area Shadows. > Optimizations rollout

The Optimizations rollout provides additional controls for *Advanced Ray-traced Shadow (page 3–1000)* and *Area Shadow (page 3–911)* generation.

See also

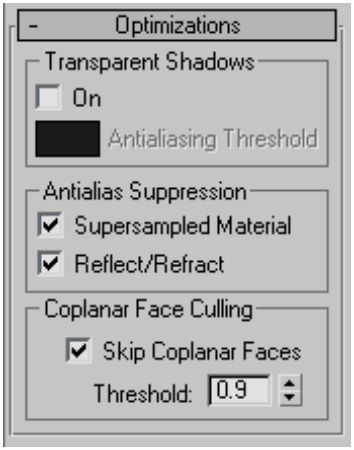
General Lighting Parameters (page 2–1331)

Shadow Parameters (page 2–1337)

Area Shadows Rollout (page 2–1357)

Advanced Ray-Traced Parameters Rollout (page 2–1356)

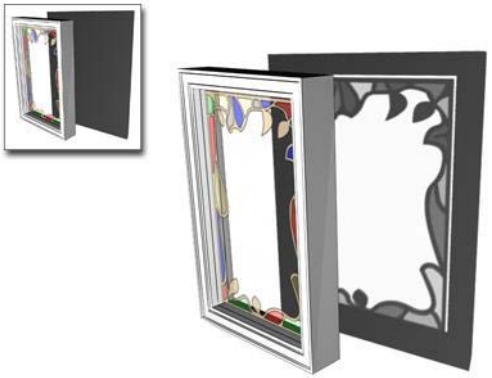
Interface



Transparent Shadows group

On—When checked, transparent surfaces will cast a colored shadow. Otherwise, all shadows are black.

Note: Shadows will generate faster with this value turned off.



A stained glass appears solid with the transparency optimization turned off.

Antialiasing Threshold—The maximum color difference allowed between transparent object samples before *antialiasing (page 3–907)* is triggered. Increasing the value of this color will make the shadow less sensitive to aliasing artifacts and improve speed, decreasing the value will increase the sensitivity, improving quality.

Antialias Suppression group

Supersampled Material—When on, only pass 1 is used during 2-pass antialiasing when shading a *supersampled (page 3–1018)* material.

Note: When off, rendering time can increase without resulting in a better image.

Reflect/Refract—When on only pass 1 is used during 2-pass antialiasing when shading reflections or refractions.

Note: When off, rendering time can increase without resulting in a better image.

Coplanar Face Culling group

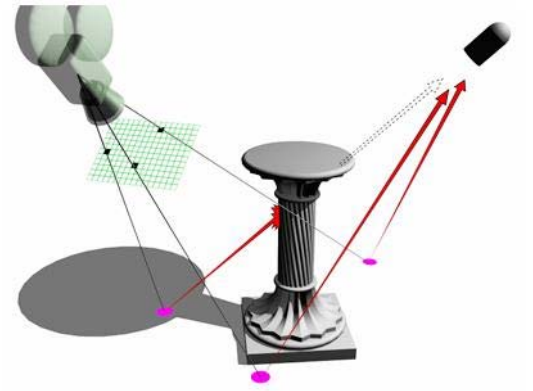
Skip Coplanar Faces—Prevents adjacent faces from shadowing each other. This is of particular concern at the terminator on curved surfaces such as spheres.

Threshold—The angle between adjacent faces.

Range = 0.0 (perpendicular) to 1.0 (parallel).

Ray-Traced Shadow Parameters Rollout

Create a light. > General Parameters rollout > Ray Traced Shadows chosen > Ray Traced Shadow Params rollout

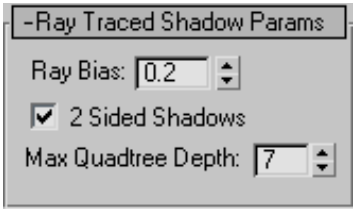


Ray-traced shadows

The Ray-Traced Shadow Parameters rollout appears when you have chosen raytracing as the shadow-generation technique for a light. You select this in the *General Parameters rollout (page 2–1331)*.

For greater control of ray-traced shadows, use *Advanced Ray-Traced Shadows (page 2–1356)*.

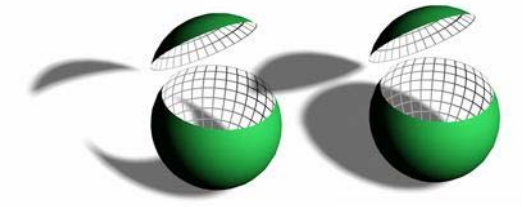
Interface



Bias—*Ray-Trace Bias (page 3–1000)* moves the shadow toward or away from the shadow-casting object (or objects).

If the Bias value is too low, shadows can "leak" through places they shouldn't, producing moire patterns or making out-of-place dark areas on meshes. If Bias is too high, shadows can "detach" from an object. If the Bias value is too extreme in either direction, shadows might not be rendered at all.

2-Sided Shadows—When on, backfaces are not ignored when calculating shadows. Objects seen from the inside are not lit by lights from the outside. This costs a bit more render time. When off, backfaces are ignored. Rendering is quicker, but outside lights illuminate object interiors. Default=on.



The faces inside the sliced sphere do not cast shadows if 2-Sided Shadows is not selected.

Max Quadtree Depth—Adjusts the depth of the *quadtree (page 3–999)* used by the ray-tracer. Greater quadtree depth values can improve ray-tracing time at the cost of memory use. However, there is a depth value where the performance improvement is offset by the time it takes to generate the quadtree itself. This depends on the geometry of the scene. Default=7.

Tip: Omni lights can generate up to six quadtrees, so they generate ray-traced shadows more slowly than spotlights. Avoid using ray-traced shadows with omni lights unless your scene requires this.

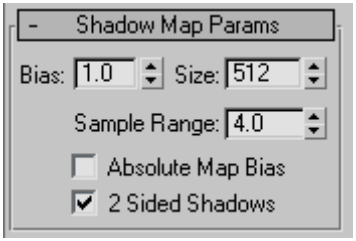
Shadow Map Parameters Rollout

Create a light. > General Parameters rollout > Shadow Map chosen > Shadow Map Params rollout

The Shadow Map Parameters rollout is displayed when you have chosen shadow mapping as the

shadow-generation technique for a light. You select this in the *General Parameters rollout (page 2–1331)*.

Interface



Bias—*Map Bias (page 3–966)* moves the shadow toward or away from the shadow-casting object (or objects).



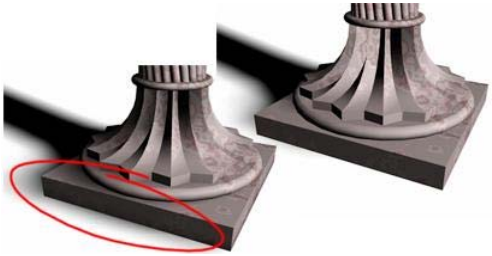
Left: Default shadows

Right: Increasing the Bias value separates the shadow from the object.

If the Bias value is too low, shadows can "leak" through places they shouldn't, produce moire patterns or making out-of-place dark areas on meshes. If Bias is too high, shadows can "detach" from an object. If the Bias value is too extreme in either direction, shadows might not be rendered at all.

This value depends on whether Absolute Map Bias is on or off:

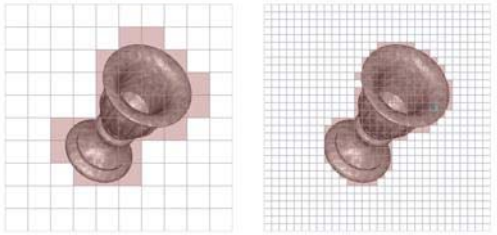
- When Absolute is off (the default), Bias is calculated based on the scene extents, and then normalized to one. This provides similar default shadow results, regardless of scene size. User adjustments to Bias are typically low decimal values near 1.0 (for example, 1.2).
- When Absolute is on, Bias is a value in 3ds Max units. User adjustments to Bias depend on the size of the scene, and can range from values close to zero to values in the hundreds (see the Tip at the end of this topic).



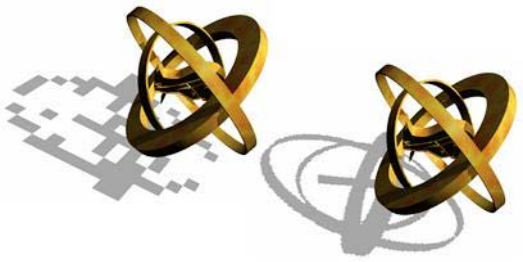
Left: Too small a Bias value causes shadow "leaks."

Right: Increasing the Bias value fixes the problem.

Size—Sets the size (in pixels squared) of the shadow map that's computed for the light.



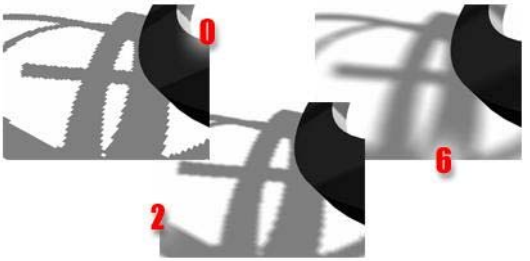
The shadow map size specifies the amount of subdivisions for the map. The greater the value, the more detailed the map will be.



Left: Size set to 32.

Right: Size set to 256.

Sample Range—The *sample range (page 3–1004)* determines how much area within the shadow is averaged. This affects how soft the edge of the shadow is. Range=0.01 to 50.0.



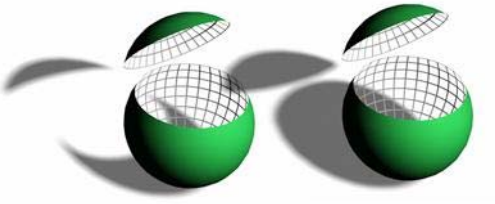
Increasing the Sample Range blends the shadow edges and creates a smooth effect, hiding the granularity of the map.

Absolute Map Bias—When on, the bias for the shadow map is not normalized, but is instead based on a fixed scale expressed in 3ds Max units. This value does not change during an animation. You must choose the value, based on the size of the scene extents.

When off, the bias is computed relative to the rest of the scene, and then normalized to 1.0. This provides a common starting bias value in scenes of any size. If the scene extents change, this internal normalization can vary from frame to frame. Default=off.

Tip: Leaving Absolute Map Bias off gives good results in most situations, because the bias is internally balanced to match the scene size. In animations, however, if moving objects cause a large change in the scene extents (or if objects are unhidden, and so on), the normalized bias value might become inappropriate, causing shadows to flicker or disappear. If this happens, turn on Absolute Map Bias. You will have to set the Bias control to a value appropriate for the scene. As a rule of thumb, try a Bias value that is the distance between the light and the target object, divided by 100.

2-Sided Shadows—When on, backfaces are not ignored when calculating shadows. Objects seen from the inside are not lit by lights from the outside. When off, backfaces are ignored, which can cause outside lights to illuminate object interiors. Default=on.



The faces inside the sliced sphere do not cast shadows if 2-Sided Shadows is not selected.

Cameras

Create panel > Cameras

Create menu > Cameras

Cameras present a scene from a particular point of view. Camera objects simulate still-image, motion picture, or video cameras in the real world.

With a *Camera viewport (page 3–745)* you can adjust the camera as if you were looking through its lens. Camera viewports can be useful for editing geometry as well as setting up a scene for rendering. Multiple cameras can give different views of the same scene.

The *Camera Correction modifier (page 2–1392)* lets you correct a camera view to 2-point perspective, in which vertical lines remain vertical.

If you want to animate the point of view, you can create a camera and animate its position. For example, you might want to fly over a landscape or walk through a building. You can animate other

camera parameters as well. For example, you can animate the camera's field of view to give the effect of zooming in on a scene.

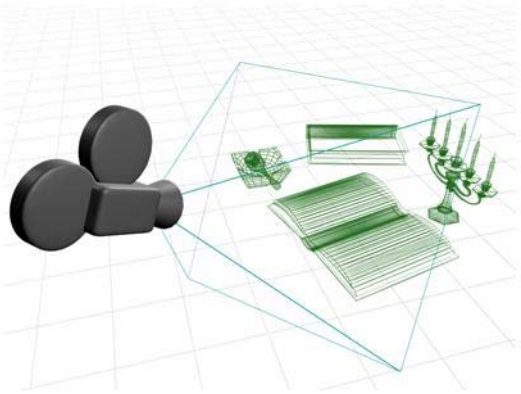
The Display panel's *Hide By Category (page 1–52)* rollout has a toggle that lets you turn the display of camera objects on and off.

A convenient way to control the display of camera objects is to create them on a separate *layer (page 3–655)*. You can hide them quickly by turning off the layer.

Tip: The *Camera Match utility (page 2–1387)* allows you to start with a background photograph and create a camera object that has the same point of view. This is useful for site-specific scenes.

There are two kinds of camera objects:

-  *Target cameras (page 2–1371)* view the area around a target object. When you create a target camera, you see a two-part icon representing the camera and its target (a white box). The camera and the camera target can be animated independently, so target cameras are easier to use when the camera does not move along a path.
-  *Free cameras (page 2–1370)* view the area in the direction the camera is aimed. When you create a free camera, you see a singe icon representing the camera and its field of view. The camera icon appears the same as a target camera icon, but there is no separate target icon to animate. Free cameras are easier to use when the camera's position is animated along a path.



An example of a camera in a scene.



The result after rendering through the camera.

You can create cameras from the Create menu > Cameras submenu, or by clicking the Cameras button on the Create panel. You can also create a camera by activating a Perspective viewport, and then choosing Views menu > Create Camera From View.

After you have created a camera, you can change viewports to display the camera's point of view. While a camera viewport is active, the navigation buttons change to *camera navigation buttons (page 3–745)*. You use the Modify panel in conjunction with a camera viewport to change the camera's settings.

While you use the navigation controls for a camera viewport, you can constrain Truck, Pan, and Orbit

movement to be vertical or horizontal only with the Shift key.

You can move a selected camera so its view matches that of a Perspective, Spotlight, or another Camera view.

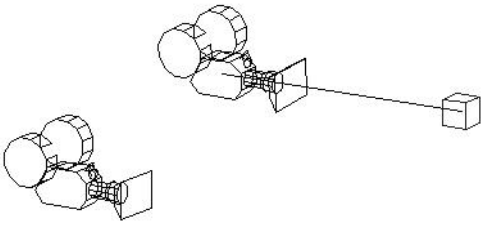
Choosing a Camera for Vertical Views

If you need an animated camera to look vertically upward or downward, use a free camera. If you use a target camera you might run into a problem of unexpected movement. The program constrains a target camera's up-vector (its local positive Y axis) to be as close as possible to the world positive Z axis. This is no problem when you are working with a static camera. However, if you animate the camera and put it in a nearly vertical position, either up or down, the program flips the Camera view to prevent the up-vector from becoming undefined. This creates sudden changes of view.

Camera Object Icons

Camera objects are visible in viewports unless you choose not to display them. However, the geometry that appears in the viewport is only an icon meant to show you where the camera is located and how it is oriented.

Target cameras (page 2–1371) create a double icon, representing the camera (a blue box intersecting a blue triangle) and the camera target (a blue box). *Free cameras (page 2–1370)* create a single icon, representing the camera and its field of view.



A free camera has no target. A target camera has a target sub-object.

You cannot shade camera objects. However, you can render their icons using Animation menu > Make Preview and turning on Cameras in the Display In Preview group.

The display of camera object icons is not scaled when you change the scale of the viewport. When you zoom in on a camera, for example, the icon size does not change. To change the size of camera object icons, you can use the *Viewports panel (page 3–821)* of the Preferences dialog, and change the value of Non-Scaling Object Size.

Scale transforms have the following effects on a camera object:

- Uniform Scale has no effect on a target camera, but does change the free camera's Target Distance setting.
- Non-Uniform Scale and Squash change the size and shape of the free camera's FOV cone. You see the effect in the viewport, but the camera's parameters do not update. Non-Uniform Scale and Squash will change the size and shape of a target camera's icon, but have no visible effect in the viewport.

mental ray Camera Shaders

When you use the *mental ray renderer (page 3–78)*, you can apply shaders to the camera used to render the scene. Specifically, you can assign shaders to modify the camera's lens, its output, or

its volume (effectively making a volume out of the entire scene).



A dimly lit scene



The scene's colors desaturated using the Night shader

Lens shader: lume Night shader with Multiplier set to 0.5

You assign camera shaders using the Render Scene dialog's *Camera Effects rollout (page 3–101)* while the mental ray renderer is active.

Note: No camera output shaders are provided with 3ds Max. You might have access to light map shaders if you have obtained them from other shader libraries or custom shader code.

See also

Common Camera Parameters (page 2–1373)

Characteristics of Cameras (page 2–1372)

Using Transforms to Aim a Camera (page 2–1379)

Using Clipping Planes to Exclude Geometry (page 2–1379)

Using the Horizon to Match Perspective (page 2–1380)

Animating Cameras (page 2–1381)

Procedures

To render a scene using a camera:

1. Create the camera and aim it at the geometry you want to be the subject of your scene. To aim a target camera, drag the target in the direction you want the camera to look. To aim a free camera, rotate and move the camera icon.
2. With one camera selected, or if only one exists in the scene, set a Camera viewport for that camera by activating the viewport, then press C. If multiple cameras exist and none or more than one are selected, the software prompts you to choose which camera to use.

   You can also change to a Camera viewport by right-clicking the viewport label and then selecting Views and the camera of choice.
3. Adjust the camera's position, rotation, and parameters using the Camera viewport's navigation controls. Simply activate the viewport, then use the Truck, Orbit, and Dolly Camera buttons. Alternately you can select the camera components in another viewport and use the move or rotate icons.

   If you do this while the *Auto Key button (page 3–717)* is on, you animate the camera.
4. Render the camera viewport.

To change a viewport to a Camera view:

1. Right-click the viewport label.

   The Viewport Properties menu is displayed.
2. Choose Views.

The name of each camera is displayed at the top of the Views submenu.

3. Choose the name of the camera you want.

   The viewport now shows the camera's point of view.

   The default keyboard shortcut for camera viewports is **C**.

   Making a camera viewport active does not automatically select the camera. To adjust a camera by using its viewport and the Modify panel at the same time, select the camera and then make the Camera viewport active.

As in other viewports, in Camera viewports you can opt to see a display of safe frame areas to help you compose the final rendered output.

To control the display of camera objects, do one of the following:

-  Go to the Display panel and in the Hide By Category rollout, turn Cameras on or off.

  Choose Views > Hide > Hide By Category, and toggle the menu item Hide Cameras.

- Choose Tools menu > Display Floater, and on the Object Level tab turn Cameras on or off.

  Cameras appear in viewports if Cameras is off; if Cameras is on, they don't appear.

  When camera icons are displayed, the *Zoom Extents commands (page 3–737)* include them in views. When camera icons are not displayed, the Zoom Extents commands ignore them.

To change the display size of camera icons:

- Choose Customize > Preferences > Viewports, and set Non-Scaling Object Size (default=1.0 in current units).

  Note: This also changes the size of light icons, helper objects, and other non-scaling objects in the scene.

To use the Modify panel in conjunction with a Camera viewport:

1. Select the camera in any viewport.
2. Right-click the Camera viewport to activate the viewport without deselecting the camera.

   The Camera viewport becomes active, but the camera is still selected in the other viewports.

3.  Adjust the camera using its Parameters rollout in the Modify panel and the navigation buttons.

   The Camera viewport updates as the parameters are changed.

To constrain Pan and Orbit to be vertical or horizontal:

- Hold down Shift as you drag in the viewport.

  The initial direction of the drag sets the constraint. If you drag vertically at first, the pan or orbit is constrained to be vertical; if you drag horizontally at first, the constraint is horizontal.

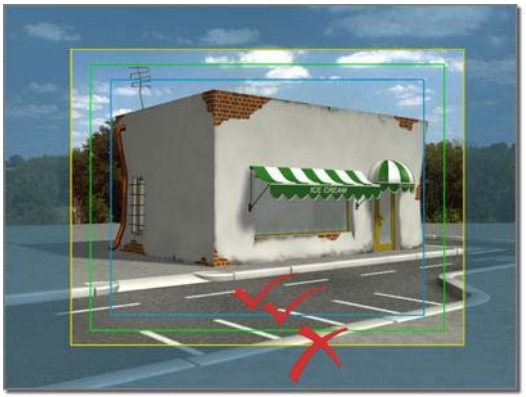
The Zoom Extents All flyout and the Min/Max toggles remain visible. These controls aren't specific to camera views. Clicking Zoom Extents All affects other kinds of viewports, but does not affect Camera viewports.

To see the safe frame:

- Right-click the viewport label and choose Show Safe Frame.

  The *safe frames (page 3–1030)* are displayed in three concentric boxes. The outermost safe frame matches the render output resolution.

  The *safe frame (page 3–1030)* matches the render output resolution.



Boxes in the viewport indicate safe frames.

To match a camera to a viewport:

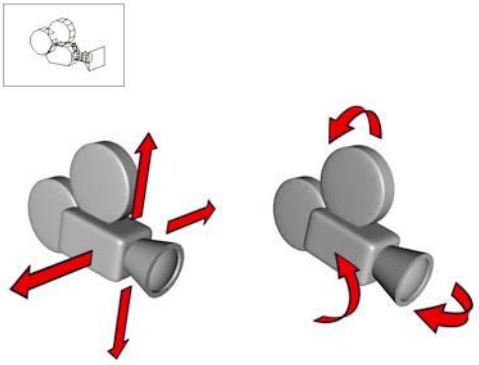
1. (Optional) Select a camera.
2. Activate a Perspective viewport.
3. If no camera was selected, 3ds Max creates a new target camera whose field of view matches the viewport. If you first selected a camera, the camera is moved to match the Perspective view. 3ds Max also changes the viewport to a camera viewport for the camera object, and makes the camera the currently selected object.

Free Camera

Create panel > Cameras > Free

Create menu > Cameras > Free Camera

Free cameras view the area in the direction where the camera is aimed. Unlike target cameras, which have two independent icons for the target and the camera, free cameras are represented by a single icon, making them easier to animate. Free cameras can be used when the camera's position is animated along a *trajectory (page 3–1025)*, as in a walkthrough of a building or when the camera is attached to a moving vehicle. The free camera can bank as it travels along the path. If the camera needs to be directly overhead in a scene, use a free camera to prevent it from spinning.



A free camera can move and be oriented without restrictions.

Initial Direction of a Free Camera

A free camera's initial direction is along the negative Z axis of the active construction grid of the viewport you click.

In other words, if you click in an orthogonal viewport, the initial camera direction is directly away from you. Clicking the Top viewport aims the camera downward, clicking the Front viewport aims the camera at the scene from the front, and so on.

Clicking in a Perspective, User, Light, or Camera viewport aims the free camera downward, along the negative Z axis of the World Coordinate System.

Because the camera is created on the active construction plane, where you also create geometry, you might have to move the camera before you can see objects in its Camera viewport. Check the camera's position from several viewports to correct.

Procedure

To create a free camera:

1.  From the Create menu, choose Cameras > Free Camera, or click Cameras on the Create panel, then click Free on the Object Type rollout.
2. Click the viewport location where you want the camera to be.

   The kind of viewport you click determines the free camera's initial direction.

   The camera is now part of the scene.
3. Set the creation parameters.
4. Rotate and move the camera to adjust the point of view.

Interface

See *Common Camera Parameters (page 2–1373)*.

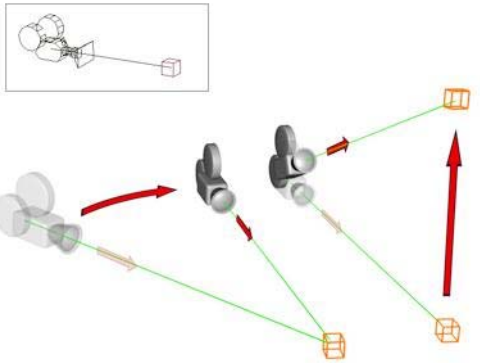
Target Camera

Create panel > Cameras > Target
Create menu > Cameras > Target Camera

A target camera "views" the area around the target icon that you place when you create the camera. A target camera is easier to aim than a free camera because you simply position the target object at the center of interest.

You can animate both the target camera and its target to create interesting effects. To animate both the target and camera along a path, it is best to link them both to a dummy object, then animate the dummy.

Note: When you add a target camera, 3ds Max automatically assigns a *Look At controller (page 2–344)* to it, with the camera's target object assigned as the Look At target. You can use the controller settings on the Motion panel to assign any other object in the scene as the Look At target.



Target cameras always face their target.

Procedure

To create a target camera:

1. Do one of the following:
   - Click Cameras on the Create panel, then click Target in the Object Type rollout.
   - Choose Create menu > Cameras > Target Camera.
2. Drag in a Top or Perspective viewport. The initial point of the drag is the location of the camera, and the point where you release the mouse is the location of the target.

   The camera is now part of the scene. It is aimed at the target, which is a separate object.
3. Set the creation parameters.

Interface

See *Common Camera Parameters (page 2–1373)* for a description of the common camera parameters.

The distance from the camera to the target is displayed at the bottom of the Parameters rollout. You can animate this parameter, or directly animate the target object's location.

When you rename a target camera, the target is automatically renamed to match. For example, renaming *Camera01* to *Rolli* causes *Camera01.Target* to become *Rolli.Target*. The target's name must have the extension *.Target*. Renaming the target object does not rename the camera object.

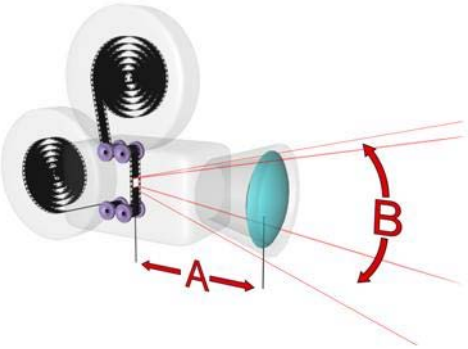
Clicking the line that connects the camera and its target selects both objects. However, region selection doesn't recognize the link line.

If a target camera is already selected, you add its target to the selection by right-clicking the camera, and then choosing Select Target from the quad menu > Tools1 quadrant. Or you can hold down the Ctrl key and click the target to add it to the selection set.

Using Cameras

Characteristics of Cameras

Real-world cameras use lenses to focus the light reflected by a scene onto a *focal plane* that has a light-sensitive surface.



Real-world camera measurements.
A: Focal length
B: Field of view (FOV)

Focal Length

The distance between the lens and the light-sensitive surface, whether film or video electronics, is called the *focal length* of the lens. Focal length affects how much of the subject appears in the picture. Lower focal lengths include more of the scene in the picture. Higher focal lengths include less of the scene but show greater detail of more distant objects.

Focal length is always measured in millimeters. A 50mm lens is a common standard for photography. A lens with a focal length less than 50mm is called a *short* or *wide-angle* lens. A lens with a focal length longer than 50mm is called a *long* or *telephoto* lens.

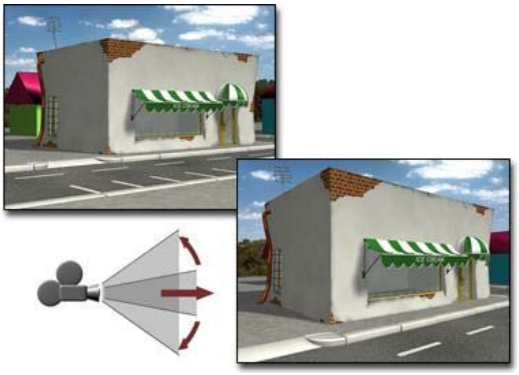
Field of View (FOV)

The field of view (FOV) controls how much of the scene is visible. The FOV is measured in degrees of the horizon. It is directly related to the focal length of the lens. For example, a 50mm lens shows 46 degrees of the horizon. The longer the lens, the narrower the FOV. The shorter the lens, the wider the FOV.

Relationship Between FOV and Perspective

Short focal lengths (wide FOV) emphasize the distortions of perspective, making objects seem in-depth, looming toward the viewer.

Long focal lengths (narrow FOV) reduce perspective distortion, making objects appear flattened and parallel to the viewer.



Upper left: Long focal length, narrow FOV

Lower right: Short focal length, wide FOV

The perspective associated with 50 mm lenses appears normal, partly because it is close to what the eye sees, and partly because such lenses are so widely used for snapshots, news photos, cinema, and so on.

Differences Between Camera Objects and Real-World Cameras

Many other controls on real-world cameras (such as those for focusing a lens, and advancing film) aren't needed for computer rendering and have no counterpart in the camera objects.

The program does have counterparts for the camera movements used in movie making, such as truck, dolly, and pan. See *Camera Viewport Controls (page 3–745)*.

Procedure

To match a real-world camera frame proportion:

1. Choose Rendering > Render.

   The Render Scene dialog is displayed.
2. In the Output Size group, click the arrow to display the list of real world output sizes.
3. Select the type you want (both film and video output sizes are available).
4. Right-click the Camera viewport label, and turn on Show Safe Frame.

   The Safe Frame proportions will match those of the selected output size.

Common Camera Parameters

Create panel > Cameras > Target button or Free button > Parameters rollout

Most of the camera controls are common to both kinds of cameras. This topic describes those controls.

Procedures

To view a wider area, do either of the following:

- Use the FOV spinner to increase the camera's field of view.
- Click a button with a shorter focal length. Use the Lens spinner to give the focal length a value other than the preset "stock" values on the buttons.

To view a narrower area, do either of the following:

- Change the FOV parameter to decrease the camera's field of view.
- Click a button with a longer focal length. Use the Lens parameter to give the focal length a value other than the preset "stock" values on the buttons.

In a camera viewport, the FOV button lets you adjust the field of view interactively.

The camera viewport Perspective button also changes the FOV in conjunction with dollying the camera.

Note: Only the FOV value is saved with the camera. The focal length value is merely an alternative way to express and select the FOV.

To set the camera lens size:

1. In the Stock Lenses group, click a button to choose a stock focal length.
2. Set the Lens spinner to a custom focal length.

   Tip: If you want to maintain the same lens, avoid using the FOV or Perspective controls among the navigation icon buttons, and don't change the FOV spinner.

   Important: When a camera viewport is active, changing the Output Size or (custom) Aperture Width in the *Render Scene dialog (page 3–2)* will change the camera's Lens setting.

To match a camera to a film or video format:

1. In the *Render Scene dialog (page 3–2)*, choose the type of output you want.
   - If it's a preset, choose the preset from the Output Size drop-down list. The Aperture Width will be locked to the preset's values.
   - If it's Custom, you can set your own Aperture Width, and should do so at this point. (You can adjust the other output values at any later time. They have no affect on the camera lens settings, although they do affect the cropping of the scene.)
2. Once the Aperture Width is set, set the Lens spinner to the type of camera lens you want to emulate (for example, 50mm).

   If you want to maintain the same lens, avoid using the FOV or Perspective controls among the navigation icon buttons.

To find a lens's focal length:

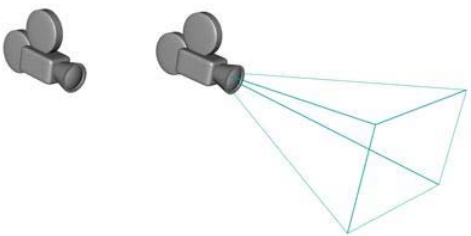
- To find the focal length of a lens based on changes in aperture width, display the *Render Scene dialog (page 3–2)*, choose Custom from the Output Size drop-down list, and specify a value in the Aperture Width spinner. The new value of the camera's Lens spinner is based on the new Aperture Width value.

To display a camera's cone:

- Turn on Show Cone.

  The camera's field-of-view cone appears outlined in light blue.

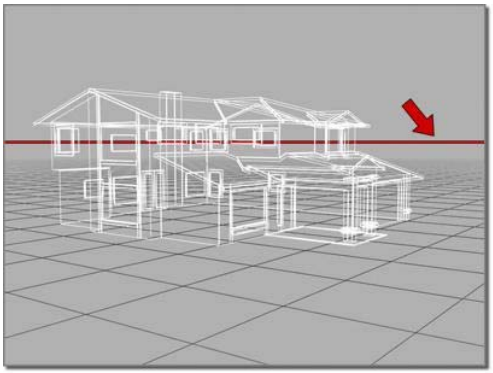
  Note: A camera's cone is always visible while the camera object is selected, regardless of the Show Cone setting.



To display a camera's horizon line:

- Turn on Show Horizon.

  A dark gray line appears at the level of the horizon in the camera's viewport.



The horizon line shown in the viewport.

The horizon line might not be visible if the horizon is beyond the camera's field of view, or if the camera is tilted very high or low.

To change the environment range:

- Adjust the value of Near Range or Far Range.

  By default, the Near Range=0.0 and the Far Range equals the Far clipping plane value.

  Environment ranges determine the near and far range limits for atmospheric effects you set in the Environment dialog.

To see the environment ranges in viewports:

- Turn on Show.

  The environment range displays as two planes. The plane closest to the camera is the near range and the one farthest from the camera is the far range.

To set clipping planes:

1. Turn on Clip Manually.

   When Clip Manually is off, the camera ignores the location of the Near and Far clipping planes, and their controls are unavailable. The camera renders all geometry within its field of view.

2. Set the Near Clip value to position the near clipping plane.

   Objects closer to the camera than the Near distance are not visible to the camera and aren't rendered.

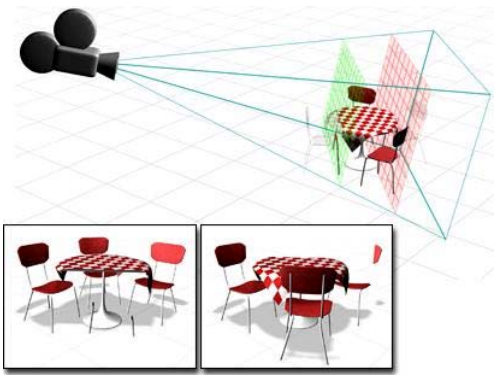
3. Set the Far Clip value to position the far clipping plane.

   Objects farther from the camera than the Far distance are not visible to the camera and aren't rendered.

   You can set the Near clipping plane close to the camera so that it doesn't exclude any geometry, and still use the Far plane to exclude objects. Similarly, you can set the Far clipping plane far enough from the camera that it doesn't exclude any geometry, and still use the Near plane to exclude objects.

   The Near value is constrained to be less than the Far value.

   If the clipping plane intersects an object, it cuts through that object, creating a cutaway view.



The effect of clipping planes

To apply a multi-pass rendering effect to a scene:

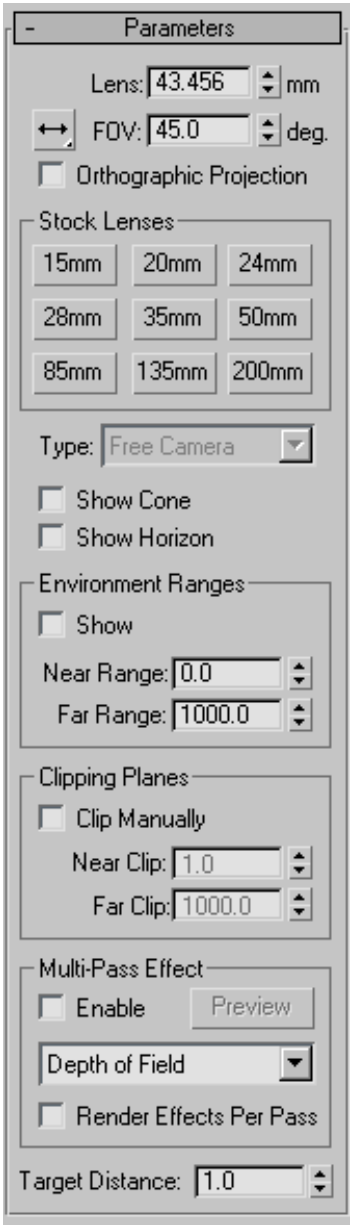
1. In the Multi-Pass Effect group, turn on Enable and choose either Depth Of Field or Motion Blur.

2. Use the *Depth Of Field Parameters rollout (page 2–1383)* or the *Motion Blur Parameters rollout (page 2–1386)* to set the values for the effect you chose.

3. Activate a camera viewport.
4. In the Multi-Pass Effect group, click Preview to preview the effect in the camera viewport.

   The Preview button has no effect if a camera viewport isn't active.
5. Render the scene or animation.

Interface



Lens—Sets the camera's focal length in millimeters. Use the Lens spinner to give the focal length a value other than the preset "stock" values on the buttons in the Stock Lenses group box.

When you change the Aperture Width value in the Render Scene dialog, you also change the value in the Lens spinner field. This doesn't change the view through the camera, but it does change the relationship between the Lens value and the FOV value, as well as the aspect ratio of the camera's cone.

FOV Direction flyout—Lets you choose how to apply the field of view (FOV) value:

-  **Horizontal**—(The default.) Applies the FOV horizontally. This is the standard way to set and measure the FOV.
-  **Vertical**—Applies the FOV vertically.
-  **Diagonal**—Applies the FOV diagonally, from one corner of the viewport to the other.

FOV—Determines how wide an area the camera views (*field of view (page 3–937)*). When FOV Direction is horizontal (the default), the FOV parameter directly sets the arc of the camera's *horizon (page 3–953)*, measured in degrees. You can also set the FOV Direction to measure FOV vertically or diagonally.

You can also adjust the field of view interactively in a *camera viewport (page 3–745)* by using the FOV button.

Orthographic Projection—When on, the camera view looks just like a User view. When off, the camera view is the standard perspective-like view. While Orthographic Projection is in effect, the viewport navigation buttons behave as they ordinarily do, except for Perspective. Perspective function still moves the camera and changes the FOV, but the Orthographic Projection cancels the two out, so you don't see any change until you turn off Orthographic Projection.

Stock Lenses group

15mm, 20mm, 24mm, 28mm, 35mm, 50mm, 85mm, 135mm, 200mm—These preset values set the camera's focal length in millimeters.

Type—Changes the camera's type from a *Target camera (page 2–1371)* to a *Free camera (page 2–1370)*, and vice versa.

Note: When you switch from a target camera to a free camera, any animation applied to the camera's target is lost, because the target object goes away.

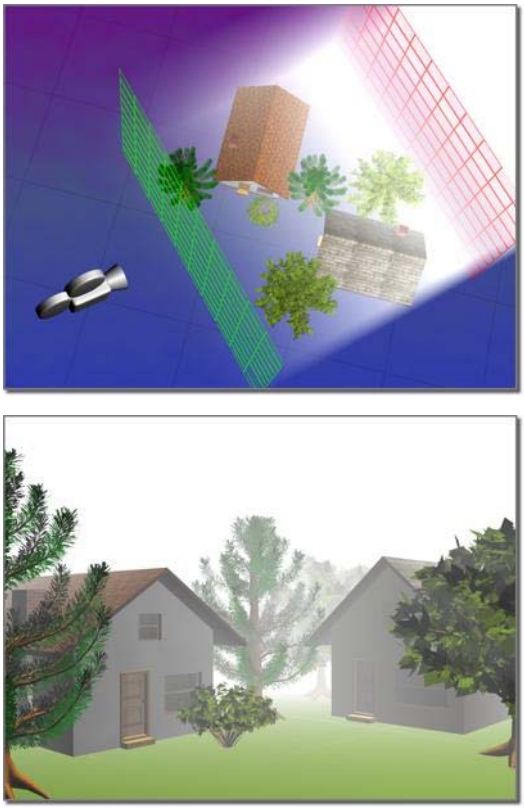
Show Cone—Displays the cone (actually a pyramid) defined by a camera's field of view. The cone appears in the other viewports but does not appear in a camera viewport.

Show Horizon—Displays the horizon line. A dark gray line appears at the level of the horizon in the camera's viewport.

Environment Ranges group

Near Range and Far Range—Determine the near and far range limits for the atmospheric effects set on the *Environment panel (page 3–272)*. Objects between the two limits fade between the Far % and Near % values.

Show—Displays rectangles within the camera's cone to show the Near and Far range settings.



Top: Conceptual image of the Near and Far ranges.
Bottom: Result after rendering.

Clipping Planes group

Sets options to define *clipping planes (page 3–921)*. In viewports, clipping planes are displayed as red rectangles (with diagonals) within the camera's cone.

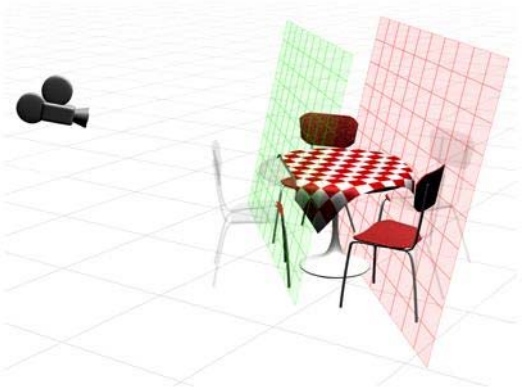
Clip Manually—Turn on to define clipping planes.

When Clip Manually is off, geometry closer to the camera than 3 units is not displayed. To override this, use Clip Manually.

Near Clip and Far Clip—Sets near and far planes. Objects closer than the near clipping plane or farther than the far clipping plane are invisible to the camera. The limit of the Far Clip value is 10 to the power of 32.

With manual clipping on, the near clipping plane can be as close to the camera as 0.1 unit.

Warning: Extremely large Far Clip values can produce floating-point error, which can cause Z-buffer problems in the viewport, such as objects appearing in front of other objects when they shouldn't.



Conceptual image of Near and Far clipping planes.

Multi-Pass Effect group

These controls let you assign a depth-of-field or motion blur effect to the camera. When generated by a camera, these effects generate blurring by rendering the scene in multiple passes, with offsets. They increase rendering time.

Tip: The depth-of-field and motion blur effects are mutually exclusive. Because they rely on multiple rendering passes, applying both to the same camera could be prohibitively slow. If you want to use both depth-of-field and motion blurring in the same scene, use multi-pass depth-of-field (using these camera parameters) and combine it with *object motion blur (page 3–981)*.

Enable—When on, previewing or rendering uses the effect. When off, the effect is not rendered.

Preview—Click to preview the effect in an active camera viewport. This button has no effect if the active viewport is not a camera view.

Effect drop-down list—Lets you choose which multi-pass effect to generate, *Depth Of Field (page 2–1383)* or *Motion Blur (page 2–1386)*. These effects are mutually exclusive. Default=Depth Of Field.

This list also lets you choose *Depth of Field (mental ray) (page 2–1383)*, which lets you use the mental ray renderer's depth of field effect.

Note: The rollout for the chosen effect appears, by default, after the Parameters rollout.

Render Effects Per Pass—When on, applies *rendering effects (page 3–218)*, if any are assigned, to each pass of the multi-pass effect (depth of field or motion blur). When off, applies rendering effects only after the passes that generate the multi-pass effect. Default=off.

Turning off Render Effects Per Pass can improve the render time for multi-pass effects.

Target Distance—With a free camera, sets a point to use as an invisible target so that you can orbit a camera around that point. With a target camera, indicates the distance between the camera and its target.

Using Transforms to Aim a Camera

You can use transforms to aim a camera and change its orientation in the scene.

 *Move (page 1–439)* adjusts the position of the camera object or the position of a *target camera's target (page 2–1371)*.

Because the target is displayed as a small square, and because it is often in the same area as objects that are the subject of the camera, it can be hard to select by clicking. Select the camera object, right-click, then choose Select Target from the Tools 1 (upper-left) quadrant of the quad menu.

You can also choose Cameras from the Selection Filters list on the toolbar, and then click the target.

 *Rotate (page 1–439)* adjusts the orientation of the camera object. This transform is most useful with *free cameras (page 2–1370)*.

You can't rotate a Target camera about its local X and Y axes, because it is constrained to aim at its target. Use Move to move the camera or its target.

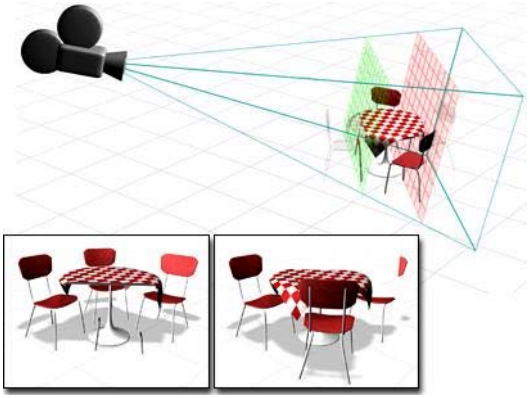
Also, if you rotate a Target camera to a nearly vertical position, either up or down, the program must flip the Camera view to prevent the up-vector from becoming undefined. If you need a camera to look vertically upward or downward, use a Free camera.

In a Camera viewport, you can also use the *navigation buttons (page 3–745)* to adjust the camera interactively. Some navigation buttons, such as Dolly and Orbit actually move the camera or its target.

Using Clipping Planes to Exclude Geometry

Clipping planes let you exclude some of a scene's geometry and view or render only certain portions of the scene. Each camera object has a near and a far clipping plane. Objects closer than the near clipping plane or farther than the far clipping plane are invisible to the camera.

Clipping planes are useful for rendering selected portions of a scene that has a lot of complex geometry. They can also help you create cutaway views.



Left: Clipping plane excludes the foreground chair and the front of the table.

Right: Clipping plane excludes the background chair and the rear of the table.

Clipping plane settings are part of the camera's *parameters (page 2–1373)*. The location of each clipping plane is measured along the camera's line of sight (its local Z axis) in the current units for the scene.

You can set the near clipping plane close to the camera so that it doesn't exclude any geometry, and still use the far plane to exclude objects. Similarly, you can set the far clipping plane far enough from the camera that it doesn't exclude any geometry, and still use the near plane to exclude objects.

The near value is constrained to be less than the far value.

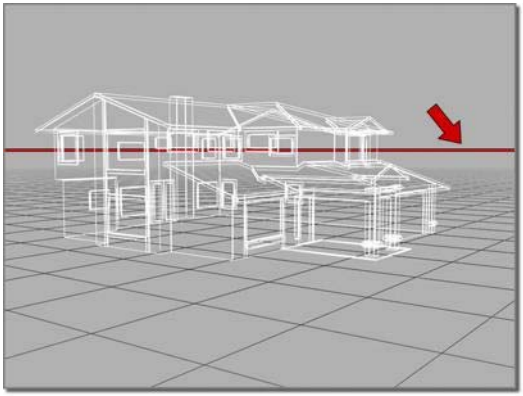
If the clipping plane intersects an object, it cuts through that object, creating a cutaway view. (How much of the cutaway object is visible depends on whether the object's material is two-sided or not.)

You can also use clipping planes in non-camera viewports. Simply right-click the viewport label and choose *Viewport Clipping (page 3–731)*.

Warning: If you use the *mental ray renderer (page 3–78)*, geometry outside the clipping planes might still appear in renderings.

Using the Horizon to Match Perspective

The horizon of a scene is the edge of vision at the height of the camera, parallel with the world coordinate plane. You can view the horizon in camera viewports.



Horizon line displayed in a camera viewport

A camera is level when it and its target are the same height from the world coordinate plane. In other words, the camera's local Z axis is parallel to the world plane. When the camera is level, the horizon line is centered in the viewport. As the camera tilts up, the horizon line lowers; as it tilts down, the horizon line raises.

The horizon line control is in the camera's *Parameters rollout (page 2–1373)*.

The horizon line can help you match the perspective of your scene to the perspective of a still image. In general, matching perspective involves the following steps:

1. Display the horizon line. Use it to help you adjust the camera and target so they are level.
2. Display the image in the camera viewport. Use Views > Viewport Background and choose Files in the Background Source group.

If the image's horizon and the camera horizon don't match, you have to offset the image, perhaps by using a paint program.

3.  Use *Orbit (page 3–749)* to move the camera until the perspective of the scene roughly matches that of the still image.

4.  Adjust the camera's *perspective (page 3–747)* to fine-tune the perspective match.

5.  Use *Move (page 1–439)* with the camera or target to position the scene against the background.

   If you raise or lower the camera, raise or lower the target by an equal amount, in order to keep them level and maintain the horizon.

   You can also use the *Camera Match utility (page 2–1387)* to match the perspective of a camera to a photograph. You will need accurate scene measurements to do so.

Animating Cameras

You animate a camera by using transforms or changing its creation parameters in different keyframes while the Set Key or Auto Key button is on. The program interpolates camera transforms and parameter values between keyframes, as it does for object geometry.

See *Auto Key Animation Mode (page 3–717)* and *Track View (page 2–501)* for further descriptions of animation. This topic summarizes some possibilities and suggests some techniques.

In general, it's best to use a *free camera (page 2–1370)* when the camera is to move within the scene; use a *target camera (page 2–1371)* when camera position is fixed.

Moving a Camera Along a Path

Having a camera follow a path is a common way to create architectural walkthroughs, roller coaster rides, and so on.

- If the camera must bank or tilt close to the vertical (as on a roller coaster), use a free camera. Assign the *Path constraint (page 2–398)* directly to the camera object. The camera follows the path, and you can adjust its point of view by adding pans or rotate transforms. This is comparable to filming with a hand-held camera.
- For a target camera, link both the camera and its target to a *dummy object (page 2–16)*, then assign the path constraint to the dummy object. This is comparable to mounting the camera on a tripod on a dolly. It is easier to manage than having, for example, separate paths for the camera and its target.

Following a Moving Object

You can use a *LookAt constraint (page 2–406)* to have the camera automatically follow a moving object.

- The LookAt constraint makes the object replace the camera's target.

  If the camera is a target camera, its previous target is ignored.

  If the camera is a free camera, it effectively becomes a target camera. While the LookAt constraint assignment is in effect, the free camera cannot rotate around its local X and Y axes, and can't be aimed vertically because of the up-vector constraint.
- An alternative is to *link (page 2–422)* a target camera's target to the object.

Panning

You can animate the pan of any camera very easily by following these steps:

1. Select the camera.
2. Activate the Camera viewport.
3. Turn on the Auto Key button and advance the time slider to any frame.
4.  Use the Pan button (in the viewport navigation tools) and pan.

Orbiting

You can animate the orbiting of any camera very easily by following these steps:

1. Select the camera.
2. Activate the Camera viewport.
3. Turn on the Auto Key button and advance the time slider to any frame.
4.  Use the Orbit button (in the viewport navigation tools) and orbit.

   The target camera revolves around its target; the Free camera revolves around its target distance.

Zooming

Zooming moves toward or away from the camera's subject matter by changing the focal length of the lens. It differs from dollying, which physically moves the camera but leaves the focal length unchanged. You can zoom by animating the value of the camera's *FOV parameter (page 2–1373)*.

Creating Animated Cutaway Views

You can animate the creation of a cutaway view by animating the location of the near or far *clipping planes (page 2–1379)*, or both.

Multi-Pass Rendering Effects

Create panel > Cameras > Target or Free > Parameters rollout > Multi-Pass Effect group



Motion blur applied to wings of the flying dragon

Cameras can create two kinds of rendering effects:

- *Depth of field (page 2–1383)*
- *Motion blur (page 2–1386)*

Multi-pass rendering effects use multiple renderings of the same frame, with slight camera movement between each rendering. The multiple passes simulate the blurring that film in a camera would register under certain conditions.

Although it is not a multi-pass effect, the choices in the drop-down list also let you specify the depth-of-field value for the *mental ray renderer (page 3–78)*. See *Depth of Field Parameter (mental ray Renderer) (page 2–1383)*.

Depth of Field Parameter (mental ray Renderer)

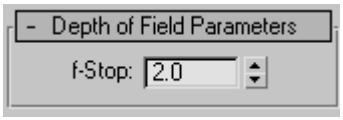
Create panel > Cameras > Target button or Free button > Parameters rollout > Multi-Pass Effect group. > Turn on Enable and choose Depth Of Field (mental ray). > Depth of Field Parameters rollout

On the *Parameters rollout (page 2–1373)*, a "Depth Of Field (mental ray)" choice has been added to the Multi-Pass Effect drop-down list to support the mental ray renderer's depth-of-field effects. To use this, turn on both Enable in the camera's Multi-Pass Effect group, and Depth Of Field on the Camera Effects rollout of the *Render Scene dialog (page 3–2)*.

The mental ray depth-of-field is exclusive of the multi-pass version of the depth-of-field effect. The mental ray renderer also supports motion blur for cameras, but the controls are not on the camera's Parameters rollout: use the Motion Blur toggle on the Object Properties dialog for camera objects. This setting has no effect on the default 3ds Max scanline renderer.

Note: When you use the mental ray renderer, reflected or refracted light rays do not always respect a camera's clipping planes (set in the Clipping Planes group of the Parameters rollout). Also, large clipping-plane values can cause poor quality in the rendering of shadow maps. To fix this, narrow the clipping range or use ray-traced shadows.

Interface

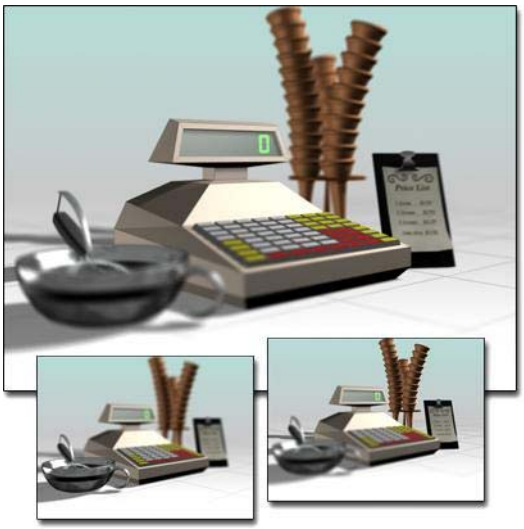


f-Stop—Sets the camera's f-Stop. Increasing the f-Stop value narrows the depth of field, and decreasing the f-Stop value broadens the depth of field. Default=2.0.

The f-Stop can have a value less than 1.0. This is not realistic in terms of an actual camera, but it can help you adjust the depth of field for scenes whose scale does not use realistic units.

Multi-Pass Depth of Field Parameters for Cameras

Create panel > Cameras > Target button or Free button > Parameters rollout > Multi-Pass Effect group > Choose Depth Of Field effect. > Depth of Field Parameters rollout



Multi-pass depth of field

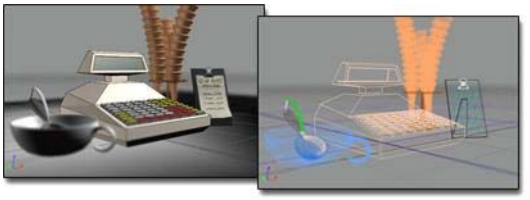
Top: Focus is in the middle distance, near and far objects are blurred.

Bottom left: Focus on near objects, far objects are blurred.

Bottom right: Focus on far objects, near objects are blurred.

Cameras can generate depth-of-field effects. Depth of field is a *multi-pass effect (page 3–77)*. You turn it on in the *Parameters rollout (page 2–1373)* for cameras. Depth of field simulates a camera's depth of field by blurring areas of the frame at a distance from the camera's focal point (that is, its target or target distance).

You can preview depth of field in viewports.



Previewing multi-pass depth of field in a shaded and a wireframe viewport

Important: This effect is for the default scanline renderer. The *mental ray renderer (page 3–78)* has its own depth-of-field effect. See *Depth of Field Parameter (mental ray Renderer) (page 2–1383)*.

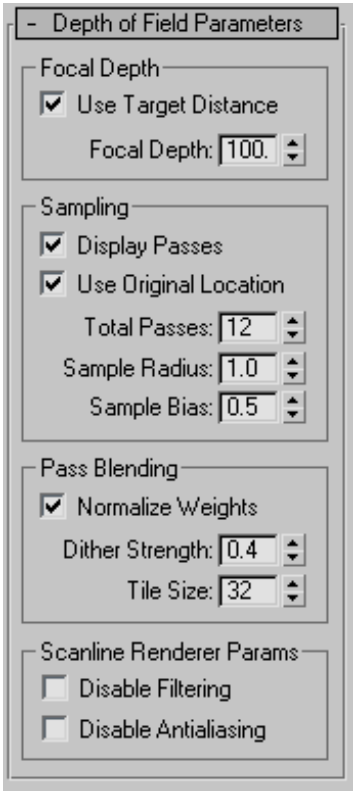
Tip: To reduce the visible effect of multiple camera passes, try setting the antialiasing filter to Blend, with a Width value in the range 4.0 to 5.0, and a Blend value in the neighborhood of 0.1. (You choose the antialiasing filter and adjust its settings in the *Default Scanline Renderer rollout (page 3–38)*.) Also, try reducing the Dither Strength value, in the effect's Pass Blending group, to the neighborhood of 0.2.

Note: There is also a *depth-of-field rendering effect (page 3–269)*.

See also

Multi-Pass Motion Blur Parameters for Cameras (page 2–1386)

Interface



Note: The multi-pass depth-of-field parameters are animatable.

Focal Depth group

Use Target Distance—When on, uses the camera's target distance as the point about which to offset the camera for each pass. When off, uses the Focal Depth value to offset the camera. Default=on.

Focal Depth—When Use Target Distance is off, sets the depth from which the camera is offset. Can range from 0.0 to 100.0, where 0.0 is at the camera's location and 100.0 is in the extreme distance (effectively, infinity). Default=100.0.

Low values of the Focal Depth give wildly blurry results. High Focal Depth values blur the distant portions of the scene. In general, using Focal

Depth instead of the camera's Target Distance tends to blur the entire scene.

Sampling group

Display Passes—When on, the rendered frame window displays the multiple rendering passes. When off, the frame window displays only the final result. This control has no effect on previewing depth of field in camera viewports. Default=on.

Use Original Location—When on, the first rendering pass is in the camera's original location. When off, the first rendering pass is offset like all subsequent passes. Default=on.

Total Passes—The number of passes used to generate the effect. Increasing this value can increase the effect's accuracy, but at a cost of rendering time. Default=12.

Sample Radius—The radius by which the scene is shifted to generate blurriness. Increasing this value increases the overall blurriness of the effect. Decreasing it reduces the blurriness. Default=1.0.

Sample Bias—Weights the blurring toward or away from the Sample Radius. Increasing this value increases the magnitude of depth-of-field blurring, giving a more even effect. Decreasing it decreases the magnitude, giving a more random effect. Can range from 0.0 to 1.0. Default=0.5.

Pass Blending group

The multiple depth-of-field passes are blended by dithering, which you can control by the parameters in this group.

These controls apply only to renderings of the depth-of-field effect, not to previews in viewports.

Normalize Weights—Passes are blended with random weighting to avoid artifacts such as streaking. When Normalize Weights is on, the weights are normalized, giving a smoother result. When off, the effect is a bit sharper but usually grainier. Default=on.

Dither Strength—Controls how much dithering is applied to the rendered passes. Increasing this value increases the amount of dithering, and can make the effect grainier, especially at the edges of objects. Default=0.4.

Tile Size—Sets the size of the pattern used in dithering. This value is a percentage, where 0 is the smallest tile, and 100 is the largest. Default=32.

Scanline Renderer Params group

These controls let you disable antialiasing or antialias filtering when you render the multi-pass scene. Disabling these rendering passes can improve render time.

These controls apply only to renderings of the depth-of-field effect, not to previews in viewports.

Disable Filtering—When on, disables the filtering pass. Default=off.

Disable Antialiasing—When on, disables antialiasing. Default=off.

Multi-Pass Motion Blur Parameters for Cameras

Create panel > Cameras > Target button or Free button > Parameters rollout > Multi-Pass Effect group > Choose Motion Blur effect. > Motion Blur Parameters rollout

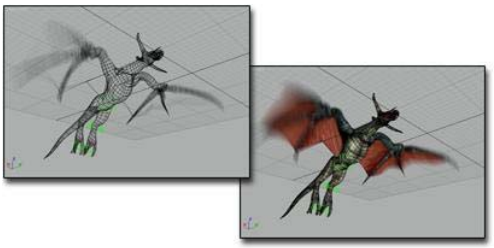


Above: Motion blur applied to wings of the flying dragon

Below: Multiple passes appear in successive refreshes of the rendered frame window.

Cameras can generate motion blur effects. Motion blur is a *multi-pass effect (page 3–77)*. You turn it on in the *Parameters rollout (page 2–1373)* for cameras. Motion blur simulates the motion blur of a camera by offsetting rendering passes based on movement in the scene.

You can preview motion blur in viewports.

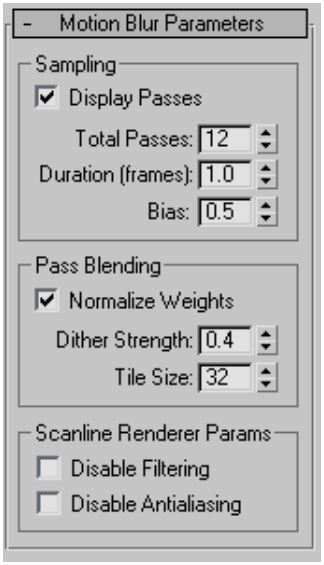


Previewing multi-pass motion blur in a wireframe and a shaded viewport

Important: This effect is for the default scanline renderer. The *mental ray renderer (page 3–78)* has its own depth-of-field effect. See *Motion Blur with the mental ray Renderer (page 3–89)*.

Tip: To reduce the visible effect of multiple camera passes, try setting the antialiasing filter to Blend, with a Width value in the range 4.0 to 5.0, and a Blend value in the neighborhood of 0.1. (You choose the antialiasing filter and adjust its settings on the *Default Scanline Renderer rollout (page 3–38)*.) Also, try reducing the Dither Strength value, in the effect's Pass Blending group, to the neighborhood of 0.2.

Interface



Note: The multi-pass motion blur parameters are animatable.

Sampling group

Display Passes—When on, the rendered frame window displays the multiple rendering passes. When off, the frame window displays only the final result. This control has no effect on previewing motion blur in camera viewports. Default=on.

Total Passes—The number of passes used to generate the effect. Increasing this value can increase the effect's accuracy, but at a cost of rendering time. Default=12.

Duration (frames)—The number of frames in the animation to which the motion blur effect will be applied. Default=1.0.

Bias—Changes the blurring so that it appears to derive more from frames before or after the current frame. Range=0.01 to 0.99. Default=0.5.

By default, the blurring comes equally from frames before and after the current frame; that is, a blurred object appears at the center of the blurred area. This is the closest to what an actual camera would capture. Increasing the Bias value moves the blurring behind the blurred object, in relation to its direction of motion. Decreasing it moves the blurring in front of the blurred object.

Extreme values move most of the blurring very close to the blurred object, which makes it difficult to see. For best results, use intermediate Bias values in the range 0.25 to 0.75.

Pass Blending group

The multiple motion blur passes are blended by dithering, which you can control by the parameters in this group.

These controls apply only to renderings of the motion blur effect, not to previews in viewports.

Normalize Weights—Passes are blended with random weighting to avoid artifacts such as streaking. When Normalize Weights is on, the weights are normalized, giving a smoother result. When off, the effect is a bit sharper but usually grainier. Default=on.

Dither Strength—Controls how much dithering is applied to the rendered passes. Increasing this value increases the amount of dithering, and can make the effect grainier, especially at the edges of objects. Default=0.4.

Tile Size—Sets the size of the pattern used in dithering. This value is a percentage, where 0 is the smallest tile, and 100 is the largest. Default=32.

Scanline Renderer Params group

These controls let you disable antialiasing or antialias filtering when you render the multi-pass scene. Disabling these rendering passes can improve render time.

These controls apply only to renderings of the motion blur effect, not to previews in viewports.

Disable Filtering—When on, disables the filtering pass. Default=off.

Disable Antialiasing—When on, disables antialiasing. Default=off.

Camera Match Utility

Tools menu > Camera Match

Utilities panel > Utilities rollout > Camera Match button

The Camera Match utility uses a bitmap background photo and five or more special "CamPoint" objects to create or modify a camera so that its position, orientation, and field-of-view matches that of the camera that originally created the photo.

Procedures

The general process to follow in using Camera Match is described in the following procedures.

To use camera matching:

1. Load a bitmap as a background for the renderer.

2. Load a bitmap as a background for the viewport.
3. Identify on the bitmap at least five features that will be used for the match. These should be objects or corners of objects in the scene that can be identified and tracked. They should remain visually throughout the scene, and should not change their shape too much or they won't work.
4. Create CameraPoints, which are helper objects found in these locations: Create panel > Helpers > Camera Match > Object Type rollout > CamPoint, and Create menu > Helpers > Camera Point.

   You must have accurate measurements of distances between at least five features in your scene, which can't be all on a single plane. Try to use points that are distributed throughout the scene rather than features that are all clustered in the front or back. This will give the most accurate results.
5. Position these CameraPoints to correspond to points in your photo. You can use the *Transform Type-In (page 1–431)* to position the points in the correct locations in 3D space.
6. Use the Camera Match utility to assign the CameraPoints to pixel locations on the bitmap.
7. Choose Create Camera and a camera is created to match the one that took the picture.

   If there are errors and the camera cannot be created, readjust the point positions in 3D space, and reassign them to the bitmap. It's easy to make a mistake doing either, but once you get it right the match should work.

To establish the bitmap background for the renderer:

1. Reset the software and enlarge the Perspective viewport to full screen.
2. From the Rendering menu choose Environment.
3. In the Common Parameters rollout, under Environment Map, click the gray box marked None. This launches the Material/Map Browser.
4. Be sure to set Browse From to New; if it isn't set already, then choose Bitmap from the list and click OK. The Select Bitmap Image File dialog appears.
5. Navigate and select the appropriate bitmap and then choose Open.

   Use Map is turned on automatically.
6. Render the viewport to verify that the background appears in the rendering. Press Shift+Q to render.

To establish the bitmap background to display in the viewport:

1. Choose Views menu > *Viewport Background (page 1–38)*.

   This opens the Viewport Background dialog
2. In the Background Source group, click Files. This launches the Select Background Image dialog.
3. Navigate and select the appropriate bitmap and then choose Open.
4. Set Aspect Ratio to Match Bitmap.
5. Make sure Display Background is turned on, and then choose OK.

   The background appears in the viewport.

   Note: You can click the Use Environment Background button, but then you don't have the Aspect Ratio controls.

To create CamPoint objects:

1. In the Helpers group of the Create panel, open the list and select Camera Match.
2. On the Object Type rollout, click the CamPoint button.
3. Open the Keyboard Entry rollout.

4. Enter the coordinates of the first CamPoint object (0,0,0), click the Create button, and then enter the name in the name field.

   Tip: To use the keyboard, first click in the X field, enter its value, and then press Tab to move to the next field and enter its value. Continue this until you tab to the Create button, and then press Enter to create the CamPoint, followed by Tab to move back to the X field, where you can start again. You can create all the CamPoint objects with default names, and then use the Select By Name floater (Tools > Selection Floater) to select and rename the six pointers.

5. Repeat the above steps for the remaining CamPoint objects.

   Note: An alternative method is to create the CamPoint objects anywhere in the scene, and then reselect each of them and enter their absolute coordinates using the Transform Type-In.

You now have the CamPoint objects occupying real-world coordinate positions that correspond to the structure in the bitmap image. The last sequence of steps involves using the Camera Match utility to specify the screen coordinate points, one for each CamPoint object, and generating a camera position based on the data.

To use the Camera Match utility:

1. On the Utilities panel, click the Camera Match button.

   The Camera Match utility appears, listing the CamPoint objects.

2. Select the first CamPoint object and click the Assign Position button.

3. Place the cursor over the corresponding feature in the bitmap and click.

   A small, red cross appears.

4. If the dot is not in the right position, you can either click again with the mouse or adjust the Input Screen Coordinates to tweak its position.

5. Select the second object in the list, and repeat steps 3 and 4.

6. Repeat for all the points. By setting the red crosses, you've indicated the X and Y pixel position on the bitmap and correlated it with each CamPoint as it exists in 3D space.

7. Once you have all of the points set, click the Create Camera button.

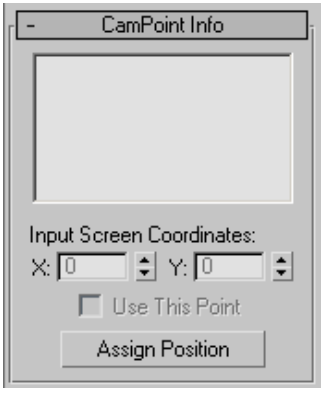
   A camera is created in the scene based on the location of the CamPoint objects and the specifications of the screen coordinates points.

   Note: If the Current Camera Error reading is greater than five, at least one of your screen coordinate points is placed wrong. Check each of them, and review the description following step 5. After reassigning the points, select the existing camera and click the Modify Camera button to recalculate the camera position.

8. Press C to switch the Perspective view to that of the new camera.

Interface

CamPoint Info rollout



List window—Displays a list of the CamPoint helper objects in the scene. You select the CamPoint

objects from this list to assign screen coordinate points. Note that if you select a CamPoint object in the viewport, it's highlighted in this list as well.

Input Screen Coordinates—

X/Y—Fine-tunes the position of the screen coordinate points in 2D space.

Use This Point—Turns off a specific coordinate point without deleting it. Select the corresponding CamPoint in the list, and then turn off Use This Point. This feature is typically used for troubleshooting when the Current Camera Error is too high (greater than five, for example).

Assign Position—Click a location on the viewport bitmap to place a screen coordinate point visually against the background image. The point you place corresponds to the currently selected CamPoint object. After activating the Assign Position button, select a CamPoint object from the list, and then click in the viewport at a position on the bitmap background that corresponds with where the associate CamPoint object should be in the 3D scene. After repeating this process with each CamPoint object in the list, you can click the Create Camera button to create a camera that matches the placed coordinates with their associate CamPoint objects.

Camera Match rollout



Create Camera—Creates a camera in the scene whose position, orientation, and field of view is based on the current location of the CamPoint helpers and the assigned screen coordinate points.

Modify Camera—Modifies the position, orientation and FOV of an existing, selected camera based on the CamPoint helpers and assigned screen coordinate points.

Iterations—Maximum number of iterations used to calculate the camera position. Default is 500, though a stable solution is usually found in less than 100 iterations.

Freeze FOV—Prevents the FOV (field of view) of the camera from being changed when using the Create Camera or Modify Camera buttons. Use if the FOV of the camera that took the photograph is known and you want to preserve it.

Current Camera Error—Displays the total error that remains between the placed screen coordinate points, the CamPoint helpers, and the camera position after the final computation. The calculations involved in the camera match are seldom perfect. A good error range is about 0 to 1.5.

Close—Exits the Camera Match utility.

Camera Match Helper Objects

Camera Match Helper

Create panel > Helpers > Camera Match (from drop-down list)

Create menu > Helpers > Camera Point

Camera Match helper objects (CamPoints) are used by the Camera Match utility to reproduce in a camera the same settings (position, roll, and FOV) that were used by a real-world camera to shoot a background image. This allows you to view your scene from the same perspective as the background image, which is a key step to mixing computer-generated and photographically-generated images in a single composited shot.

Camera point objects are placed in the scene at locations that will be visible in the background photograph. By comparing the X/Y/Z position of several camera point objects to their analogous positions in a background image, the Camera Match utility is able to determine the real-world camera's settings, and use these settings to either create a new camera or move/reset an existing camera to match them.

To see the *CamPoint (page 2–1391)* object type, choose CameraMatch from the Object set list.

See also

CamPoint Helper (page 2–1391)

Camera Match Utility (page 2–1387)

CamPoint Helper

Create panel > Helpers > Camera Match > CamPoint

Create menu > Helpers > Camera Point

The CamPoint helper lets you create camera points in your scene. These points are used by the *Camera Match (page 2–1387)* utility to match the camera with a photographic background. The virtual camera is positioned to mimic the position of the actual camera that was used to take the photograph.

Clicking CamPoint displays the Keyboard Entry and Camera Match Point rollouts, which you use to create and name the camera points.

Procedures

To create a camera match point:

1.  On the Create panel, click the Helpers tab.
2. Choose Camera Match from the list and click the CamPoint button in the Object Type rollout.

   The button turns yellow when it is active. You are now ready to create a camera point.
3. Choose Create menu > Helpers > Camera Point.
4. Click in the viewport to create a CamPoint helper object. Alternatively, you can use the keyboard to create a CamPoint object by expanding the Keyboard Entry rollout, entering values in the fields, and then clicking Create.

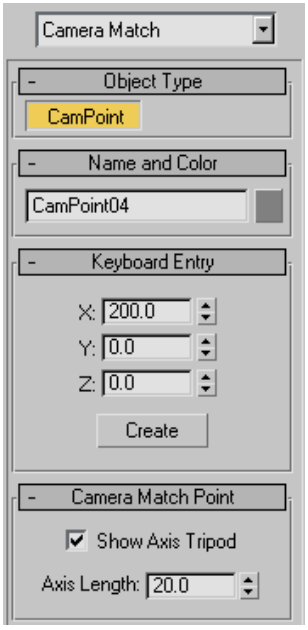
To position camera match points in your scene, do one of the following:

- Create your camera match points using keyboard entry. Use this if you have accurate actual measurements of the locations of the points.
- Build geometry that matches your scene geometry, and then snap the points to the

geometry using snaps. Use this technique when you don't have measurements and can approximate the geometry.

- Create your camera match points interactively, and then use Transform Type-In to move them to the correct locations.

Interface



Name and Color rollout

Lets you name the CamPoint object before you create it. If you want to name it after you create it, change its name in the Modify panel.

Keyboard Entry rollout

X, Y, Z—Let you enter the X, Y, and Z coordinates where a CamPoint object will be inserted. Field values reflect the current Unit setup.

Create—Inserts the CamPoint object into the scene.

Camera Match Point rollout

Show Axis Tripod—Controls whether an axis tripod is displayed with the Camera point object. Default=on.

Axis Length—Controls the length of the axis tripod. When Show Axis Tripod is on you can watch the tripod change length in the viewport while you adjust the spinner arrows. When Show Axis Tripod is off, you can still make adjustments the Axis Length value, but the axis tripod won't be displayed.

Two-Point Perspective

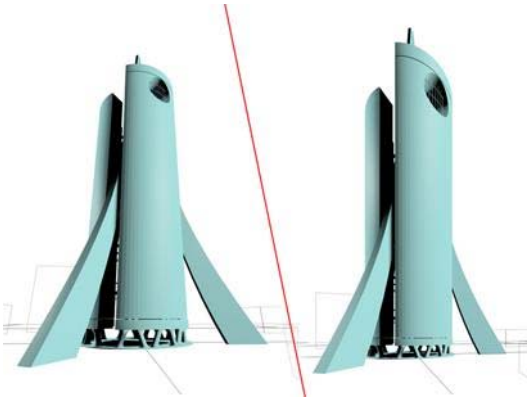
Camera Correction Modifier

Select a camera. > Right-click. > Tools 1 (upper-left) quadrant of the quad menu > Apply Camera Correction Modifier

The Camera Correction modifier applies two-point perspective to a camera view. By default, camera views use three-point perspective, in which vertical lines appear to converge with height. In two-point perspective, vertical lines remain vertical.

The amount of correction you need to use depends on how steeply the camera tilts. For example, a camera that looks up from ground level to the top of a high building will need more correction than a camera that looks toward the horizon.

Note: This modifier doesn't appear on the Modify panel's Modifier List.



Normal camera view (left) and corrected camera view (right)

Procedure

To apply two-point perspective to a camera:

1. Select a camera.

   Tip: For best results, set a viewport to this camera's view. The change in perspective appears both in the viewport and when you render this view.

2. Apply the Camera Correction modifier.
3. On the 2-Point Perspective Correction rollout, click Guess.

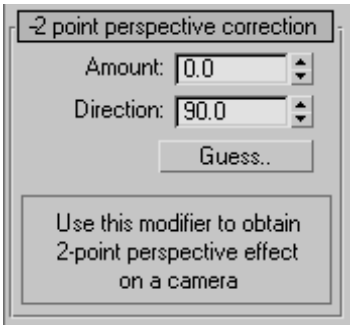
   The Camera Correction modifier creates a first-guess Amount value for the two-point perspective.

4. Adjust the Amount and Direction to get the effect you want.

   In the viewports, the camera's field-of-view "cone" distorts or moves to show the perspective adjustments.

5. Render the view.

Interface



Amount—Sets the amount of correction for two-point perspective. Default=0.0.

Direction—Biases the correction. Default=90.0.

Setting Direction greater than 90.0 biases the correction to the left. Setting it less than 90.0 biases it to the right.

Guess— Click to have the Camera Correction modifier set a first-guess Amount value.

16 Material Editor, Materials, and Maps

Materials describe how an object reflects or transmits light. Within a material, maps can simulate textures, applied designs, reflections, refractions, and other effects. (Maps can also be used as environments and projections from lights.) The Material Editor is the dialog you use to create, alter, and apply the materials in your scene.



Image by Michael McCarthy

See also

Designing Materials (page 2–1395)

Types of Materials (page 2–1457)

Material Editor (page 2–1409)

Types of Maps (page 2–1617)

Designing Materials



Materials make objects more convincing.

These topics give you an overview of using the Material Editor to design *materials (page 3–971)*. The *Material Editor (page 2–1427)* gives you a huge number of options in material design. It also has a large number of controls. If you aren't already familiar with using it, read this topic for a general idea of working with materials, and what the most important options are.

For more details, follow the links in the workflow outline.

Workflow Outline

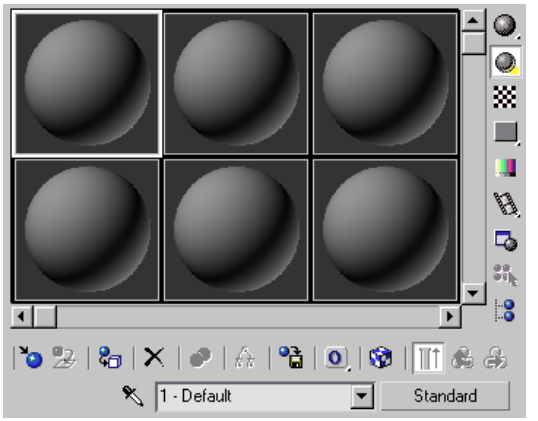
In general, when you create a new material and apply it to an object, you follow these steps:

1. Make a *sample slot (page 2–1396)* active, and enter a name for the material you are about to design.
2. Choose the *material type (page 2–1397)*.

   Tip: 3ds Max provides two renderers, the *default scanline renderer (page 3–38)* and the *mental ray renderer (page 3–78)*. Each has its own capabilities. In general, for each scene, you will decide which renderer you want to use. It is a good idea to design materials with a particular renderer in mind. The *mental ray Connection rollout (page 2–1461)* lets you add features unique to the mental ray renderer to basic 3ds Max materials.
3. For a Standard or Raytrace material, choose the *shading type (page 2–1398)*.
4. Enter settings for the various material *components (page 2–1399)*: diffuse color, glossiness, opacity, and so on.

   Note: *Lights and Shading (page 2–1399)* describes how lights affect the appearance of a material. *Choosing Colors for Realism (page 2–1400)* gives guidelines on getting good results from unmapped materials.
5. Assign *maps (page 2–1403)* to the components you want to map, and adjust their parameters.
6. *Apply (page 2–1405)* the material to the object.
7. If necessary, adjust the *UV mapping coordinates (page 2–1405)* in order to orient maps with objects correctly.
8. *Save (page 2–1406)* the material.

Sample Slots and Material Name



The *sample slots (page 2–1420)* display previews of materials. They are the most prominent feature of the Material Editor interface. Below and to the right of the sample slots are various *tool buttons (page 2–1427)* for the Material Editor. Below the tool buttons is a name field that shows the name of the material.

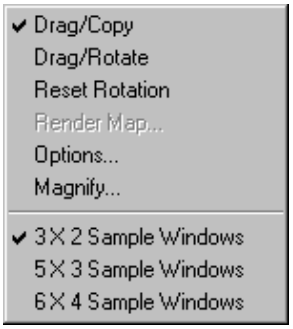
Tip: Always give a material a unique, intelligible name as soon as you begin to work on it.

By default, six sample slots are visible at once. The Material Editor actually holds 24 materials at one time. You can use the scroll bars to move among the sample slots, or you can change the number of sample slots visible at once to 15 or 24 slots. Seeing more slots at once can be helpful if you are working with a complex scene.

Important: While the Material Editor can edit no more than 24 materials at a time, a scene can contain an unlimited number of materials. When you are through editing one material, and have applied it to objects in the scene, you can use that sample slot to get a different material from the scene (or create a new one) and then edit it.

To increase the number of sample slots visible at once, right-click a slot and then choose 5 X 3

Sample Windows or 6 X 4 Sample Windows from the pop-up menu.



Note: The right-click menu also has an *Options choice (page 2–1436)*. This displays a dialog with various options for sample display. Exploring these options can help you learn to preview materials effectively. However, keep in mind that these settings affect the sample display *only*. They change nothing in the 3ds Max scene.

When more sample slots are visible, the images are smaller, but you can display a larger, floating, and resizable material sample by double-clicking the slot you want to see better.

Click a sample slot to make it active. Now you can design a new material from scratch, or you can load a previously stored material by clicking *Get Material (page 2–1439)*, which displays the Material/Map Browser. The Browser is a dialog that lets you choose materials and maps from a material library, from the scene, and so on.

You can also copy a material from one sample slot to another. Drag the slot with the material to another slot. To avoid confusion, rename the copy in the new sample slot before you begin to make changes to it.

If the *\matlibs* subdirectory contains a material library called *medit.mat,* the sample slots show the first 24 materials in this library file. If the library contains fewer than 24 materials, the remaining slots contain Standard materials of various colors, as they do if no *medit.mat* library is found.

Material Type

Every material has a type. The default is *Standard (page 2–1465)*, which is the material type you will probably use most often. In general, other material types are for special purposes. The other material types are:

- *Advanced Lighting Override (page 2–1601)*

  Used to fine-tune the effects of a material on *Advanced Lighting (page 3–44)*, including *light tracing (page 3–44)* and *radiosity solutions (page 3–51)*. Radiosity Override is not required for calculating advanced lighting, but it can enhance the result.

- *Blend (page 2–1588)*

  Mixes two other materials together. Can use a mask or a simple amount control.

- *Composite (page 2–1589)*

  Mixes up to 10 materials.

- *Double-Sided (page 2–1591)*

  Contains two materials, one for the front and one for the back faces of an object.

- *Ink 'n Paint (page 2–1605)*

  Creates cartoon effects with flat shading and "inked" borders.

- *Lightscape (page 2–1604)*

  Supports *import (page 3–573)* and *export (page 3–572)* of data from the Lightscape product.

- *Matte/Shadow (page 2–1584)*

  Displays the environment but receives shadows. This is a special-purpose material. The effect is similar to using a matte in filmmaking.

- *Morpher (page 2–1592)*

Lets you morph between materials using the *Morpher modifier (page 1–729)*.

- *Multi/Sub-Object (page 2–1594)*

  Lets you apply multiple sub-materials to a single object's sub-objects.

- *Raytrace (page 2–1512)*

  Supports the same kind of diffuse mapping as Standard material, but also provides fully raytraced reflections and refractions, along with other effects such as fluorescence.

- *Shell (page 2–1600)*

  Contains a material that has been *rendered to a texture (page 3–144)*, as well as the original material upon which the texture is based.

- *Shellac (page 2–1597)*

  Mixes two materials by applying a "shellac" material to another.

- *Top/Bottom (page 2–1599)*

  Contains two materials, one for faces that point upward, the other for faces that point downward.

  Standard materials let you set values for components such as color, glossiness, and opacity. They also let you apply maps to the components, which can produce an enormous variety of effects. Some other material types have these features as well. Some materials, such as Multi/Sub-Object or Double Sided, have controls only for combining other materials.

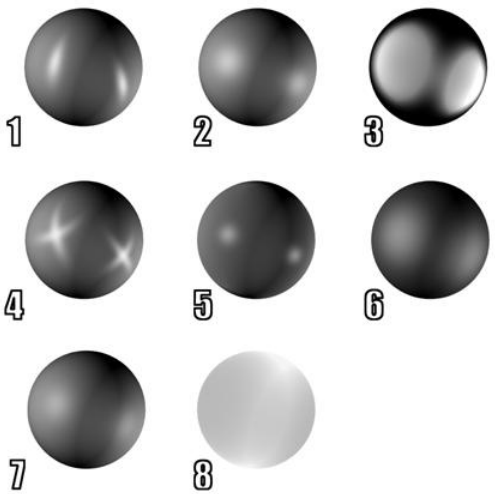
mental ray Materials

A group of materials is provided for use with the *mental ray renderer (page 3–78)*. See *mental ray Materials (page 2–1543)*.

Shading Type

The Standard and Raytrace materials let you specify a shading type. Shading types are handled by a "shader," which describes how the surface responds to light.

Warning: When you change the shading type of a material, you lose the settings (including map assignments) for any parameters that the new shader does not support. If you want to experiment with different shaders for a material with the same general parameters, copy the material to a different *sample slot (page 2–1420)* before you change its shading type. That way, you can still use the original material if the new shader doesn't give you the effect you want.



Samples of different shading for a standard material

1. *Anisotropic*
2. *Blinn*
3. *Metal*
4. *Multi-layer*
5. *Oren-Nayar-Blinn*
6. *Phong*
7. *Strauss*
8. *Translucent*

Several different shaders are available. Some of these are not available for the Raytrace material, as indicated below. Blinn is the most general-purpose of these shaders. The others have special purposes, especially regarding how the material creates highlights.

- *Anisotropic (page 2–1480)*

  Creates surfaces with noncircular, "anisotropic" highlights; good for modeling hair, glass, or metal.

- *Blinn (page 2–1480)*

  Creates smooth surfaces with some shininess; a general-purpose shader.

- *Metal (page 2–1481)*

  Creates a lustrous metallic effect.

- *Multi-Layer (page 2–1481)*

  Creates more complex highlights than Anisotropic by layering two anisotropic highlights.

  Not available for Raytrace material.

- *Oren-Nayar-Blinn (page 2–1482)*

  Creates good matte surfaces such as fabric or terra-cotta; similar to Blinn.

- *Phong (page 2–1480)*

  Creates smooth surfaces with some shininess; similar to Blinn, but doesn't handle highlights (especially glancing highlights) as well.

- *Strauss (page 2–1483)*

  Creates both nonmetallic and metallic surfaces; has a simple set of controls.

  Not available for Raytrace material.

- *Translucent (page 2–1484)*

  Translucent shading is similar to Blinn shading, but it also lets you specify translucency, where light is scattered as it passes through the material. You can use translucency to simulate frosted and etched glass.

  Not available for the Raytrace material.

Material Components

A material's components describe its visual and optical properties. The components in the *Architectural material (page 2–1535)* are based on physical qualities; for example, diffuse color, shininess, transparency, and so on. The components in a *Standard material (page 2–1465)* include color components, highlight controls, self-illumination, and opacity. Like the Standard material, the *Raytrace material (page 2–1512)* uses a nonphysical model to describe surfaces.

Standard and Raytrace material components vary depending on which *shader (page 2–1398)* you use.

You can assign maps to most components, including color components such as Diffuse, and value components such as Transparency or Opacity. Maps can increase the complexity and realism of the material's appearance.

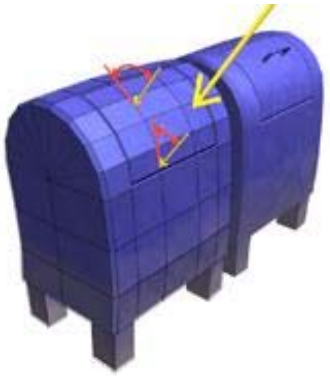
Lights and Shading

Materials work in combination with *lights (page 2–1272)*. The intensity of light that falls on a surface determines the intensity of color to display. Three factors contribute to the intensity of light where it falls on an object:

- **Light intensity:** A light's original intensity at its point of origin.
- **Angle of incidence:** The more a surface inclines away from the light source, the less light it receives and the darker it appears. The angle between a ray of light and the *face normal (page 3–980)* of a surface is the angle of incidence for that face.

  When the angle of incidence is 0 degrees (that is, the light strikes the face perpendicularly), the face is illuminated at full intensity unless

the light is attenuated. Full intensity is the light's Multiplier value times the value of the face's surface color. The Multiplier value is 1.0 by default; the surface value is the Value component of the surface color's *HSV description (page 3–1001)*. As the angle of incidence increases, the intensity of the face illumination decreases.



Angle of incidence affects intensity.

- **Distance:** Light diminishes over distance. This effect is known as *attenuation (page 3–912)*. By default, attenuation is turned off, but you can turn it on and specify the distance over which it operates.

Lights and the Component Colors of a Standard Material

As the names of a standard material's *color components (page 2–1399)* imply, the kind of light that strikes a surface with a material determines how the surface appears when it is shaded.

- **Ambient color** appears where the surface is lit by ambient light alone (where the surface is in shadow).
- **Diffuse color** appears where light falls directly on the surface. It is called "diffuse" because light striking it is reflected in various directions. Highlights, on the other hand, are reflections of light sources.
- **Specular highlights** appear where the viewing angle is equal to the angle of incidence. Glancing highlights appear where the angle of incidence is high, relative to the observer or camera (that is, the light ray is nearly parallel to the surface). Shiny surfaces usually have specular highlights. Glancing highlights are characteristic of metallic surfaces.

  Some surfaces are completely reflective, or nearly so. These reflect their environment as well as the light sources that illuminate them. To model such surfaces, you need to use *reflection mapping (page 2–1699)* or ray tracing (see *Raytrace Material (page 2–1512)*).

The three color components blend at the edges of their regions. Between ambient and diffuse, the blending is calculated by the shader. Between diffuse and specular, you set the amount of blending by using the standard material's highlight controls.

Choosing Colors for Realism

Materials add greater realism to a scene only if you choose their colors and other properties to appear like real-world objects. This topic presents some general guidelines for choosing standard material colors. When possible, you should also observe colors in the objects you are modeling, especially under different lighting conditions.

For objects on which you want the viewer to focus attention, an unmapped standard material doesn't often provide the level of realistic detail you probably want. However, for distant and peripherally visible objects, as well as some kinds of real-world materials, such as molded plastic, an unmapped standard material can work well. Keeping the number of maps to a minimum can help keep down the file size.

Indoor and Outdoor Lighting

Whether a scene is indoors or outdoors affects your choice of material colors, just as it affects the way you set up *lights (page 2–1272)*. Full sunlight is bright and unidirectional. Most indoor lighting is less intense and more even (that is, multidirectional) than daylight. However, some special indoor lighting (and nighttime outdoor lighting), as for the stage, also features intense, directional light.

Direct sunlight has a yellow tint. Materials for objects to appear in daylight should have a specular color of a pale, unsaturated yellow (for example, RGB values of 240, 240, 188). The ambient color should be the complement of the specular: a deep, dark purple with a hint of the diffuse color.

Materials for objects to appear under normal interior lighting should have a specular color that is close to white. (Our perception compensates for the yellow or green tint that is often present in artificial light.) The ambient color can often have the same hue as the diffuse color, but with a darker value.

Materials for objects to appear under spotlights should follow the general guidelines for daylight materials. The specular color should match the spotlight's color, and the ambient color should be a very dark value of the spotlight color's complementary hue, mixed with a bit of the material's diffuse color.

If you want to render an object under changing lighting conditions, you can choose colors that are a compromise between the optimal colors for each kind of lighting, or you can *animate (page 2–1449)* the material so that its colors change to suit the changing light.

Representing Natural Materials



Outdoor scene with natural materials

Most natural materials have a matte surface with little or no specular color. For natural materials such as these, use the following guidelines:

- **Ambient color:** The ambient color depends on whether the scene is indoors or outdoors, as previously described.
- **Diffuse color:** Choose a color found in nature. It is best to use the observed color of the object itself, or a similar object.
- **Specular color:** Make the specular color the same hue as the diffuse, but with a higher value and a lower saturation.
- **Glossiness:** Set the Glossiness to a low value.

  Some foliage, bird feathers, fish scales, and so on, are shiny. For materials such as these, set the Glossiness to higher values. You might also want to change the specular color so it's closer to the lighting color than the surface's diffuse color.

Water is reflective, and is best modeled by a color component in combination with a *reflection map (page 2–1699)* or a *water map (page 2–1683)*.

While metal is a natural material, its special visual characteristics are most apparent when it has been polished. Standard material represents this by

using a special shading type, described later in this topic.

Representing Manufactured Materials



Indoor scene with manufactured materials

Manufactured materials often have a synthetic color rather than an "earth tone." Also, many manufactured materials, such as plastics and porcelain glazes, are very shiny. For manufactured materials, use the following guidelines:

- **Ambient color:** The ambient color depends on whether the scene is indoors or outdoors, as previously described.
- **Diffuse color:** Although the diffuse color doesn't have to be an "earth tone," as with natural materials you should used the observed color of the object or a similar object.
- **Specular color:** Make the specular color close to white, or to the color of the light source. White is especially characteristic of plastic materials.
- **Glossiness:** Set the glossiness to a high value.

Representing Metallic Objects



Metallic cup and ice cream scoop

Polished metal has a characteristic "glancing" highlight that appears where the light is at a high angle of incidence. To generate this effect, Metal shading uses the Cook/Torrance illumination model.

For metallic materials, you can use the Metal shading type. This disables the specular color and highlight controls. The Metal shader calculates its own specular color, which can vary between the diffuse color and the color of the light.

In the diffuse region of a metal material, the ambient component is greater than it is for other kinds of materials.

The Anisotropic, Multi-Layer, and Strauss shaders give you further options for modeling polished metal.

If the metallic object is the focus of the scene, you can improve realism by using a *Blend material (page 2–1588)* to combine metallic shading with a *reflection map (page 2–1699)*.

Tip:  When you preview metallic surfaces, it is useful to turn on a backlight. This displays the metal's glancing highlight. The Backlight button is to the right of the sample slots.

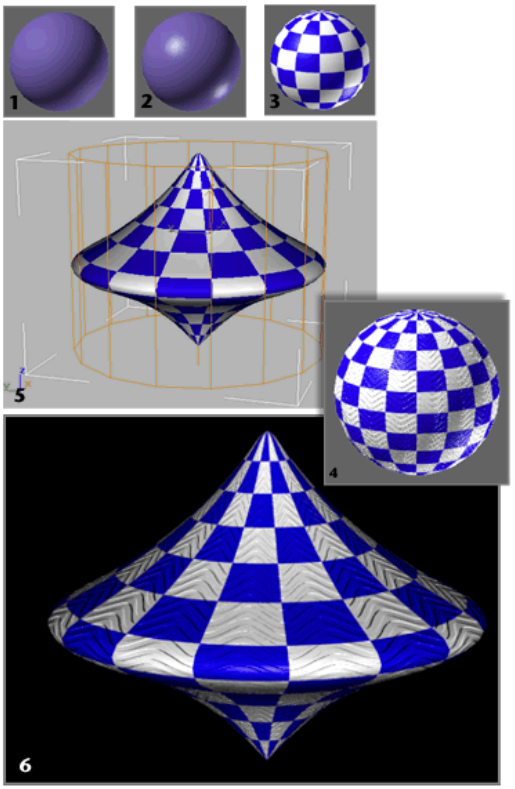
Using Maps to Enhance a Material

Maps provide images, patterns, color adjustments, and other effects you can apply to the visual/optical components of a material. Without maps, material design in 3ds Max is limited. Maps give the Material Editor its full flexibility, and can give you dramatic results.



Spheres with various maps applied to them (as well as a reflection map applied to the surface beneath them)

The simplest use of a map is to assign a pattern to a material's Diffuse color. *Diffuse mapping (page 2–1498)* is also known as "texture mapping." It applies an image or pattern to geometry the material is applied to.



Example of designing a mapped material:

1. *Choose a sample slot.*
2. *Increase the highlight.*
3. *Apply a checker map to the material's diffuse component.*
4. *Apply a bump map to give the material ridges.*
5. *The checker map displays in viewports, but the bump map does not.*
6. *Rendering the material shows the full effect of mapping.*

Warning: When you change the *shading type (page 2–1398)* of a standard material, you lose the settings (including map assignments) for any parameters that the new shader does not support. If you want to experiment with different shaders for a material with the same general parameters, copy the material to a different *sample slot (page 2–1420)* before you change its shading type. That way, you can still use the original material if the new shader doesn't give you the effect you want.

Map Terminology

The term "material map" is sometimes used to describe a map assigned in the material editor. A material map applies a color or pattern to a surface. This is different from maps used for displacement mapping with the *Displace modifier (page 1–629)*, environment mapping for backgrounds, or projection mapping from lights.

The term "texture map" is sometimes used as well. It is interchangeable with "diffuse map"; that is, with a map that applies colors to geometry, as opposed to a map that create reflections, bumps, and so on.

In the *Material/Map Browser (page 2–1412)*, maps are categorized according to how the map software functions. The categories are:

- **2D maps**

  A *bitmap (page 2–1631)* is the prototypical 2D map. 2D maps apply pictures and patterns to the surface of objects.

- **3D maps**

  3D maps are generated procedurally. 3D maps apply patterns throughout an object's geometry.

- **Compositors**

  Compositors combine other maps.

- **Color Modifiers**

  Color modifiers are usually composited with another map to adjust its color. The Vertex Color map is a special case that displays the colors you assign to vertices in a mesh.

- **Other**

  "Other" maps include maps that simulate reflection or refraction.

The names of individual map types describe the pattern or effect they create, such as Checker map, Bitmap, Gradient, Flat Reflection, and so on.

Note: In some cases the user interface also uses "map" to describe not the map type, but the visual component being mapped. For example, a "diffuse map" means a map of any type applied to a material's diffuse component. This is an ambiguity in the use of "map" that can be a bit confusing when you first encounter it.

Assigning Maps

For a *standard material (page 2–1465)*, you assign maps using the Maps rollout. Click the Map button in line with the name of the visual component you want to map. The *Material/Map Browser (page 2–1412)* is displayed. Select the map type (for example, Bitmap) from the list of maps, and then click OK. Double-clicking the map's name in the Browser also assigns the map type.

The Browse From group box in the Browser creates new maps by default. You can also use it to obtain maps from a library (see Saving A Material), from the current scene, from objects selected in the scene, or from elsewhere in the material editor.

In the Browser, you can turn on icons of differing sizes to preview maps before you assign them.

A Standard material's Basic Parameters rollout has shortcut buttons for assigning a map to *some* of the material's visual components. These small buttons are equivalent to the buttons in the Maps rollout. Assigning a map to a button in one rollout changes the corresponding button in the other.

Each type of map has its own set of parameters and controls. If the map is a Checker map, for example, you can choose the colors of the checkers, and whether a checker color has a map of its own. You can change tiling values to affect the scale of the checkers, adjust noise parameters to make the checkers irregular, and so on.

Note: To save loading time, if a map with the same name is in two different locations (in two different paths), it is loaded only once. This poses a problem only if your scene includes two maps that have different content but the same name.

In this case, only the first map encountered will appear in the scene.

Navigating the Material/Map Tree

When you build a material of any complexity, you are building a material/map tree. The root of the tree is the material itself. The branches are the maps you have assigned to the material's components. Some maps can themselves contain maps, as in a map applied to one color of a *Checker material (page 2–1638)*, so the tree can be more than two levels deep, and can actually be as deep as you need it to be.


The *Material/Map Navigator (page 2–1447)* is a dialog that displays the tree for the current material. It is useful for finding a map and displaying its parameters. Click the map to display its rollouts in the Material Editor. To copy a map to a different component of the same material, you can also drag the map's name from the Navigator to a map button in the Material Editor.


The Go Forward To Sibling and Go To Parent buttons also navigate the map tree. Go Forward To Sibling moves laterally in the map tree, while Go To Parent moves up the tree. (To move down the tree, click a map button that has a map assigned to it.) Another way to move between parents and children in the tree is to drop down the *material name field (page 2–1448)* and click the name of a map or material.

Applying a Material to an Object

There are two ways to apply a material to an object:

- If the sample slot is active and the object is already selected, click *Assign Material To Selection (page 2–1441)*.
- Drag from the sample slot to the object.

  As you drag, a tooltip appears over each object beneath the mouse, showing the object's name. You can apply the material whether the object is selected or not. Release the mouse to apply the material.

Applying a material overrides any previous material assignment the object might have had. Once the material is applied, while the sample slot is active, the material is "hot" and changes you make to it affect the object automatically. See *Sample Slots (page 2–1420)* for more about hot and cold materials.

The *Undo command (page 1–94)* works for material assignment.

You can apply only one material to an object. To overcome this restriction, use a *Multi/Sub-Object material (page 2–1594)*. This is a container for various sub-materials that correspond to specified sub-objects such as different faces in a mesh, NURBS surfaces in a NURBS model, and so on.

You can apply the same material to multiple objects in the scene.

See also

Dragging and Dropping Maps and Materials (page 2–1423)

Drag and Drop Sub-Object Material Assignment (page 2–1424)

Mapping Coordinates

An object assigned a 2D mapped material (or a material that contains 2D maps) must have mapping coordinates. These coordinates specify how the map is projected onto the material, and whether it is projected as a "decal," or is tiled or mirrored. Mapping coordinates are also known as UV or *UVW coordinates (page 3–1028)*. These letters refer to coordinates in the object's own

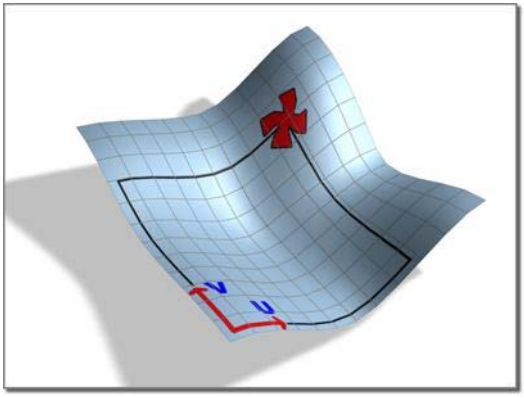
space, as opposed to the XYZ coordinates that describe the scene as a whole.

Most renderable objects have a Generate Mapping Coordinates parameter. If you don't turn this on, but apply a mapped material to the object, you get a warning when you try to render.

Some objects, such as editable meshes, don't have automatic mapping coordinates. For these types of objects, you can assign coordinates by applying a *UVW Map Modifier (page 1–922)*.

If the material appears the way you want it with the default mapping, you don't need to adjust the mapping. If you need to adjust it, use the map's Coordinates rollout. There are two typical sets of coordinates parameters: one for 2D maps such as *Bitmaps (page 2–1631)*, and another for 3D maps such as *Noise (page 2–1674)*. See *Coordinates Rollout (2D) (page 2–1625)* and *Coordinates Rollout (3D) (page 2–1663)*.

Note: The *UVW Remove utility (page 2–1408)* provides a way to remove mapping coordinates or an entire material from selected objects.



Mapping coordinates shown as U and V axes local to a surface.

Note: Mapping coordinates are local to each object, so they are termed *UV or UVW coordinates (page 3–1028)*, as opposed to the XYZ coordinates of the scene as a whole.

Saving a Material

While a material is in the Material Editor or applied to an object, it is part of the scene, and is saved with the scene. However, for complicated scenes it is inconvenient to have all materials active in the Material Editor. You can also save a material by putting it into a material library. Some libraries are provided in the *\matlibs* subdirectory. The file *3dsmax.mat* is the default library. You can add your material to this library, or create your own libraries.

Procedure

To save a material in a library:

This stores a material in the current material library. To use a different library, first open it using the *Material/Map Browser (page 2–1412)*.

1. In the Material Editor, click the sample sphere for the material to save.
2.  On the Material Editor toolbar, click *Put to Library (page 2–1443)*.

   A *Put to Library dialog (page 2–1455)* is displayed. Either change the material name or leave it as is, and then click OK.
3.  Choose Rendering > Material/Map Browser, or, on the Material Editor toolbar, click *Get Material (page 2–1439)*.

   The *Material/Map Browser (page 2–1412)* opens.
4. In the Browser > Browse From group, choose Mtl Library, if necessary.

   The stored material appears in the list.
5. In the Browser > File group, click Save to save the library with the current name (if any) or Save As to save it as a different file.

Tip: You can use the Merge function in the File group to add materials from the current library to another library.

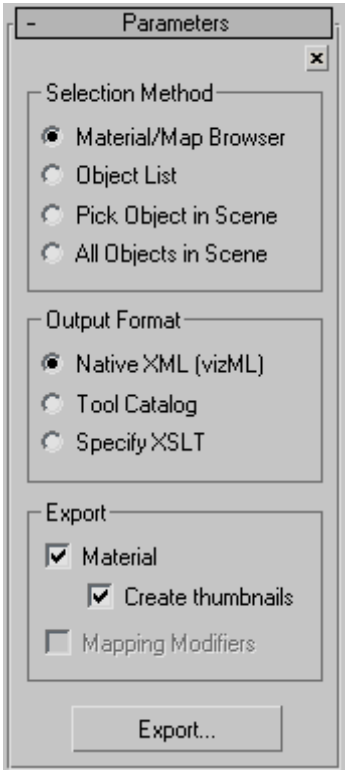
Material XML Exporter Utility

Utilities panel > Utilities rollout > More button > Utilities dialog > Material XML Export

You can export materials you create in 3ds Max to XML files, which can then be shared with other 3ds Max users or used in AutoCAD Architecture (formerly Autodesk Architectural Desktop) to modify material definitions.

Note: You can add an exported XML material to your 3ds Max scene by dragging and dropping from a Web site or Windows Explorer onto an object in your scene, or by importing it directly onto objects.

Interface



Selection Method group

The Selection Method group lets you set the method for selecting the materials you wish to export.

Material/Map Browser—Lets you select a material to export from the *Material/Map browser (page 2–1412)*.

Object List—Lets you select objects from a *Select Objects dialog (page 1–78)*. All materials assigned to the selected objects are exported.

Pick Object in Scene—Lets you select an object from your scene. Any materials assigned to the selected object are exported.

All Objects in Scene—Exports all of the assigned materials in your scene.

Output Format group

The Output Format group defines the format of the XML Material output.

Native XML (vizML)—Materials are exported as raw XML.

Tip: Use this format for sharing XML material files within 3ds Max.

Tool Catalog—Materials are exported to the ATC (Autodesk Tool Catalog) format.

This file type is suitable for display in the AutoCAD Content Browser and the Autodesk VIZ Content Browser.

Note: This file type cannot be imported to 3ds Max unless you also have Autodesk VIZ Render installed on the same system.

Specify XSLT—This option lets you apply your own XSL transform to the XML output.

Export group

The Export group lets you define the parts of the material assignments to export.

Material—Exports the material definitions.

Create Thumbnails—Exports thumbnails for each material.

Note: Thumbnail images are referenced by the ATC and displayed in the AutoCAD Content Browser and Autodesk VIZ Content Browser.

Mapping Modifiers—Exports the mapping modifiers applied to specific objects.

Export—This button begins the XML export process, using the defined selection method, output format, and export parameters.

Upon clicking Export, you are prompted to set the path or URL to store in the XML file as the path to any referenced bitmap files. The default is an empty string, which means no path will be prepended to the bitmap filenames when written to XML. When the material is later imported, it will be assumed that the referenced bitmap file can be found in the bitmap search path.

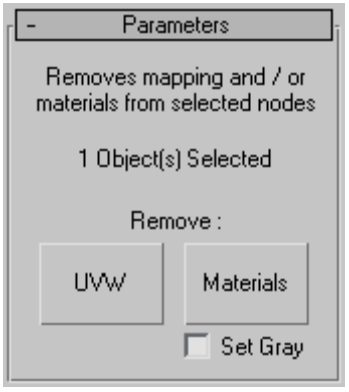
If you have elected to export thumbnails or to apply your own XSL transform, you are then prompted to set a path to store the thumbnail files and to locate your XSLT file, respectively.

UVW Remove Utility

Utilities panel > Utilities rollout > More button > Utilities dialog > UVW Remove

The UVW Remove utility removes mapping coordinates or materials from the currently selected objects.

Interface



UVW—Click to remove UVW mapping from this object.

Materials—Click to remove material assignment from the selected objects.

Set Gray—If this is on when you click the Materials button, the object color is set to a neutral gray. Default=off.

Material Editor

Main toolbar > Material Editor

Rendering menu > Material Editor

Keyboard > **M**

The Material Editor provides functions to create and edit *materials (page 3–971)* and *maps (page 3–968)*.

Materials create greater realism in a scene. A material describes how an object reflects or transmits light. Material properties work hand-in-hand with light properties; shading or rendering combine the two, simulating how the object would look in a real-world setting.

You apply materials to individual objects or selection sets; a single scene can contain many different materials.

Note: Creating a new material clears the Undo/Redo lists.

Procedures

To view the Material Editor:

-  Click the Material Editor button on the main toolbar.

  The Material Editor dialog has *sample slots (page 2–1420)* for viewing previews of materials. When you first view the Material Editor, the material previews have a uniform default color.

To give a material a different name:

- Edit the name field that appears below the Material Editor toolbar.

  The name of the active material appears in the title bar of the Material Editor dialog. The name of the material is not a file name: it can contain spaces, numbers, and special characters.

  The name field displays only 16 characters, but a material name can be longer than that.

To make a copy of a preview material:

-  On the Material Editor toolbar, click *Make Material Copy (page 2–1442)*.

To get a material from a scene:

If a material that you want to change has been saved in the scene, but not in the Material Editor, you can load the material by getting it from the scene.

1. Click a sample slot to make it active.

   Be careful not to click the sample slot of a material you want to use later.

2.  On the Material Editor toolbar, click *Get Material (page 2–1439)*.

   A modeless *Material/Map Browser (page 2–1412)* is displayed.

3. In the Browse From group box at the upper left, make sure that either Selected or Scene is chosen.

   The Selected option lists only materials in the current selection. If no objects are selected, the list of materials is blank.

   The Scene option lists all the materials currently in the scene.

4. In the list of materials, double-click the name of the material you want.

   You can also drag the material name to the sample slot.

   The material you chose replaces the previous material in the active sample slot.

Warning: When you get a material from a scene, it is initially a *hot material (page 3–953)*.

To apply a material to objects in a scene:

1. Select the sample slot that contains the material you want to apply.
2. Select the object you want to apply the material to.
3. Drag from the sample slot to the object. If you selected more than one object, you are asked if you want to apply to the single object only, or to the whole selection.

 You can also apply materials by clicking *Assign Material To Selection (page 2–1441)* on the Material Editor toolbar.

Warning: When you apply a material to an object or selection, that material becomes a *hot material (page 3–953)* (its sample slot is displayed with white corner brackets). When you change the properties of a hot material, the scene immediately updates to reflect those changes. *Any* object with that material will change its appearance, not just the objects in the current selection.

To remove a material from an object:

1.  On the Material Editor toolbar, click Get Material.

   The Material/Map Browser appears.
2. Drag the entry NONE from the top of the list in the Browser to the object.

   The object now has no material applied to it.

 To make a material no longer hot so it doesn't change the current scene, click *Make Material Copy (page 2–1442)*.

To put a material back into a scene:

-  On the Material Editor toolbar, click *Put Material To Scene (page 2–1440)*.

  The material in the active sample slot is now a *hot material (page 3–953)*.

The Put Material button is available only when (1) the material in the active sample slot has the same name as a material in a scene, and (2) the material in the active sample slot is not hot. In other words, this command is meant to fit into the following overall sequence of handling materials:

- You create a hot material either by applying it to objects in the scene or by getting it from the scene.
- You make a copy of the material.
- You make changes to the copy of the material.
- You update the scene by putting the changed material back into the scene.

  These steps are not as immediate as changing a material while it is hot, but they help you avoid changing the scene's materials unintentionally or in unexpected ways.

  When a material in the Material Editor is applied to objects in the scene, you can select the objects from the Material Editor.

To select objects that have the same material applied:

When a material in the Material Editor is applied to objects in the scene, you can select the objects from the Material Editor.

1. Click a sample slot that contains a material in the scene.

   White corner brackets indicate materials that are in the scene.
2.  Click *Select By Material (page 2–1439)*.

   This button is unavailable unless the active sample slot contains a material in the scene.

   A *Select Objects dialog (page 1–78)* is displayed. The names of objects with the active material applied are highlighted when the dialog appears.

3. Click Select to select objects that have the active material applied to them.

   You can also change the selection by choosing other objects. If you change the selection, you must then click *Assign Material To Selection (page 2–1441)* to apply the active material to newly selected objects.

To get a material from a library:

1.  On the Material Editor toolbar, click *Get Material (page 2–1439)*.

   A modeless *Material/Map Browser (page 2–1412)* is displayed.

2. In the Browse From group box at the upper left, make sure that Material Library is chosen.

   If you have opened a library, the list of materials shows the contents of the library.

   If you haven't opened a library, click Open in the file area of the Browser. In the file dialog that is displayed, you can choose a material library to open. After you open the library, the list of materials updates to show the library's contents.

3. In the list of materials, double-click the name of the material you want.

   You can also drag the name of the material to the sample slot.

   The material you chose replaces the previous material in the active sample slot.

To save a material in a library:

1. Click to select the sample slot that has the material you want to save.

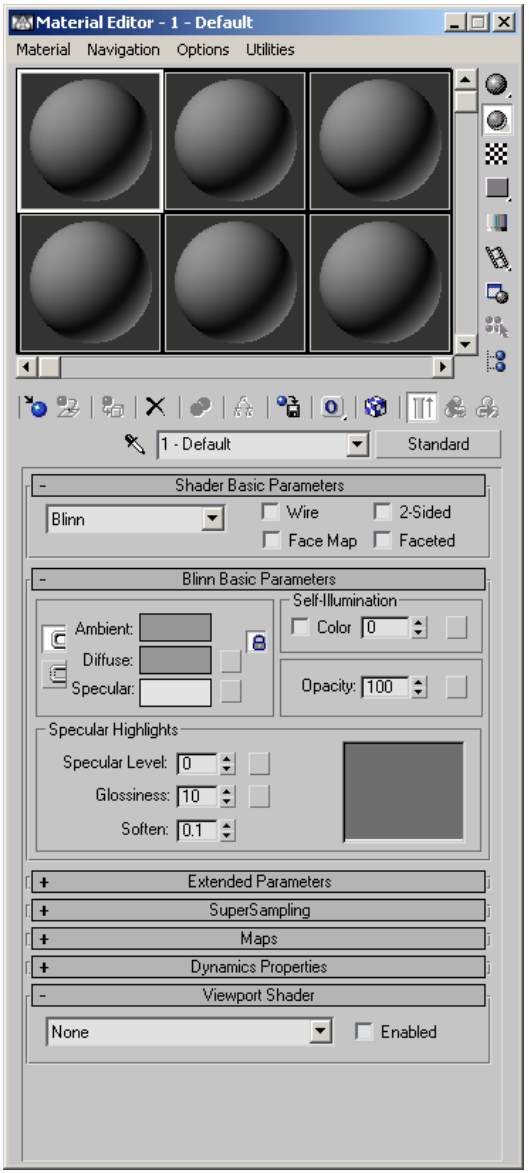
2.  On the Material Editor toolbar, click *Put To Library (page 2–1443)*.

3. A *Put To Library dialog (page 2–1455)* appears.

4. Either change the material name or leave it as is, and then click OK.

   The material is saved in the currently open library. If no library is open, a new library is created. You can save the new library as a file using the *Material/Map Browser (page 2–1412)* file controls.

Interface



For choosing materials, see *Material/Map Browser (page 2–1412)*. For applying materials using drag and drop techniques, see *Dragging and Dropping Maps and Materials (page 2–1423)*.

For an overview of how to use the Material Editor, see *Designing Materials (page 2–1395)*.

For the user interface elements in Material Editor and the different materials and map types, see the following topics:

Material Editor Menu Bar (page 2–1428)

Sample Slots (page 2–1420)

Material Editor Tools (page 2–1427)

Types of Materials (page 2–1457)

Types of Maps (page 2–1617)

For information about how to animate materials, see *Animating Materials (page 2–1449)*.

Material/Map Browser

Main toolbar > Material Editor > Get Material > Material/Map Browser
Rendering menu > Material/Map Browser
Material Editor > Click Material Type button or Map Type button. > Material/Map Browser

Procedure (page 2–1413) Interface (page 2–1414)

The Material/Map Browser lets you choose a *material (page 3–971)*, *map (page 3–968)*, or *mental ray shader (page 2–1710)*.

When you click *Get Material (page 2–1439)*, the Browser that is displayed is modeless (you can leave it displayed while you do other work). However, when you display the Browser by clicking the *Type (page 2–1449)* button, a map assignment button in the *Environment dialog (page 3–272)*, or from a projector light (see *Advanced Effects Rollout (page 2–1341)*), it appears as a modal dialog with OK and Cancel buttons.

You can leave the modeless Browser displayed, and drag materials from its listings to material or map sample slots and buttons in the user interface. When the Browser displays a Material Library, you can also add materials to the library by dragging them from the Material Editor *sample slots (page 2–1420)*.

When you double-click a material, map, or shader in the Browser, it places that material, map, or shader in the Material Editor's active sample slot. It automatically chooses between an instance or a copy, as follows:

- Browsing New Materials: Creates a new material.
- Browsing a Library: Makes a copy.
- Browsing the Material Editor, Scene, or Selected: Depends on the status of the map or material.

When you browse the Material Editor, the Scene, or Selected objects, the choice between making a copy or an instance depends on the status of the material, as follows:

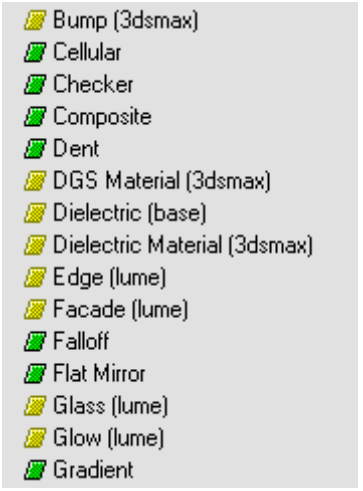
- If the material or map is already in the active slot, the Browser does nothing.
- If the material or map is in some other sample slot, the Browser puts a copy in the active slot.
- In all other cases, the Browser makes an instance of the material or map.

Browsing mental ray Materials and Maps

When you use the *mental ray renderer (page 3–78)*, you might want to use the materials and shaders that provide effects for this renderer only. (The default scanline renderer renders these materials and shaders only as black or white, or it simply ignores their effects.) The Material/Map Browser lists mental ray maps and materials only if you

assign the mental ray renderer as the currently active renderer.

Once you have enabled the renderer, when you use the Browser, it shows mental ray materials and shaders. Materials are displayed with a yellow sphere, rather than blue for standard materials, and shaders are displayed with a yellow parallelogram, rather than green as for standard maps.



mental ray maps in the browser's list are shown with yellow icons.

When you use the *mental ray Connection rollout (page 2–1461)*, or other shader buttons specific to mental ray materials and shaders, the shaders that appear in the Browser's list are restricted to those that the mental ray renderer allows for that particular shader component. By default, only shaders that ship with 3ds Max are listed. If you have acquired other shader libraries, you might see the names of shaders that are not mentioned in this reference.

Note: You can see the listing of materials, maps, or shaders that are incompatible with the current renderer, if you turn on the Incompatible toggle in the Show group, as described under "Interface," below.

See also

Material/Map Navigator (page 2–1447)

Procedures

To navigate materials with the Browser:

Tip: Use the Browser primarily in Root Only mode, to see only the top levels of the materials. This provides a simpler view of your materials, and speeds redraws when you're using any of the icon display modes. (You can also create thumbnails to speed up redraws, as described below.)

1. In the Material Editor, choose any sample slot you want that contains a complex, multilevel material.
2. In the Browser > Browse From group, choose Active Slot mode to display all levels of the active sample slot.
3. Click any of the items in the Browser's material/map list to move to that level of the current material.
4. When you want to switch to a different material, select its sample slot in the Material Editor, and its hierarchy will appear in the Browser.
5. Again, click the items in the Browser to change levels.

To delete an assigned map:

1. While viewing the map parameters, click the Map Type button.
2. In the Material/Map Browser, choose NONE as the new map type.

   Tip: You can also remove a material or map by dragging the NONE item from the Browser over to the object or map button.

To merge material libraries:

1. In the Browse From group, choose Mtl Library, and then click the Merge button.

2. In the *Merge Material Library dialog (page 2–1454)*, select a material library other than the current library, or select a 3ds Max or **VIZ Render** (*DRF (page 3–527)*) scene.

   A *Merge dialog (page 2–1453)* is displayed, listing all materials in the specified library, or all materials assigned to the 3ds Max or **VIZ Render** file. Below the list are All and None buttons to help in the selection.

3. Select the materials in the list that you want to merge, and then click OK.

   The selected materials are merged into the current material library.

4. Save the library to save your changes.

To save the sample spheres as thumbnail images:

1. Open the Browser. In the Browse From group, choose Mtl Library.
2.  Choose View Small Icons.
3. Display all of the icons in the library by either scrolling through all of them, or by enlarging the Browser so that all of the icons have been displayed at least once.

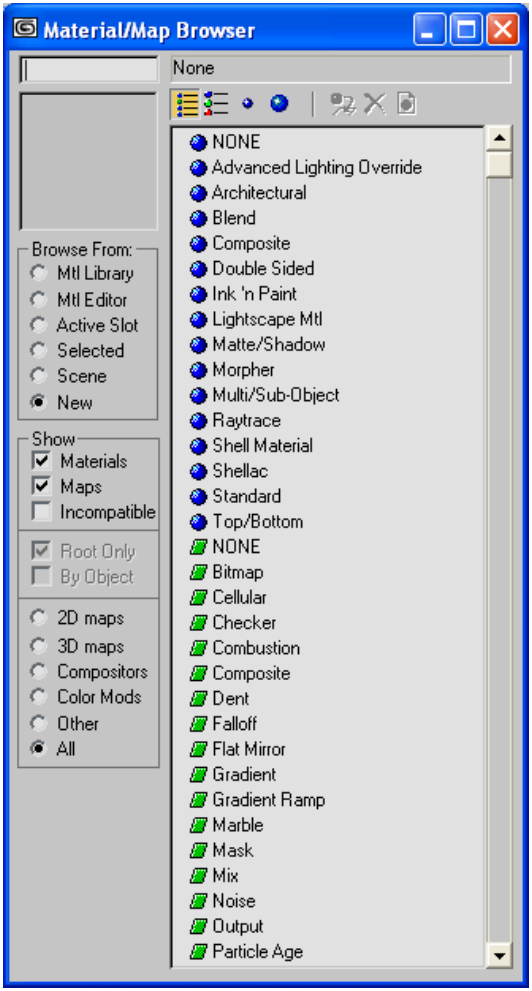
   The action of displaying the icons automatically creates thumbnails in memory.

   Important: If you want to include thumbnails of the sub-materials and maps, be sure to turn off Root Only.

4. Save the library.

   When you save the library, you save the thumbnail images of the samples as they appeared in the Browser at that time. If you change any of the materials or maps later, you must re-save the library in order to update the thumbnails. If you do not re-save the library after altering or adding a materials, the icon of the material will still appear correctly, but it will be rerendered when it first appears in the Browser, while all the other icons will appear immediately.

Interface



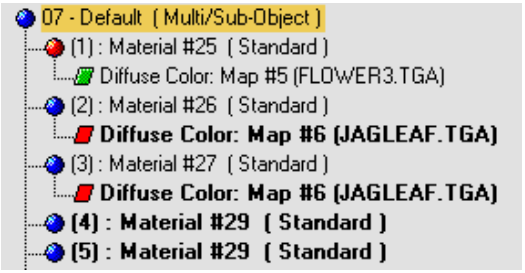
The Material/Map Browser contains the following controls:

Material/Map list—The main part of the Material/Map Browser dialog is a scrollable list of materials and maps. The list indicates a material with a blue sphere, and indicates a map with a green parallelogram. When you list both materials and maps, the materials are listed first.

Note: Icons of materials and maps for which *Show Map in Viewport (page 2–1445)* is on are red.

Also, the names of *instanced (page 3–957)* materials and maps appear in boldface.

Both of these effects are shown in the following illustration:



Text entry—As you enter a material name in this field, the first matching text item is selected in the list. Press Enter to select the next matching name, and so on.

For example, if you enter **ch** when the list includes the material names Cherry Red, Chrome Blue, and Chrome Zinc, Cherry Red is selected first. Press Enter, and the Chrome Blue is selected. Press Enter again, and Chrome Zinc is selected.

The search is not case-sensitive.

Sample slot—Below the text-entry field is a single sample slot. This displays a sample of the current selection. You can drag the sample to any other sample slot or material button. The sample slot display is interruptible, so you can quickly click from one list item to the next without waiting. In addition, if you complete the display of one sample, then move on to another sample, when you return to the first sample, it displays instantly.

Tool buttons

The first part of this row of buttons controls how you view the list. The second part is for managing material libraries.

To speed up the display of the sample spheres in the Browser, the smaller of the sample spheres (those displayed when you choose View Small Icons or View List + Icons) can be saved as thumbnail images in the material library file. (See the Procedures for this topic, above.)

Keep in mind that the saved thumbnails increase the size of the material library file.

View List—Displays the materials and maps in list format. Blue spheres are materials. Green parallelograms are maps. The green parallelograms turn red if Show Map in Viewport is on for a material.

View List + Icons—Displays the materials and maps in a list with small icons.

View Small Icons—Displays the materials and maps as small icons. As you move the mouse over the icons, tooltip labels pop up, showing you the name of the material or map.

View Large Icons—Displays the materials and maps as large icons.

The large icons are labeled with the name of the material or map and are displayed using progressive refinement—samples are rendered quickly, with large pixels, then rendered a second time in greater detail.

Update Scene Materials from Library—Updates materials in the scene with the materials of the same name stored in the library.

When you click Update Scene Materials from Library, the *Update Scene Materials dialog (page 2–1457)* is displayed. This dialog lists materials in the library that have the same name as materials in the scene. In the list, select the materials you want to update in the scene, and then click OK.

If no materials exist in the scene that match the names in the library, an alert informs you of this.

This button is available only when the Browser is viewing a library.

 **Delete from Library—**Removes the selected material or map from the library display. The library on disk is not affected until you save it. Use Open to reload the original library from disk. This button is only active when you select a named material that exists in the current library.

This button is available only when the Browser is viewing a library.

 **Clear Material Library—**Removes all materials from the library display. The library on disk is not affected until you save it. Use Open to reload the original library from disk.

This button is available only when the Browser is viewing a library.

Browse From group

The controls in this group choose the source of the materials displayed in the material/map list.

Material Library—Displays the contents of a material library file from disk. When you set this option, the buttons under File become active (see below).

You can also load a library from a MAX file. When browsing from the Material Library in the Material/Map Browser, choose Open, and then choose 3ds Max (*.*max*) from Files of type. Select and load a *.max* file. All materials assigned in that scene are listed in the Browser. To convert the collection of materials to a library file, click Save, and save it as a MAT (*.mat*) file.

Material Editor—Displays the contents of the sample slots.

Active Slot—Displays the contents of the currently active sample slot.

This option is unavailable in the modal version of the Browser.

When you choose this mode, all check boxes in the Show group box are made available. The entire material and map tree of the active material is displayed, regardless of the state of these check boxes in other Browse From modes.

You can also use Active Slot mode to navigate the hierarchy of the active material. When Active Slot is chosen, clicking an item in the material/map list moves Material Editor controls to that level of the material.

Selected—Displays the material applied to the selected objects.

Scene—Displays all materials applied to objects in the scene.

All maps assigned to the scene, including Environment Background or spotlight projector maps, are displayed in the Browser list.

New—Displays the set of material/map types for you to create a new material.

Show group



These options filter what is displayed in the list. Either Materials or Maps is always on, and both can be on at the same time. The first two options can be unavailable, depending on the active Browse From and View settings.

Materials—Turns display of materials and sub-materials on or off.

This is always unavailable in the modal version of the Browser.

Maps—Turns display of maps on or off.

This is always unavailable in the modal version of the Browser.

Incompatible—When on, displays materials or maps and shaders that are incompatible with the currently active renderer. The incompatible materials are displayed in gray. You can still assign incompatible materials, maps, or shaders to buttons where they would be legal, but if you use the current renderer, the results might not be correct. Default=off.

Root/Object group

Root Only—When on, the material/map list displays only the root of the material hierarchy. When off, the list displays the full hierarchy.

The default state of Root Only depends on how you display the Browser. Generally, when you display the modeless Browser, you're selecting materials rather than maps (to begin with), so Root Only is on. However, when you display the modal Browser (by clicking a map button anywhere in the user interface), Root Only is off so you can see all the maps.

By Object—This is available only when you're browsing from either Scene or Selected. When on, the list displays materials by their object assignment in the scene. At the left are the names of the objects arranged alphabetically, with a yellow cube icons as in *Track View (page 2–501)*. Applied materials are shown as children of the objects. When off, the list displays only material names.

File group



This button group is displayed when you've chosen Material Library, Material Editor, Selected, or Scene in the Browse From group. All four buttons are displayed only when browsing from the Material Library; otherwise, only the Save As button appears.

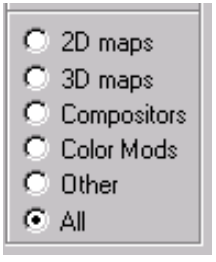
Open—Opens a material library.

Merge—Merges materials from another material library or scene. When you click Merge, the *Merge Material Library dialog (page 2–1454)* is displayed. This file dialog lets you choose a material library or a scene. When you choose a library or scene to merge, the *Merge dialog (page 2–1453)* is displayed. This lets you select which materials to merge. If there are duplicate names among the materials you're merging, the *Duplicate Name dialog (page 2–1453)* is displayed so you can resolve the name conflicts.

Save—Saves the open material library.

Save As—Saves the open material library under another name.

Display group



This group of radio buttons is displayed only when you've chosen New under Browse From. It controls what types of maps the Browser displays in the material/map list. (The Browser displays materials regardless of this setting.)

2D Maps—Lists only 2D map types.

3D Maps—Lists only 3D *(procedural) (page 3–997)* map types.

Compositors—Lists only compositor map types.

Color Mods—Lists only color modifier map types.

Other—Lists reflection and refraction map types.

All—(The default.) Lists all map types.

Copying and Pasting: Right-Click Menu for Materials, Maps, Bitmaps, and Colors

Material Editor > Right-click a Type button, sub-materials button, map button, bitmap button, or color swatch.

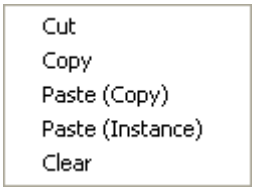
Elsewhere in the user interface > Right-click a map button or color swatch.

A set of right-click pop-up menus in the Material Editor (and elsewhere in the program's user interface) lets you copy and paste, and otherwise manage materials, maps, bitmaps, and colors.

You see these menus only if copy and paste actions are appropriate. For example, if you copy a material and then right-click a map button, nothing happens.

Material Right-Click Menu

When you right-click a button that represents a material, this menu appears. This includes the *Type button (page 2–1449)* for a material, and sub-material buttons such as you find in the *Multi/Sub-Object material (page 2–1594)*, the *Blend material (page 2–1588)*, and others.



Cut—Makes a copy of the material. If you right-click a sub-material button, this also removes the material from that sub-material component. If you right-click the Type button, Cut is equivalent to Copy.

Copy—Makes a copy of the material.

Paste (Copy)—Pastes a copy from the copy buffer.

This item doesn't appear if you haven't yet copied a material.

Paste (Instance)—Pastes an instance from the copy buffer.

This item doesn't appear if you haven't yet copied a material.

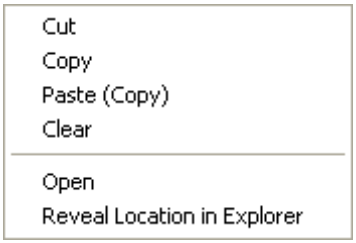
Clear—If you right-click a sub-material button, Clear removes the material from that sub-material component without making a copy of it. If you right-click the Type button, Clear has no effect.

Map Right-Click Menu

When you right-click the button for a map component (a "map slot"), the menu you see depends on whether a map has been assigned yet.

When a Map Has Been Assigned

The following menu appears if a map has been assigned:



Cut—Removes the assigned map, and puts a copy of it in the copy buffer.

Copy—Copies the map without removing it.

Paste (Copy)—Pastes a copy from the copy buffer.

This item doesn't appear if you haven't yet copied a map.

Clear—Removes the assigned map without copying it.

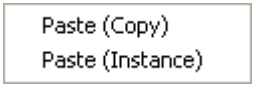
Open—For *bitmaps (page 2–1631)*, launches whichever application is associated with the currently assigned 2D map.

This item does not appear for other map types such as procedural maps.

Reveal Location in Explorer—Launches a copy of Windows Explorer to display the folder where the map is saved.

When No Map Has Been Assigned

If no map has been assigned, all you can do is paste another:



Paste (Copy)—Pastes a copy from the copy buffer.

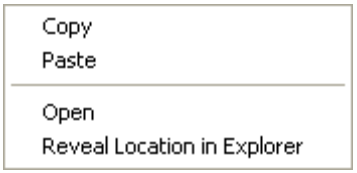
Paste (Instance)—Pastes an instance from the copy buffer.

When No Map Has Been Either Assigned or Copied

If no map has been assigned and you haven't copied a map yet, then no right-click menu appears at all.

Bitmap Right-Click Menu

This menu appears when you click a button that specifies an external *bitmap (page 3–917)*. See *Bitmap 2D Map (page 2–1631)*.



Copy—Copies the bitmap.

Paste—Pastes the bitmap from the copy buffer.

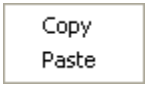
Open—Launches whichever application is associated with the bitmap.

Reveal Location in Explorer—Launches a copy of Windows Explorer to display the folder where the bitmap is saved.

Color Swatch Right-Click Menu

This menu appears when you right-click a color swatch.

Note: This menu does not apply to the VertexPaint modifier's *Color Palette (page 1–950)*, which has its own right-click menu.

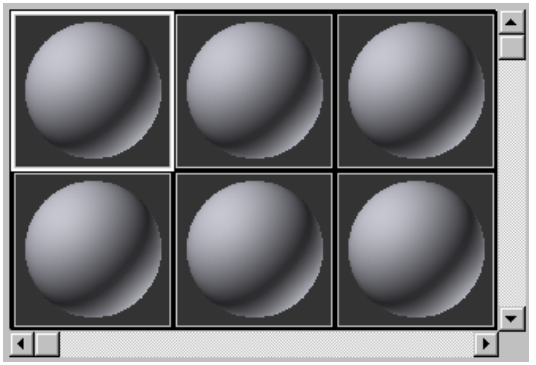


Copy—Copies the color.

Pastes—Pastes a color from the copy buffer.

Sample Slots

Material Editor > Sample slots display



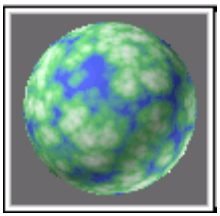
The sample slots let you maintain and preview *materials (page 3–971)* and *maps. (page 3–968)* Each slot previews a single material or map. You can change the material by using the Material Editor controls, and you can apply the material to objects in the scene. The easiest way to do this is to drag the material from the sample slot to objects in viewports. See *Dragging and Dropping Maps and Materials (page 2–1423)*.

Important: While the Material Editor can edit no more than 24 materials at a time, a scene can contain an unlimited number of materials. When you are through editing one material, and have applied it to objects in the scene, you can use that sample slot to get a different material from the scene (or create a new one) and then edit it.

You can display a sample slot in a window of its own. This magnifies the sample slot, which can make it easier to preview the material. You can resize the magnified window to make it even larger. To magnify a sample slot, double-click it, or right-click and choose Magnify from the pop-up menu. See *Sample Slot Right-Click Menu (page 2–1422)*.

The Material Editor has 24 sample slots. You can view them all at once, six at a time (the default), or 15 at a time. When you view fewer than 24 slots at once, scroll bars let you move among them. See *Material Editor Options (page 2–1436)* and *Sample Slot Right-Click Menu (page 2–1422)*.

A material in a slot is shown on a sample object. By default, the object is a sphere. Use the *Sample Type flyout (page 2–1432)* to change the sample object.



Sample slot showing a material

By default, a standalone map in a slot fills the whole slot. This is when the slot shows *only* a standalone map at the top of a tree; when the map is assigned to a material, the slot shows it as part of the material, mapped to the sample object. See *Get Material (page 2–1439)* and *Material Editor Options (page 2–1436)*.



Sample slot showing a map

The Material Editor renders only the active sample sphere for the current frame.

If the *\matlibs* subdirectory contains a material library called *medit.mat*, the sample slots show the first 24 materials in this library file. If the library contains fewer than 24 materials, the remaining slots contain Standard materials of various colors, as they do if an *medit.mat* library is not found.

Hot and Cool Materials

A sample slot is *"hot" (page 3–953)* when the material in the slot is assigned to one or more surfaces in the scene. When you use the Material Editor to adjust a hot sample slot, the material in the scene changes at the same time.

The corners of a sample slot indicate whether the material is a hot material:

- No triangle: The material is not used in the scene.
- Outlined white triangle: The material is hot. In other words, it's instanced in the scene. Changes you make to the material in the sample slot will change the material displayed in the scene.
- Solid white triangle: The material is not only hot, but is applied to the currently selected object.



Left: "Hot" material applied to currently selected object.

Middle: "Hot" material assigned to scene but not to currently selected object.

Right: "Cool" material: active but not assigned to scene.

A material is *"cool" (page 3–924)* if it is not applied to any object in the scene.

To make a hot sample slot cool, click *Make Material Copy (page 2–1442)*. This copies the material in the sample slot on top of itself so that it's no longer used in the scene.

You can have the same material (with the same name) in more than one sample slot, but only one slot containing that material can be hot. You can have more than one hot sample slot, as long as each has a different material.

If you drag to copy a material from a hot slot to another slot, the destination slot is cool, and the original slot remains hot.

See also

Sample Slot Right-Click Menu (page 2–1422)

Dragging and Dropping Maps and Materials (page 2–1423)

Creating a Custom Sample Object (page 2–1425)

Procedures

To use a sample slot:

- Click the sample slot to make it active.

  The active sample slot is displayed with a white border around it.

  The sample slot shows a sample object shaded with a material. (By default, the sample object is a sphere.) The sample object is lit by a light source above it and slightly toward the viewer. For the sphere, the highlight is in the upper-left quadrant. The diffuse color shows most clearly above and to the left of the highlight, shading toward the ambient color at the sphere's lower right.

To change the preview shape:

1. Make sure the sample slot of the material you want to view is active.
2. Use the *Sample Type flyout (page 2–1432)* to choose the shape you want to view. The flyout gives you three options: sphere (the default), cylinder, or box.

   The new shape is displayed in the sample slot, with the material mapped to it.

To render the current mapping level:

1. Move to the level of the map hierarchy that you want to render.

2. Right-click in the sample slot, and choose Render Map from the pop-up menu.

   The *Render Map dialog (page 2–1455)* is displayed.

3. Choose Single or the range of frames you want to render.

4. In the Dimensions group box, specify the pixel resolution of the map.

5. Click the Files button, and specify a path and file name for the file. Make sure Save To File is on unless you want to see the image only in a *rendered frame window (page 3–5)*.

6. Click Render.

   A rendered frame window appears displaying the map. If Save To File is on, the image is also saved to disk.

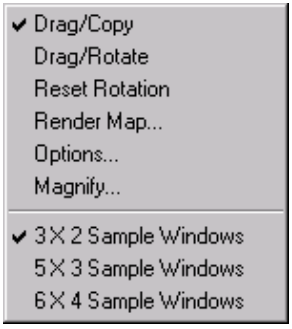
Sample Slot Right-Click Menu

Material Editor > Right-click active sample slot.

When you right-click the active sample slot, a pop-up menu is displayed. For other sample slots, click or right-click once to select them, then right-click to use the pop-up menu.

The pop-up menu is available in magnified sample slot windows. See the "Magnify" option, below.

Interface



The pop-up menu has these options:

Drag/Copy—Sets dragging a sample slot to copy mode. When on, dragging a sample slot copies the material from one sample slot to another, or from the sample slot to an object in the scene or a material button.

Drag/Rotate—Sets dragging a sample slot to rotate mode. When on, dragging in the sample slot rotates the sample object. This helps you preview the material. Drag over the object to rotate it about its X or Y axis; drag in the corners of the sample slot to rotate the object about its Z axis. In addition, if you press Shift before dragging in the center, the rotation is constrained to the vertical or horizontal axis, depending on your initial drag direction.

If you have a three-button mouse and are using Windows NT, the middle mouse button rotates the sample object while Drag/Copy mode is active.

Reset Rotation—Resets the sample object to its default orientation.

Render Map—Renders the current map, creating a bitmap or an AVI file (if the map is animated). Only the current map level renders. That is, the rendering shows the image displayed when Show End Result is turned off.

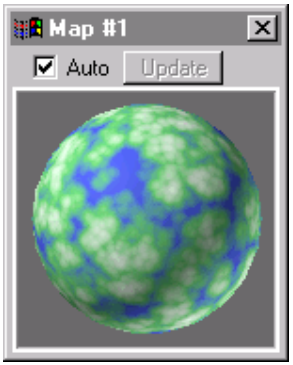
If you're at a material level instead of a map level, this menu item is disabled.

Options—Displays the *Material Editor Options dialog (page 2–1436)*. This is an alternative to clicking the Options button.

Magnify—Generates a magnified view of the current sample slot. The magnified sample is displayed in its own, floating (modeless) window. You can display up to 24 magnified windows, but you can't display the same sample slot in more than one magnified window at a time. You can resize magnify windows. Clicking a magnify window activates the sample slot, and vice versa.

Shortcut—Double-click a sample slot to display the magnified window.

The magnify window's title bar displays the contents of the editable *material name field (page 2–1448)*. It varies depending on which level of the material is active.



Auto—Turn off to prevent the magnified window from updating automatically. This can save rendering time, especially when you have resized the magnified window to make it larger. Default=on.

Update—Click to update the magnified window. This button is unavailable unless Auto is turned off.

Dragging a different sample slot to the magnify window changes the contents of the magnify window.

Sample Windows options

The Material Editor always has 24 sample slots available. You can choose to display fewer sample slots at a larger size. When you do, scroll bars let you move around among the sample slots.

3 X 2 Sample Windows—Displays a 3 x 2 array of sample slots. (The default: 6 windows.)

5 X 3 Sample Windows—Displays a 5 X 3 array of sample slots. (15 windows.)

6 X 4 Sample Windows—Displays a 6 X 4 array of sample slots. (24 windows.)

Dragging and Dropping Maps and Materials

You can move materials from sample slots to objects using a drag-and-drop operation. You can also drag to and from map and material buttons. See the following lists to determine where in the interface you can drag from and drag to.

When dragging materials, use the *Undo command (page 1–94)* to cancel material assignments.

Where You Can Drag From

- *Sample Slots (page 2–1420)*

  The content you drag from a sample slot is always at the top level of the sample slot, regardless of which level is currently displayed. If the sample slot contains a material, you can't drag from the sample slot to a map button, even if you're at the map level of the material.
- The Material Editor *Type button (page 2–1449)*
- *Material/Map Browser (page 2–1412)* lists (text or icon lists)

  You can't drag from the modal version of the Browser (when OK and Cancel buttons are present).
- The sample slot in the Browser
- Material map buttons (see below)
- A projector light map button (see *Advanced Effects Rollout (page 2–1341)*)
- The *Environment Background (page 3–272)* map button
- *Fog (page 3–282)* Color and Opacity map buttons
- *Displace modifier (page 1–629)* map buttons

Material Map Buttons

The material map buttons you can drag from include:

- The buttons in the *Maps rollout (page 2–1474)*
- The small shortcut map buttons on the *Basic Parameters rollouts (page 2–1470)*.
- Any map buttons at any level.
- Sub-material buttons, such as those found in the *Multi/Sub-Object material (page 2–1594)*.

Where You Can Drag To

- Objects in viewports.

  Drag from a material button, sample slot, or Browser listing into the viewport and over an object. When you release the mouse, the material is applied.

  If you drag a material over two or more selected objects, an alert message asks if you want to apply the material to the object or to the selection. Choose the option you want, and click OK.
- The Material Editor *Type button (page 2–1449)*.

  You can drag to the Type button only from the Browser. If the Type button shows a standalone map, you can drag only a map to it. If it shows a material type, you can drag only a material to it.
- All of the items in the previous list, except that you can drag to the Browser only when it displays a material library.

  When the Browser is set to browse from a Material Library, dragging materials and maps into the Browser adds them to the library.
- A Face, Polygon, Patch, or Element sub-object selection of an editable surfacemesh object (mesh, patch, or poly).
- A Face, Polygon, Patch, or Element sub-object selection created by the *Edit Mesh modifier (page 1–634)* or *Edit Patch (page 1–638)*; or by one of these selection modifiers: *Mesh Select (page 1–719)*, *Patch Select (page 1–751)*, or *Poly Select (page 1–762)*.

See also

Applying a Material to an Object (page 2–1405)

Drag and Drop Sub-Object Material Assignment (page 2–1424)

Drag and Drop Sub-Object Material Assignment

You can apply a material to a selection of renderable sub-objects, such as faces in a mesh. In the Material Editor, you can use *Assign Material to Selection (page 2–1441)*. You can also drag the material from the Material Editor or the *Material/Map Browser (page 2–1412)* to the selected faces. This can create a new *Multi/Sub-Object material (page 2–1594)* on the fly.

You can disable drag-and-drop of materials to sub-objects. To do so, go to the *General tab (page 3–815)* of the *Preferences dialog (page 3–815)*, and in the Sub-Materials group, turn off Assign Automatically. This check box is on by default.

Procedure

To drag materials onto sub-object selections:

1. In the Modify panel > Modifier Stack rollout, choose Face as the sub-object level.
2. Select faces of an editable mesh object.
3. Drag a material from a Material Editor sample slot to the selected faces.
4. In the Modifier Stack rollout, click to turn off Sub-Object and return to the object level.
5.  On the Material Editor, click Pick Material From Object, then use the eyedropper to get the material from the sphere.

The new Multi/Sub-Object material appears in the active sample slot.

How the Multi/Sub-Object Material Is Created

The Multi/Sub-Object material is created in one of three ways, depending on what material is already applied to the selected sub-objects:

- **No material applied**

  If the selected faces have no material applied, a new Multi/Sub-Object material is created. The dragged material becomes a sub-material in the new material. If material IDs already exist, they are preserved.

- **Existing material applied (other than Multi/Sub-Object material)**

  A new Multi/Sub-Object material is created and applied to the selected faces. The existing material is moved into the Multi/Sub-Object material and becomes the first sub-material. Unselected faces get material ID #1, the selected faces get material ID #2, and the dragged material becomes part of the Multi/Sub-Object material. Existing material IDs are not preserved.

- **Multi/sub-object material applied**

  If the existing Multi/Sub-Object material is already applied more than once in the scene, the material is copied and the new copy is applied to the selected faces.

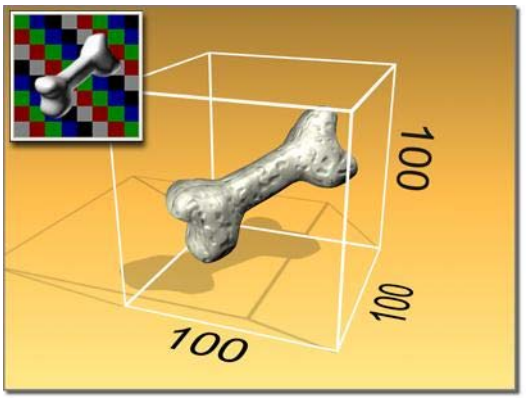
  If the Multi/Sub-Object material is only applied once in the scene, then the existing material is used. The dragged material is added to the existing Multi/Sub-Object material.

  If the dragged material already is a part of the Multi/Sub-Object material, then the selected faces receive the corresponding material ID number. If the selected faces all have the same material ID number, and no unselected faces are already using this number, then this number is used and the new material replaces the old sub-material at this ID. Otherwise, a new material ID number is assigned to the faces, and used for the dragged material. In this case, any existing material IDs are preserved.

Creating a Custom Sample Object

Material Editor > Right-click the active sample slot. > Right-click menu > Options > Material Editor Options dialog

Material Editor > Material Editor Options > Material Editor Options dialog



Custom sample object fits within a 100-unit cube

By default, the sample object in a sample slot is a sphere. You can use the *Sample Type flyout (page 2–1432)* to change this to a cube or a cylinder. You can also create a custom sample object by creating a 3ds Max scene that shows the object.

Overview of Using a Custom Sample Object

See Procedures for more details.

The scene you create should contain a single object that fits into an imaginary cube that is 100 units on each side. The object must be at the root level of the scene: it can't be linked to other objects. If more than one object is in the scene, only the first

object listed in the Track View hierarchy is used as the sample object.

If the object is of a type that doesn't have a Generate Mapping Coords check box, apply a *UVW Map modifier (page 1–922)* to it.

After you have saved the single-object scene as a *.max* file, use the Custom Sample Object group in the *Material Editor Options dialog (page 2–1436)* to specify the file. When you specify the file, a new button is displayed at the right of the Sample Type flyout. This button, which shows an object with a question mark, displays the sample object file you chose.

If the sample object scene contains only the object, sample slots display it with default lighting. If the scene also contains a camera and lights, you can use the camera to specify the view, and the lights to light the object as you choose. Turn on Load Camera and/or Lights in the Custom Sample Object group.

Procedures

To set up a custom sample object:

1. Create a scene with a single object, such as a pyramid.

   The object should fit into an imaginary cube that is 100 units on each side. Also, the object must be at the root level of the scene: it can't be linked to other objects.

   If more than one object is in the scene, the Material Editor uses the first object listed in the Track View hierarchy.

2. If the object does not have built-in mapping coordinates (via a Generate Mapping Coords check box), then assign a UVW Map modifier to it to provide mapping coordinates.

   If the object has a Generate Mapping Coords check box, it's on by default, and the Material Editor uses those coordinates. If you want to use coordinates other than those built into the object, assign a *UVW Map modifier (page 1–922)* and set up your own coordinates.

3. Save the *.max* file.

4. In the Material Editor Options dialog, click the File Name button in the Custom Sample Object group box, and choose the *.max* file that contains your object.

5. Activate the sample slot in which you want to see the custom object, then choose the button at the far right of the Sample Type flyout.

   Your custom object is displayed in the sample slot.

   If the size of your object is not quite right for the sample slot, adjust its size and save the scene again. To update the sample slot so it uses the newly saved *.max* file, open the Material Editor Options dialog, and then click OK.

To use a camera and lights with the custom sample object:

1. Create a camera in the scene that contains your sample object, and then adjust the camera to show the object as you want it seen in sample slots.

   Tip: Viewports have a different aspect ratio than sample slots, so using *Zoom Extents (page 3–740)* on the sample object usually results in the object appearing *smaller* in a sample slot. Do Zoom Extents on the object, and then before you save the file, zoom in a little farther so the object more than fills the viewport. The sample slot projection is based on the width of the sample object's geometry, not on the image in the viewport.

   If more than one camera is in the scene, the Material Editor uses the first camera listed in the Track View hierarchy.

2. If you want to use your own lighting rather than the default sample-slot lighting, set up as many lights as you need. If you want to use the

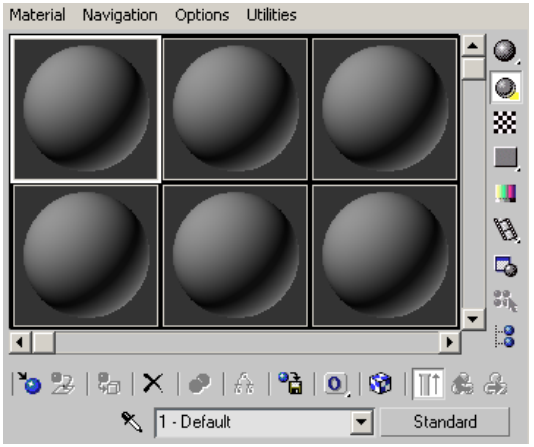
sample-slot lighting, do not add any lights to the scene.

3. Save the *.max* file.
4. In the Material Editor Options dialog, specify the file as the Custom Sample Object file.
5. Turn on Load Camera and/or Lights.

   Sample Slots set to use the custom object now display the object as seen through the camera. If lights are in the scene, those lights are used in the sample slot instead of the default lights.

Material Editor Tools

Main toolbar > Material Editor



Material Editor tools below and to the right of the sample slots

Above the *sample slots (page 2–1420)* is the menu bar. Below and to the right of the Material Editor sample slots are buttons and other controls that you use to manage and change maps and materials.

Material Editor Menu Bar (page 2–1428)

Reflectance and Transmittance Display (page 2–1430)

Note: These fields are not displayed unless you change a toggle in Preferences > Advanced Lighting.

Buttons below the sample slots (the "toolbar")

 *Get Material (page 2–1439)*

 *Put Material to Scene (page 2–1440)*

 *Assign Material to Selection (page 2–1441)*

 *Reset Map/Mtl to Default Settings (page 2–1442)*

 *Make Material Copy (page 2–1442)*

 *Make Unique (page 2–1442)*

 *Put to Library (page 2–1443)*

 *Material Effects Channel (page 2–1443)*

 *Show Map in Viewport (page 2–1445)*

 *Show End Result (page 2–1446)*

 *Go to Parent (page 2–1446)*

 *Go Forward to Sibling (page 2–1447)*

Buttons to the right of the sample slots

 *Sample Type (page 2–1432)*

 *Backlight (page 2–1432)*

 *Pattern Background (page 2–1433)*

 *Sample UV Tiling (page 2–1433)*

 *Video Color Check (page 2–1434)*

 *Make Preview, Play Preview, Save Preview (page 2–1434)*

 *Material Editor Options (page 2–1436)*

 *Select By Material (page 2–1439)*

 *Material/Map Navigator (page 2–1447)*

Controls below the toolbar

 *Pick Material from Object (page 2–1448)*

Material name field (page 2–1448)

Material Type button (page 2–1449)

Procedures

To use the Material Editor Options dialog:

1.  Click *Options (page 2–1436)* to the right of the sample slots.
2. Set the options as you want, and then click OK.

To change the preview shape:

1. Activate the sample slot of the material you want to view.
2. Use the *Sample Type flyout (page 2–1432)* to choose the shape you want to view.

   The new shape is displayed in the sample slot, with the material mapped to it.

   The flyout gives you three options: sphere (the default), cylinder, or box. An additional custom object option is available if you define a custom object as described in *Creating a Custom Sample Object (page 2–1425)*.

Material Editor Menu Bar

Material Editor > Menu bar

The Material Editor menu bar appears at the top of the Material Editor window. It provides another way to invoke the various Material Editor tools.

Material Menu

The Material menu provides the most commonly used Material Editor tools.

- *Get Material (page 2–1439)*
- *Pick from Object (page 2–1448)*
- *Assign to Selection (page 2–1441)*
- *Put to Scene (page 2–1440)*
- *Put to Library (page 2–1443)*
- **Change Material/Map Type—**Equivalent to clicking the *Material Type Button (page 2–1449)*.
- *Make Material Copy (page 2–1442)*
- **Launch Magnify Window—**Equivalent to double-clicking the active *sample slot (page 2–1420)*, or choosing Magnify on its *right-click menu (page 2–1422)*.
- *Make Preview (page 2–1434)*
- *View Preview (page 2–1434)*
- *Save Preview (page 2–1434)*
- *Show End Result (page 2–1446)*
- *Show Map in Viewport (page 2–1445)*
- **Activate All Maps—**Turns on *Show Map in Viewport (page 2–1445)* for every material used in the scene.
- **Deactivate All Maps—**Turns off *Show Map in Viewport (page 2–1445)* for every material used in the scene.
- **Reset Sample Slot Rotation—**Returns the active sample slot's object to its default orientation;

equivalent to choosing Reset Rotation on the *sample slot right-click menu (page 2–1422)*.

- **Update Active Material—**If *Material Editor Options dialog (page 2–1436)* > Update Active Only is on, choosing this updates the active material in its sample slot.

Navigation Menu

The Navigation menu provides tools that navigate a material's hierarchy.

- *Go to Parent (page 2–1446)*
- *Go Forward to Sibling (page 2–1447)*
- **Go Backward to Sibling—**Like Go Forward To Sibling, but navigates to the preceding sibling map in the tree, instead of the succeeding one.

Options Menu

The Options menu provides some additional tools and display choices.

- *Propagate Materials to Instances (page 2–1432)*
- **Manual Update Toggle—**Equivalent to the Manual Update toggle in the *Material Editor Options dialog (page 2–1436)*.
- **Copy/Rotate Drag Mode Toggle—**Equivalent to choosing either Drag/Copy or Drag/Rotate on the *sample slot right-click menu (page 2–1422)*.
- *Background (page 2–1433)*
- **Custom Background Toggle—**If you have used the *Material Editor Options dialog (page 2–1436)* to assign a custom background, this toggles its display.
- *Backlight (page 2–1432)*
- **Cycle 3X2, 5X3, 6X4 Sample Slots—**Cycles through the equivalent choices on the *sample slot right-click menu (page 2–1422)*.
- **Options—**Opens the *Material Editor Options dialog (page 2–1436)*.

Utilities Menu

The Utilities menu provides map rendering and selecting objects by material.

- **Render Map—**Equivalent to choosing Render Map on the *sample slot right-click menu (page 2–1422)*.
- *Select Objects by Material (page 2–1439)*
- *Clean MultiMaterial (page 2–1742)*
- *Instance Duplicate Map (page 2–1744)*
- **Reset Material Editor Slots—**Replaces all materials in the Material Editor with the default material type. This action is not undoable, but you can restore the previous state of the Material Editor with the Restore Material Editor Slots command (see below).
- **Condense Material Editor Slots—**Sets all unused materials in the Material Editor to the default type, retaining only materials in the scene and moving those materials to the first slots in the editor. This action is not undoable, but you can restore the previous state of the Material Editor with the Restore Material Editor Slots command (see below).
- **Restore Material Editor Slots—**When you use either of the two previous commands, 3ds Max saves the current state of the Material Editor in a buffer; using this command restores the state of the editor using the buffer contents.

  Tip: The buffer that holds the material definitions survives the File menu > Reset command. Thus, if you use either the Reset Material Editor Slots or the Condense Material Editor Slots function, and then reset 3ds Max, you can then use Restore Material Editor Slots to bring all materials back into the Material Editor. This makes it easy to use the same materials in different projects.

  For best results, follow this procedure:

  1. Save your scene.

2. Use the Reset Material Editor Slots or Condense Material Editor Slots function.
3. Reset 3ds Max. When prompted to save the scene, click No. Otherwise, you might lose material definitions in the Material Editor that were saved with the scene.
4. Open the Material Editor and choose Utilities menu > Restore Material Editor Slots.

   The Material Editor status before step 2 is restored.

Reflectance and Transmittance Display

Material Editor > Reflectance and Transmittance fields (below the sample slots)

These fields show the reflectance and transmittance of the active material. Both the average value and the maximum value are shown.

Note: These fields are not displayed unless you turn on Display Reflectance & Transmittance Information in the *Advanced Lighting panel (page 3–836)* of the Preferences dialog.

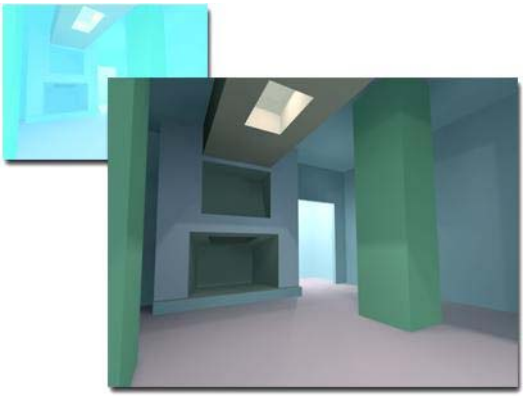


These fields can change according to the current *Show End Result setting (page 2–1446)*. When Show End Result is off, the reflectance and transmittance show the properties of the current element in the material hierarchy, not the properties of the material overall.

The reflectance and transmittance values of a material are particularly important when you are preparing a *radiosity solution (page 3–51)* to obtain a physically accurate lighting simulation. If a material's reflectance or transmittance values are not what your scene needs, you need to adjust these values.

Note:

Materials and Radiosity



Upper left: Washed-out walls due to high reflectance.

Right: Better radiosity achieved by reducing the HSV Value (V) of materials in the scene.

To get good radiosity results, keep these points in mind as you design materials:

- Before you generate a radiosity solution, check the reflectance value of all materials in the scene, to make sure it is not too high. The reflectance of a material determines how much of the light energy it receives is subsequently used in the radiosity calculation. Keep this value within the range of the physical materials you are simulating. (See the table in the description of "Reflectance," below.)
- Don't be concerned if a material preview seems too dark. For example, a white wall with the maximum recommended reflectance of 80% appears gray. The color balance will be adjusted correctly by the *exposure control (page 3–293)* in the final rendering.
- Bitmaps used as diffuse textures have already been illuminated by the scanner, digital camera, or paint program in which you created them. To bring them into the proper reflectance range, you might have to dim them by reducing the

RGB Level value in the bitmap's *Output rollout (page 2–1621)*.



Left: Wood grain bitmap as originally photographed has too high a reflectance.

Right: Reducing the RGB Level value reduces the map's reflectance.

Interface

Reflectance—Reflectance is the percentage of diffuse light energy that is reflected from a material. When you increase a material's HSV value (V), the material reflects more diffuse light. Decreasing a material's Opacity also decreases its reflectance.

Typically, the reflectance of a material should never be greater than 85%. this is an unusually high value that will lead to poor-quality renderings. In the real world, even the whitest wall reflects no more than 80% of the light it receives.

One source of high reflectance can be a map assigned to the material's diffuse component. For example, a white tile bitmap might create high reflectance. In this case, you can reduce reflectance by reducing the RGB Level in the bitmap's Output rollout.

An alternate way to reduce a bitmapped material's reflectance is to set the diffuse color of the material to black, and then reduce the diffuse map's Amount (in the parent material's *Maps rollout (page 2–1474)*). You can use this method to reduce the reflectance of *3D procedural maps (page 2–1662)* as well.

Here are some typical reflectance ranges for common materials:

Material	Minimum	Maximum
Ceramic	20%	70%
Fabric	20%	70%
Masonry	20%	50%
Metal	30%	90%
Paint	30%	80%
Paper	30%	70%
Plastic	20%	80%
Stone	20%	70%
Wood	20%	50%

Transmittance—Transmittance is the amount of light energy transmitted through a material. A completely opaque material has 0% transmittance.

When the material is transparent (like clear glass), the energy transmitted is specular, and light passes directly through the material (subject to refraction). The value of specular transmittance is an interaction between the material's *Opacity value (page 2–1488)* and its *filter color (page 3–939)*. When the filter color is pure white, specular transmittance is the inverse of Opacity.

When the material is translucent (like frosted glass), the energy transmitted is diffuse, and scattered in all directions. The value of diffuse transmittance is based on the HSV value (V) of the *translucent color (page 2–1491)*. Diffuse transmittance happens only for materials that use the *Translucent shader (page 2–1484)*, or *Raytrace materials (page 2–1512)*, that specify translucency.

Usually, when a material has high transmittance, it has low reflectance, and the other way around.

Sample Type

Material Editor > Sample Type



The Sample Type flyout lets you choose which geometry to display in the active *sample slot (page 2–1420)*. This flyout has three buttons:

- **Sphere (the default)—**Displays the material on a sphere.
- **Cylinder—**Displays the material on a cylinder.
- **Cube—**Displays the material on a cube.
- **Custom—**Displays the material on a custom object. This button appears only if you have used the *Material Editor Options dialog (page 2–1436)* to specify a custom object for sample slots.

See also

Creating a Custom Sample Object (page 2–1425)

Propagate Materials to Instances

Material Editor > Options menu > Propagate Materials to Instances

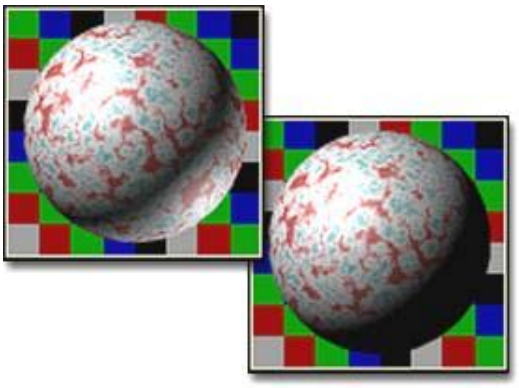
When Propagate Materials To Instances is on, any material assignment you make will be propagated to all *instances (page 3–957)* of the object in your scene, including imported AutoCAD blocks and ADT style-based objects; these object types are common in *DRF files (page 3–527)*. Assignments are also propagated to instances of Revit objects and of other instances that you've made in the current scene.

When Propagate Materials To Instances is off, materials are assigned in traditional 3ds Max fashion; each object has a unique material assignment.

Backlight

Material Editor > Backlight

Material Editor menu > Options menu > Backlight



Left: Backlight on
Right: Backlight off

Turning on Backlight adds a backlight to the active sample slot. This button is on by default.

The effect is most easily seen with the sample spheres, where the backlight highlights the lower-right edge of the sphere.

Backlight is especially useful whenever you're creating *metal (page 2–1481)* and *Strauss (page 2–1483)* materials. Backlight lets you see and adjust the specular highlight created by glancing light, which is much brighter on metals.

Sample Slot Background

Material Editor > Background

Material Editor menu > Options menu > Background



Left: Background on
Right: Background off

Turning on Background adds a multicolored checkered background to the active sample slot. The pattern background is helpful when you want to see effects of opacity and transparency.

The *Material Editor Options dialog (page 2–1436)* also lets you assign a bitmap to use as a custom background.

Sample UV Tiling

Material Editor > Sample UV Tiling



1x1, 2x2, and 3x3 sample tiling



The buttons on the Sample UV Tiling flyout adjust the repetition of the *map (page 3–968)* pattern on the sample object in the active sample slot.

The tiling pattern you set with this option affects *only* the sample slot. It has no effect on the tiling on the geometry in the scene, which you control with parameters in the map's own coordinates rollout.

- **1 x 1 (the default)—**Tile once in the U dimension and once in V. This is equivalent to no tiling at all.
- **2 x 2—**Tile twice in U and twice in V.
- **3 x 3—**Tile three times in U and three times in V.
- **4 x 4—**Tile four times in U and four times in V.

The buttons specify the number of times the pattern repeats over the surface of the sample object. Because maps are mapped spherically around the sample sphere, the *tiling (page 3–1022)* repetition covers the entire surface of the sphere. The sample cylinder maps cylindrically. The sample cube uses box mapping: the tiling appears on each side of the cube. Custom sample objects use the default mapping coordinates for that kind of object, unless the object has a *UVW Map modifier (page 1–922)* applied to it, in which case the modifier controls the mapping.

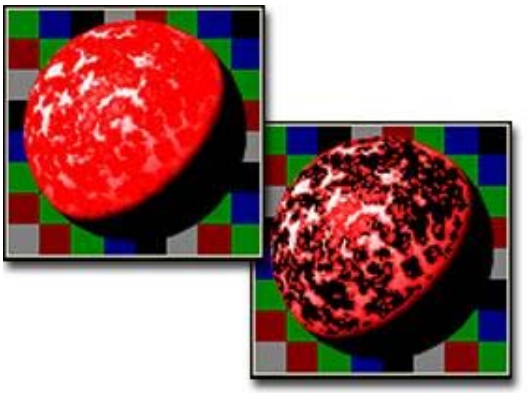
This flyout is unavailable when the sample slot displays a standalone (top-level) map.

See also

Creating a Custom Sample Object (page 2–1425)

Video Color Check

Material Editor > Video Color Check



Left: Material with oversaturated colors

Right: Video color check shows black areas that are beyond the video threshold

Video Color Check checks the material on the sample object for colors that are beyond the safe *NTSC (page 3–980)* or *PAL (page 3–988)* threshold.

Such colors tend to blur or fuzz when transferred from computer to video. Pixels containing these "illegal" or "hot" colors are flagged on the sample object.

You can have 3ds Max correct illegal colors automatically when you render, depending on the settings in the *Rendering tab (page 3–826)* of the *Customize/Preferences dialog (page 3–815)*.

Use this option as a guide only. The colors in a rendered scene depend not only on the material color but also on the intensity and color of the lighting. A material that shows as safe in the sample slot might become illegal if rendered under several bright lights. A safe practice for video is to use colors whose saturation is less than 80 to 85 percent.

Procedures

To check for illegal video colors:

1. On the Material Editor, turn on *Video Color Check (page 2–1434)*.

   The active sample slot now renders "illegal" pixels as black. Illegal pixels have a color that is beyond the safe video threshold.

   If Video Color Check detects illegal colors, try reducing the saturation of the material colors in question.

2. You can also check for legal video colors when you render a scene.

To change the video system:

NTSC (page 3–980) is the default video system to check. You can change the video checking to *PAL (page 3–988)*.

- In the Tools menu > Options > Options dialog > *Rendering panel (page 3–826)* > Video Color Check group, change the setting to PAL.

  The sample slot does not automatically update when you change the video system preference. Turn video checking off and back on to see the change.

  Note: The Rendering panel also shows options for choosing alternate ways to display illegal pixels. These apply to renderings only, not to the sample slots in the Material Editor.

Make Preview, Play Preview, Save Preview

Material Editor > Make/Play/Save Preview flyout

Material Editor menu > Material menu > Make Preview, View Preview, or Save Preview

You can use animated maps to add movement to a scene. For example, to simulate a sky view,

you could map an animation of moving clouds to a clerestory window. The Make Preview options let you experiment with the effect in the Material Editor before you apply it to your scene.

The buttons on the Make Preview, Play Preview, Save Preview flyout let you preview the effect of an animated map on the object in a sample slot (see *Sample Type (page 2–1432)* for alternatives to the default sphere.) You can use an *AVI file (page 3–609)*, or *IFL file (page 3–616)* as a source for the animation. The finished preview is saved as a new AVI file and automatically played back. You can also view the preview in a sample slot by dragging the time slider.

This flyout has three options:

-  **Make Preview—**Displays a *Create Material Preview dialog (page 2–1452)* to create an AVI file of the animated material.

  When the preview is complete, it is saved as *_medit.avi* in the *\previews* subdirectory. The file is then opened in the Windows Media Player and played back.

  Without this option, the only other way to preview a complex animated material in real time is to use the *renderer (page 3–27)* to render an animated sequence and save it to an AVI file, a digital disk recorder, or a video tape recorder.

-  **Play Preview—**Uses the Windows Media Player to play the current *_medit.avi* preview file in the *\previews* subdirectory.

-  **Save Preview—**Saves the *_medit.avi* preview to an AVI file of a different name in the *\previews* subdirectory.

Procedures

To create a preview of an animated material:

1. Choose Make Preview from the flyout.
2. In the Create Material Preview dialog, set the preview conditions you want, and then click OK.

   The preview is created, and then plays.

To play back a preview of an animated material:

- Choose Play Preview from the flyout.

  The Material Editor starts the Windows Media Player. The Player then loads the current *_medit.avi* preview file in the *\previews* subdirectory and plays it in a separate window.

To view an animated material in a sample slot:

- Drag the time slider through the range of frames.

  By default, all sample slots with animated maps update at the same time. To change this, see *Material Editor Options Dialog (page 2–1436)*.

To save a preview under a different name:

1. Choose Save Preview from the flyout.

   The Material Editor displays a File Save dialog.
2. Enter a new name for the preview, and then click OK to save the file in the *\previews* subdirectory.

To play a renamed preview:

1. Choose Tools menu > Display Image.
2. In the View File dialog, change to the *\previews* subdirectory.
3. Select the renamed preview file and click Open.

   The preview is played in a separate window.

Material Editor Options

Material Editor > Material Editor Options

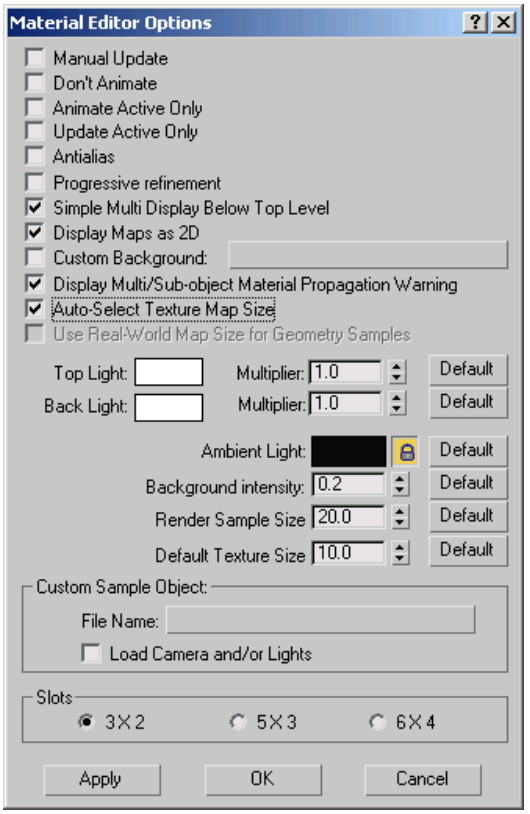
Material Editor menu > Options menu > Options

This button displays the Material Editor Options dialog to let help you control how materials and maps are displayed in the *sample slots (page 2–1420)*.

These settings are "sticky"; they survive a reset, and even quitting and restarting 3ds Max.

Note: The controls to assign a renderer for the sample slots are on the *Assign Renderer rollout (page 3–35)*.

Interface



Manual Update—When on, the sample slots don't update their contents until you click them. This option affects only the updating of the sample slots; it doesn't affect the icon displays in the Browser. Default=off.

Don't Animate—When on, animated maps are not updated in the sample slots while you play an animation or drag the time slider. However, the animation is updated to the current frame when you stop the animation or release the time slider. An animated map can use an *AVI file (page 3–609)* or *IFL file (page 3–616)*as a source. Default=off.

Animate Active Only—When on, only the active sample slot is animated when you play an animation or drag the time slider. This option is good for situations where you have multiple animated materials in the Material Editor, but you only need to see one at a time. This check box is unavailable when Don't Animate is on. Default=off.

Update Active Only—When on, sample slots do not load or generate maps until you make one sample slot active. This can save time while you use the Material Editor, especially when your scene uses a lot of materials with maps. Default=off.

Antialias—Turns on *antialiasing (page 3–907)*in the sample slots. Default=off.

If the mental ray renderer is used to render sample slots, this setting is disregarded. The mental ray renderer's sampling values control antialiasing.

Progressive Refinement—Turns on progressive refinement in the sample slots. When on, samples are rendered quickly, with large pixels, then rendered a second time in greater detail. Default=off.

Simple Multi Display Below Top Level—When on, the sample sphere for a Multi/Sub-Object material displays the multiple patches only at the top level of the material. The sub-materials are displayed over the entire sphere. When you use nested

Multi/Sub-Object materials, the multiple patches again appear at the top level of the nested material, but the sample sphere is again whole when displaying any of the sub-materials. Default=on.

Display Maps as 2D—When on, sample slots display maps, including standalone maps, in 2D. The map fills the entire slot. When off, maps are displayed on the sample object, as materials are. Default=on.

Custom Background—Lets you specify a custom background for the sample slots, instead of the default checkers background. Click the file-assignment button to display a file dialog from which you can select the custom background. This can be any bitmap format supported by 3ds Max. Turn on Custom Background to use the new background instead of the checkered background. The custom background is stored in the *3dsmax.ini (page 1–18)* file, so it is available from session to session. Default=off.

Display Multi/Sub-Object Material Propagation Warning—Toggles display of warning dialog when you apply a *Multi/Sub-Object material (page 2–1594)* to an instanced ADT style-based object.

Auto-Select Texture Map Size—When on, and you have a material that uses a texture map set to Use Real-World Scale, ensures that the map will be displayed correctly on the sample sphere. Turn off to be able to enable Use Real-World Map Size For Geometry Samples (see following).

Note: If a material uses several texture maps at different levels, and only one is set to Use Real-World Scale, the sample sphere will render with real-world size coordinates.

Use Real-World Map Size For Geometry Samples—This is a global setting that allows you to manually choose which style of texture coordinates are used. When on, real-world coordinates are used for the sample slot display. Otherwise, the old style of 3ds Max mapping coordinates is active. When off, you must turn on Use Real-World Scale on the map's Coordinates rollout to see the sample sphere as you'd expect. Available only when Auto-Select Texture Map Size (above) is off. Default=off.

Top Light color/Back Light color—Specify the two lights used in the sample slots. Click the color swatch to alter the color of either light. Adjust the Multiplier spinners to multiply the values (intensity) of the lights.

Use the Default buttons to return to the initial settings.

Ambient Light—Shows the color of *ambient light (page 3–908)* used in the sample slots. Click the color swatch to change the color. When the lock button is on, changing the Ambient Light color here or on the *Environment panel (page 3–272)* changes both; when off, changing one setting does not affect the other.

Use the Default button to return to the initial setting.

Background Intensity—Sets the background intensity in the sample slots. The range is from 0 (black) to 1 (white). Default=0.2.

Use the Default button to return to the initial setting.

Render Sample Size—Sets the scale of the sample sphere to any size, making it consistent with the object or objects in the scene that have the texture on them. This setting affects how 2D and 3D maps are displayed providing that the sample spheres are set to display real-world scale.

Note: The size is scaled to use the current units.

This is a global option that affects all the sample slots. Default=100.0 (Imperial units) and 2.54m (Metric units).

Use the Default button to return to the initial setting.

Default Texture Size—Controls the initial size (both height and width) of a newly created real-world texture. You see the result of changing this option only when you create a new texture in a material; the change appears in the *Coordinates rollout (page 2–1625)*. Default=48.0 (Imperial units) and 1.219m (Metric units).

Note: This setting applies to real-world textures only. For the default size to be applied to newly created textures, the Preferences dialog > General panel > Use Real-World Texture Coordinates check box must be on.

Use the Default button to return to the initial setting.

DirectX Shader group

These options affect the viewport behavior of the *DirectX 9 Shader material (page 2–1613)*.

Force Software Rendering—When on, forces DirectX 9 Shader materials to use the selected software render style for viewports. When off, the FX file specified in the DirectX 9 Shader is used unless the material's local Force Software Rendering toggle is on. Default=off.

Shade Selected—When Force Software Rendering is on, selected objects, and only selected objects, are shaded by the DirectX 9 Shader material. This toggle is unavailable unless Force Software Rendering is on. Default=off.

Custom Sample Object group

Controls in this group let you specify a *custom sample object (page 2–1425)* to use in the *sample slots (page 2–1420)*.

File Name—Selects the MAX scene file.

The scene should contain a single unlinked object that fits in an imaginary cube 100 units on a side. The object must be either a primitive with a Generate Mapping Coords. check box, or have a *UVW Map modifier (page 1–922)* applied to it. The scene can contain a camera and lights.

Load Camera and/or Lights—Turn on to have sample slots use the camera and lights in the scene, instead of the default sample slot lighting.

Slots group

These options let you choose how many sample slots to display at a time.

The Material Editor always has 24 sample slots available. You can choose to display fewer sample slots at a larger size. When you do, scroll bars let you move around among the sample slots.

3 X 2—Specifies a 3 x 2 array of sample slots. (The default: 6 windows.)

5 X 3—Specifies a 5 X 3 array of sample slots. (15 windows.)

6 X 4—Specifies a 6 X 4 array of sample slots. (24 windows.)

Apply—Applies the current settings, except for changes to the Slots group, without leaving the Material Editor Options dialog. This is useful when you adjust lighting values for the sample slots.

OK—Closes the dialog and applies any changes you made.

Cancel—Closes the dialog and cancels any changes you made, including changes you applied with the Apply button.

Select By Material 

Material Editor > Select By Material

Material Editor menu > Utilities menu > Select Objects by Material

Select By Material allows you to select objects in the scene, based on the active material in the Material Editor. This button is unavailable unless the active sample slot contains a material used in the scene.

Clicking this button displays a *Select Objects dialog (page 1–78)*. All objects that have the selected material applied to them are highlighted in the list.

Note: Hidden objects don't appear in this list, even if the material is applied to them. However, in the *Material/Map Browser (page 2–1412)*, you can choose Browse From: Scene, turn on By Object, and then browse from the scene. This lists all objects in the scene, hidden or unhidden, along with their assigned materials.

Procedure

To select objects that have the same material applied:

1. Click a sample slot that contains a material in the scene.

   White corner brackets indicate materials that are in the scene.

2.  Click *Select By Material (page 2–1439)* in the Material Editor.

   This button is unavailable unless the active sample slot contains a material in the scene.

   A *Select Objects dialog (page 1–78)* is displayed. The names of objects with the active material applied are highlighted.

3. Click Select to select objects with the active material applied.

   You can also change the selection by choosing other objects. If you change the selection, you must then click *Assign Material To Selection (page 2–1441)* to apply the active material to newly selected objects.

Get Material 

Material Editor > Get Material

Material Editor menu > Material menu > Get Material

Get Material displays the *Material/Map Browser (page 2–1412)* to allow you to choose a *material (page 3–971)* or *map (page 3–968)*.

Procedures

To get a material from a scene:

1. Click a sample slot to make it active.

   Be careful not to click the sample slot of a material you want to use later.

2.  On the Material Editor toolbar, click *Get Material (page 2–1439)*.

   A modeless *Material/Map Browser (page 2–1412)* is displayed.

3. In the Browse From group box at the upper left, make sure that either Selected or Scene is chosen.

   The Selected option lists only materials in the current selection. If no objects are selected, the list of materials is blank.

   The Scene option lists all the materials currently in the scene.

4. In the list of materials, double-click the name of the material you want to get.

   You can also drag the material name to the sample slot.

The material you chose replaces the previous material in the active sample slot.

Warning: When you get a material from a scene, initially it is a *hot (page 3–953)* material.

To get a material from a library:

1.  On the Material Editor toolbar, click *Get Material (page 2–1439)*.

   A modeless *Material/Map Browser (page 2–1412)* is displayed.

2. In the Browse From group box at the upper left, make sure that Material Library is chosen.

   If you have opened a library, the list of materials shows the contents of the library.

   If you haven't opened a library, click Open in the file area of the Browser. A file dialog is displayed. Choose a library. After you open the library, the list of materials updates to show the library contents.

   Note: Open also lets you get materials from a 3ds Max scene (a *.max* file).

3. In the list of materials, double-click the name of the material you want to get.

   You can also drag the name of the material to the sample slot.

   The material you chose replaces the previous material in the active sample slot.

To create a standalone map tree:

1. Activate a sample slot.
2.  On the Material Editor toolbar, click Get Material.
3. In the *Material/Map Browser (page 2–1412)*, make sure Browse From is set to New.
4. In the Show group box, turn off Materials so only maps are displayed in the list.
5. Double-click the name of the map type (not a material type) you want to use, or drag the map to a sample slot.

   The sample slot now contains a standalone map not associated with material parameters.
6. Use the Material Editor to modify the map as you would any other map.

   By default, the sample slots distinguish maps from materials by displaying maps as 2D surfaces without lighting or shading.

To remove a material from an object:

1.  On the Material Editor toolbar, click Get Material.

   The Material/Map Browser appears.
2. Drag the entry NONE from the top of the list in the Browser to the object.

   The object now has no material applied to it.

Put Material to Scene 

Material Editor > Put Material to Scene

Material Editor menu > Material menu > Put to Scene

Put Material To Scene updates a material in the scene after you edit the material. Put Material To Scene is available only when:

- The material in the active sample slot has the same name as a material in a scene.
- The material in the active sample slot is not hot.

In other words, this command is intended to fit into the overall sequence of handling materials:

- You create a hot material either by applying it to objects in the scene or by getting it from the scene.
- You make a copy of the material.
- You make changes to the copy of the material.

- You update the scene by putting the changed material back into the scene.

Notes

- If you apply a mapped material to a parametric object whose Generate Mapping Coords option is off, the software automatically turns on mapping coordinates at render time. In addition, if you apply a mapped material with Show Map in Viewport active to an object, that object's Generate Mapping Coords option is turned on if necessary.
- The Show Map In Viewport flag is now saved with individual materials, so you can drag mapped materials from the modeless Browser onto objects in your scene, and the mapping appears in the viewports.

Procedure

To put a material back into a scene:

-  On the Material Editor toolbar, click *Put Material To Scene (page 2–1440)*.

  The material in the active sample slot is now a *hot material (page 3–953)*.

Assign Material to Selection 

Material Editor > Assign Material to Selection

Material Editor menu > Material menu > Assign to Selection

Assign Material to Selection applies the material in the active sample slot to the currently selected object or objects in the scene. At the same time, the sample slot becomes *hot (page 3–953)*.

If you apply a mapped material to a parametric object whose Generate Mapping Coords option is off, the software automatically turns on mapping coordinates at render time. In addition, if you apply a mapped material with *Show Map In Viewport (page 2–1445)* active to a parametric object, that object's Generate Mapping Coords option is turned on if necessary.

The Show Map In Viewport flag is saved with individual materials, so when you drag mapped materials from the modeless *Browser (page 2–1412)* onto objects in your scene, the mapping appears in the viewports.

The Undo command works for material assignment.

Procedure

To apply a material to objects in a scene:

1. Select the sample slot that contains the material you want to apply.
2. Select the objects you want to apply the material to.
3. Drag from the sample slot to the objects. If more than one object is selected, you are asked whether you want to apply to the single object or to the whole selection.

    You can also apply materials by clicking *Assign Material To Selection (page 2–1441)* on the Material Editor toolbar.

Warning: When you apply a material to an object or selection, that material becomes a *hot material (page 3–953)*. When you change the material's properties, the scene immediately updates to reflect those changes. *Any* object with that material will change its appearance, not just the objects in the current selection. When a material is hot, its sample slot is displayed with white corner brackets.

 To make a material no longer hot so it doesn't change the current scene, click *Make Material Copy (page 2–1442)*.

Reset Map/Mtl to Default Settings 

Material Editor > Reset Map/Mtl to Default Settings

Reset Map/Mtl to Default Settings resets the values for the map or material in the active sample slot.

The material colors are removed and set to shades of gray. Glossiness, opacity, and so on are reset to their default values. Maps assigned to the material are removed.

If you are at a map level, this button resets the map to default values.

Reset changes the *name (page 2–1448) only* when this field names a material used in the scene.

Make Material Copy 

Material Editor > Make Material Copy

Material Editor menu > Material menu > Make Material Copy

Make Material Copy *"cools" (page 3–924)* the current *hot (page 3–953)* sample slot by copying the material to itself.

The sample slot is no longer hot, but the material retains its properties and name. You can adjust the material without affecting it in the scene. Once you've got what you want, you can click *Put Material to Scene (page 2–1440)* to update the material in the scene and change the sample slot to hot again.

Make Unique (Material Editor) 

Material Editor > Make Unique

Make Unique makes a map *instance (page 3–957)* into a unique copy. It also makes an instanced sub-material into a unique, standalone sub-material. It gives the sub-material a new material name. A sub-material is a material within a *Multi/Sub-Object material (page 2–1594)*.

Using Make Unique prevents changes to the top-level material instance from affecting the sub-material instance within the Multi/Sub-Object material.

You can also use Make Unique at the map level, when a map is instanced to different components of the same material.

Note: If you drag an instanced map to a Material Editor sample slot, the Make Unique button will not be available, because it is not clear from the context what it would be unique relative to. Instead, you need to bring one of the parent maps or materials into the Material Editor, browse down into the map, and then make the map unique relative to that parent.

Procedure

Example: Create an instanced sub-material:

1.  Create a box and a sphere.
2.  Open the Material Editor.
3. Choose a sample slot, click the Material Type button, choose Multi/Sub-Object in the Material/Map Browser, and then click OK.
4.  Select the box, and then apply the new Multi/Sub-Object material to it.

5. Drag one of the Sub-Material buttons from the Material Editor to the sphere.
6.  Choose a different sample slot, and use the Pick Material From Object button to get the sub-material applied to the sphere.

   At this point, the material applied to the sphere and the sub-material are instances of each other.
7. Go to the parameters for the instanced sub-material by clicking its Sub-Material button.
8.  The Make Unique button is now available. Click it to make the sub-material unique again, and assign it a new material name.

   Make Unique is not available for the top-level instance of the sub-material.

Put to Library 

Material Editor > Put to Library

Material Editor menu > Material menu > Put to Library

Put to Library adds the selected material to the current library.

A *Put to Library dialog (page 2–1455)* is displayed, which lets you enter a name for the material that's different from the one you used in the Material Editor.

The material becomes visible in the material library display in the *Material/Map Browser (page 2–1412)*. The material is saved to the library file on disk. (You can also save a library by using the Save button in the Material/Map Browser.)

Procedure

To save a material in a library:

1. Click to select the sample slot that has the material you want to save.
2.  On the Material Editor toolbar, click *Put To Library (page 2–1443)*.
3. A *Put To Library dialog (page 2–1455)* is displayed.
4. Either change the material name or leave it as is, and then click OK.

   The material is saved in the currently open library. If no library is open, a new library is created. You can save the new library as a file using the *Material/Map Browser (page 2–1412)* file controls.

Material ID Channel 

Material Editor > Material ID Channel

The buttons on the Material ID Channel flyout tag a material as a target for a *Video Post (page 3–311)* effect or a *rendering effect (page 3–218)*, or for storing with a rendered image saved in *RLA (page 3–630)* or *RPF (page 3–631)* file format (so that the channel value can be used in a post-processing application). The material ID value is the counterpart of a *G-buffer value (page 3–946)* for objects.



Zero (0), the default, indicates that no material ID channel is assigned.

A value from 1 to 15 means to apply a Video Post or rendering effect that uses this channel ID to this material.

For example, you might want a material to glow wherever it appears in the scene. The material is in the Material Editor and the glow comes from a rendering effect. First, you add a *Glow rendering effect (page 3–226)* and set it up so that it operates on ID 1. Use Material ID Channel to give the material an ID number of 1, then apply the material to objects in the scene in the usual way.

To save the channel data with the rendering, use the RLA or RPF format.

Warning: The *mental ray renderer (page 3–78)* does not recognize Z-depth with G-buffers. G-buffer data is saved on a single layer. Also, the mental ray renderer does not support the following effects:

- *Glow lens effect (page 3–226)* (rendering effect)
- *Ring lens effect (page 3–230)* (rendering effect)
- *Lens effects Focus filter (page 3–362)* (Video Post)

Procedure

To assign a material ID channel to a material:

- 

Choose a channel number from the *Material ID Channel flyout (page 2–1444)*.

  Note: Giving a material a nonzero ID channel number tells the renderer to generate a material ID channel containing that value. This information is stored in images only if you save the rendered scene in RLA or RPF format. However, the ID channel data is available to rendering effects at render time.

Material ID Channel Flyout

Material Editor > Tool buttons > Material ID Channel flyout



The *Material ID Channel (page 2–1443)* flyout in the *Material Editor (page 2–1409)* lets you assign a material ID to the material. The ID value can associate a *Video Post effect (page 3–311)* or *rendering effect (page 3–218)* to objects assigned this material. The channel ID can also be saved if you render to a file in the *RLA (page 3–630)* or *RPF (page 3–631)* format. Saving the material ID in an image file lets you use the channel in post-processing applications, such as Autodesk Combustion.

Zero (0), the default, indicates that no material ID channel is assigned.

A value from 1 to 15 says that an effect that uses this material ID will be applied to this material.

Show Map in Viewport 

Material Editor > Show Map in Viewport

Material Editor menu > Material menu > Show Map in Viewport



Left: Map shown on sample cube

Right: Map shown in a viewport

Show Map In Viewport displays mapped materials on the surfaces of objects in viewports with the *interactive renderer (page 3–1030)*.

The following conditions must be met for mapped materials to display in the viewport:

- *Mapping coordinates (page 3–967)* must be applied to the object. This is already the case with most primitive objects, which by default have Generate Mapping Coords on at creation. If an object doesn't have mapping coordinates, you can turn this on, or apply a mapped material to the object (if it has a Generate Mapping Coords check box), or apply a *UVW Map modifier (page 1–922)* or an *Unwrap UVW modifier (page 1–878)*.
- A mapped material must be applied to the object.
- Show Map In Viewport must be on, either at the level of the material that contains the *map (page 3–968)*, or at the top level of the material.

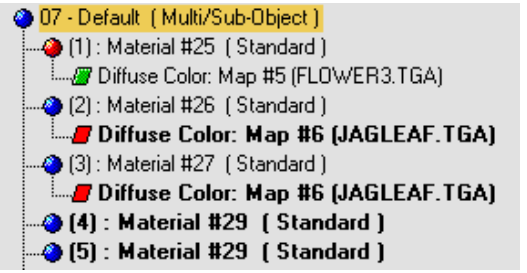
  Tip: Displaying mapped materials in a viewport can slow performance. If you don't need to view the map, turn off its viewport display.

  Note: With the *software display driver (page 3–840)*, viewports don't accurately display a map with transparency that has UV tiling or mirroring. Also, this driver cannot display a map on a faceted material.

The state of this button is saved with the material in the library. When you apply a mapped material while Show Map In Viewport is active, the object's Generate Mapping Coords. check box is turned on. This means that you can drag mapped materials from the Material Library in the Browser over objects in your scene, and have the mapped material appear in the viewports.

Tip: You can turn Show Map In Viewport on or off for *all* materials by choosing Views menu > Activate/Deactivate All Maps.

Note: In the *Material/Map Browser (page 2–1412)* and *Material/Map Navigator (page 2–1447)*, icons of materials and maps for which Show Map in Viewport is on are red, as shown in the following example:



3D Maps in Viewports

Show Map In Viewport works for 3D procedural maps as well as 2D maps.

3D map display in viewports is not necessarily accurate. To improve the 3D preview, you can use the *Material Editor Options dialog (page 2–1436)* to set the 3D Map Sample Scale to equal a main dimension of the object you are applying the map to. For example, if you want to use the planet map

on a sphere with a radius of 20, change the map scale from 100 (the default) to 20.

Particle Age and Particle MBlur maps do not preview in viewports.

Multiple Maps in Viewports

Viewports can display multiple maps. For multiple map display, the display driver must be *OpenGL (page 3–841)* or *Direct3D (page 3–844)*. The *software display driver (page 3–840)* does not support multiple map display.

The *composite map (page 2–1688)* and *mix map (page 2–1689)* support multiple map display.

In addition, turning on Show Map In Viewport at the top level of a *standard material (page 2–1465)* lets you view maps on both the diffuse and opacity components (though not on other mapped components).

Procedures

To view maps interactively:

1. Select an object.
2. In the object's creation parameters, make sure that Generate Mapping Coords. is on. (If this option isn't enabled, the object can't be mapped.)

   If the object type does not have a mapping coordinates check box, apply a *UVW Map (page 1–922)* modifier.
3. In the Material Editor, apply the mapped material to the object.
4.  Turn on Show Map In Viewport.

   The map appears on objects assigned the material in all shaded viewports. Now when you adjust a map, the viewports update to display the adjustments.

To turn off interactive texture display:

-  Turn off Show Map In Viewport.

  The object is shaded but the map no longer appears.

Show End Result 

Material Editor > Show End Result

Material Editor menu > Material menu > Show End Result

Show End Result lets you look at the material at the level you're on instead of looking at the end result of all the other *maps (page 3–968)* and settings.

When this button is off, the sample slot shows only the current level of the material. This tool is useful when you are working with *compound materials (page 3–922)*. It would be difficult to see exactly what effect you're creating on a particular level if you didn't have the ability to turn off the display of the other levels.

Go to Parent 

Material Editor > Go to Parent

Material Editor menu > Navigation menu > Go to Parent

Go To Parent moves up one level in the current material.

This button is available only when you are not at the top level of a *compound material (page 3–922)*. You can tell you're at the top level when this button is unavailable and the name in the edit field matches the name in the Material Editor title bar.

A typical situation is one in which you have a material with a Diffuse map. The Material level is the parent and the Diffuse map is the child. The Go To Parent button becomes available at the level of the Diffuse map.

Tip: You can also navigate through the levels of a material with the *Material/Map Navigator (page 2–1447)*.

Go Forward to Sibling

Material Editor > Go Forward to Sibling

Material Editor menu > Navigation menu > Go Forward to Sibling

Go Forward To Sibling moves to the next map or material at the same level in the current material.

This button is available only when you are not at the top level of a *compound material (page 3–922)*, and there is more than one map or material at the current level.

A typical situation is one in which you have a material with a Diffuse map, a Bump map, and a Glossiness map. The Material level is the parent and the Diffuse map, Bump map, and Glossiness map are its children. Go Forward To Sibling becomes available at the level of the children and allows you to go from one to another.

You can also navigate through the levels of a material with the *Material/Map Navigator (page 2–1447)*.

Material/Map Navigator

Material Editor > Material/Map Navigator

The Material/Map Navigator is a modeless dialog that provides quick navigation through the hierarchy of *maps (page 3–968)* in a *material (page 3–971)*, or of sub-materials in a *compound material (page 3–922)*.

The Navigator displays the material and maps in the currently active sample slot. You can navigate the hierarchy of the current material by clicking the material or map listed in the Navigator. Conversely, as you navigate the material in the Material Editor, the current level is highlighted in the Navigator. The selected material or map becomes active in the sample slot, while the rollouts for the selected material or map are displayed below.

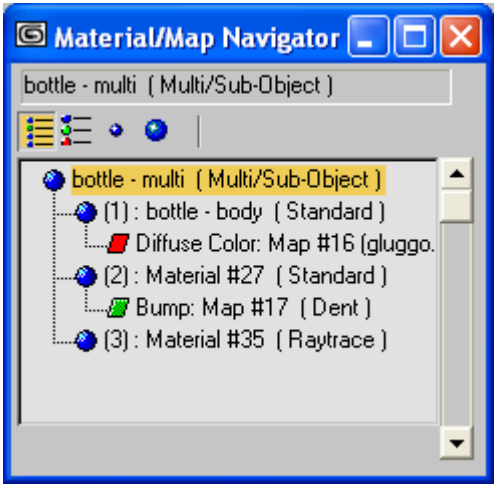
You can also drag from the Navigator to any valid sample slot or button in the user interface.

Use the view buttons at the top to change the display. The list displayed in the Navigator is similar to the one displayed in the Track View hierarchy.

See also

Material/Map Browser (page 2–1412)

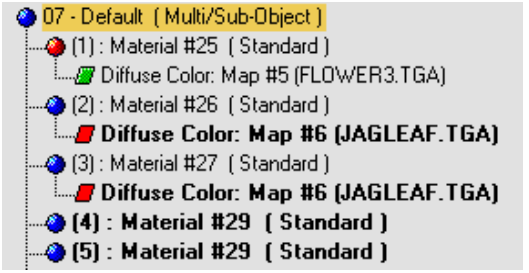
Interface



Note: Icons of materials and maps for which *Show Map in Viewport (page 2–1445)* is on are red.

Also, the names of *instanced (page 3–957)* materials and maps appear in boldface.

Both of these effects are shown in the following illustration:



The view buttons are as follows:

View List—Displays the materials and maps in list format. Blue spheres are materials. Green parallelograms are maps. The green parallelograms turn red if Show Map in Viewport is on for a material.

View List + Icons—Displays the materials and maps as small icons in a list.

View Small Icons—Displays the materials and maps as small icons. As you move the mouse over the icons, tooltip labels show you the name of the material or map.

View Large Icons—Displays the materials and maps as large icons. As you move the mouse over the icons, tooltip labels show you the name of the material or map.

The large icons are labeled with the name of the material or map and are displayed using progressive refinement. That is, samples are rendered quickly, with large pixels, then rendered a second time in greater detail.

Pick Material From Object (Eyedropper)

Material Editor > Pick Material From Object

Material Editor menu > Material menu > Pick from Object

Pick Material From Object lets you select a material from an object in the scene. Click the eyedropper button, and then move the eyedropper cursor over the objects in the scene. When the eyedropper cursor is over an object containing a material, it fills with "ink" and a tooltip with the name of the object pops up. Click the object. The material is placed in the active sample slot.

If the material is already in the active sample slot, the eyedropper has no effect.

If the eyedropper cursor is over an *editable mesh (page 1–996)* with faces selected at the sub-object level, and the mesh has a *Multi/Sub-Object material (page 2–1594)* applied to it, then the eyedropper picks up the sub-material. However, if the selected faces have more than one sub-material assigned to them, then the eyedropper picks up the entire Multi/Sub-Object material.

Name Field (Materials and Maps)

Material Editor > Material Name field



The Name field displays the name of the material or map. Default material names are "01 – Default," and so on, the number changing to reflect the material's sample slot. Maps are named "Map #1," and so on.

You can edit this field to change the name of the material in the active sample slot. You can also edit

the names of maps and sub-materials assigned at lower levels of the map or material hierarchy.

The name of the material is not a file name: it can contain spaces, numbers, and special characters. It can be of any length.

Note: In releases prior to 3ds Max **4**, material names were limited to 16 characters. As of 3ds Max **4**, you can assign material names of arbitrary length.

This field also functions as a drop-down list. At the top level, it shows only the material or map name. At lower levels of the hierarchy, drop the list down to see the names of ancestors to the map or material. The top level is at the top of the list, the current level is at the bottom, and intermediate levels appear between them.

Procedure

To give a material a different name:

- Edit the name field that appears below the Material Editor toolbar.

  The name of the active material appears in the title bar of the Material Editor dialog.

Type Button (Materials and Maps)

Material Editor > Type button

Material Editor menu > Material menu > Change Material/Map Type

Click the Type button to display the *Material/Map Browser (page 2–1412)* and choose which material type or map type to use.

When changing a material's type, the original material type is replaced unless you choose a *compound material (page 3–922)*, in which case a *Replace Material dialog (page 2–1456)* is displayed. The Replace Material dialog lets you choose between discarding the original material or using it as a sub-material within the new material.

For a standalone map (a map at the top level), clicking the Type button lets you change the map type instead of the material type. However, you can't use this button to make a map standalone. To make a standalone map, you have to click *Get Material (page 2–1439)* and choose a map from the Browser it displays.

When you change the type of a standalone map, a *Replace Map dialog (page 2–1456)* is displayed. The Replace Map dialog lets you choose between discarding the original map or using it as a sub-map within the new map.

Animating Materials

This topic provides some suggestions about how to animate materials.

Animating Basic Materials

In general, you animate a basic material by changing its parameters in different keyframes while Auto Key is active. 3ds Max interpolates values between keyframes, as it does when you animate transforms and modifiers.

Be aware that the tracks for a material assigned to an object are distinct from the material tracks that belong to the Material Editor: animating a material in the Material Editor affects the scene only if the material is *hot (page 3–953)*.

Animating Mapped Materials

As with basic materials, you can create animation keys for map parameters.

The noise parameters and the *Noise map (page 2–1674)* itself provide the Phase parameter specifically for animating the noise function.

You can also produce an animated material by applying an animated bitmap as a map. This

automated bitmap can be an *AVI (page 3–609)* or *MOV (page 3–621)* file, or an image sequence in the form of an *IFL (page 3–616)* file.

Changing One Material into Another

Probably the easiest way to animate the change of one material into another is to create a *Blend material (page 2–1588)*, make the two other materials its sub-materials, and then animate its Mix Amount parameter.

Preview and Playback

When you animate a material, or use an animated bitmap in a material, you can create and view a preview of this material before you decide to use it in a fully rendered animation.

See also

Creating and Playing Animated Material Previews (page 2–1450)

Synchronizing an Animated Bitmap with the Scene (page 2–1450)

Creating and Playing Animated Material Previews

You can create a preview movie of an animated material. Use *Make Preview (page 2–1434)* in the Material Editor. Clicking this button displays the *Create Material Preview dialog (page 2–1452)*, which like the Renderer has controls for which frames to render, playback speed, and output size.

The preview movie is saved in the *\previews* subdirectory in a file called *_medit.avi*. Each time you create a new preview, the Material Editor overwrites this file. To keep a preview movie on hand, save it under a different name by using the Save Preview button on the same flyout as Make Preview.

Procedures

To create a preview of an animated material:

1. Choose Make Preview from the flyout.
2. In the Create Material Preview dialog, set the preview conditions you want, and then click OK.

   The Material Editor creates the preview.

To play back a preview of an animated material:

1. Choose Play Preview from the flyout.

   The Material Editor starts the Windows Media Player with the active preview movie loaded.
2. Click Play to view the preview movie.

To save a preview movie under a different name:

1. Choose Save Preview from the flyout.

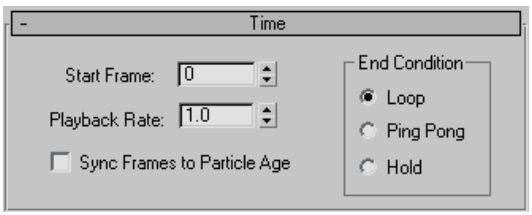
   The Material Editor displays a file save dialog.
2. Enter a new name for the preview, and then click OK.

Synchronizing an Animated Bitmap with the Scene

Material editor > Bitmap map > Time rollout

Bitmaps (page 2–1631) have a Time rollout with controls that let you synchronize an animated bitmap with scene animation.

Interface



Start Frame—The frame number of the 3ds Max scene at which the bitmap animation begins to play.

Playback Rate—Sets the bitmap's playback rate. This value is a multiplier: 1.0 is one bitmap frame per scene frame, 2.0 is twice as fast, 3.33 is 1/3 as fast, and so on. Default=1.0.

Sync Frames to Particle Age—When on, the software synchronizes the frames of a bitmap sequence to the age of particles to which the map is applied. With this effect, each particle displays the sequence from the start when it is born, rather than being assigned whichever frame is current. Default=off.

When using Particle Flow, assign the material containing the Bitmap map to a Material Dynamic operator. For more details and a procedure, see *Material Dynamic Operator (page 2–191)*.

Note: This functionality is not supported by the mental ray renderer.

End Condition group

Determines what happens after the last frame of the bitmap animation if the animation is shorter than the scene.

Loop—Causes the animation to repeat over and over again from the beginning.

Ping-Pong—Causes the animation to be played forward and then backward repeatedly, making every animated sequence "loop smoothly."

Hold—Freezes on the last frame of the bitmap animation.

Material Editor Subdialogs

Copy (Instance) Map Dialog

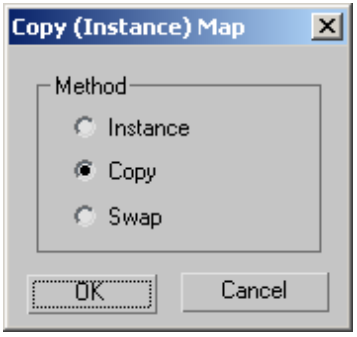
Material Editor > Maps rollout > Drag one map button to another.

The Copy (Instance) Map dialog is displayed when you copy a map by dragging and dropping a map button. It gives you the choice of copying the map, making the newly assigned map an instance of the one you dragged, or swapping maps.

You can drag map buttons in the Maps rollout, in the Basic Parameters rollout, or from one rollout to the other.

Tip: Usually when you use the same map for multiple parameters, such as both self-illumination and opacity, it is easier to work with an instance rather than a copy.

Interface



Instance—Makes an instance of the map you dragged. The newly assigned map is not independent. Adjusting the parameters of one map or the other changes both of them.

Copy—Copies the map you dragged. The newly assigned map is a copy whose parameters you can adjust independently.

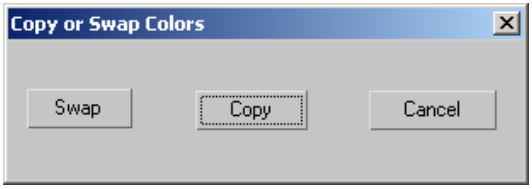
Swap—Swaps the maps. This option isn't displayed when you drag from one rollout to another.

Copy or Swap Colors Dialog

Material Editor > Basic Parameters rollout or Extended Parameters rollout or both > Drag one color swatch to another.

The Copy or Swap Colors dialog is displayed when you copy a color by dragging and dropping a color swatch. It gives you the choice of copying the color or swapping the colors, trading one for the other.

Interface



Swap—Swaps the colors.

Copy—Copies the color you dragged.

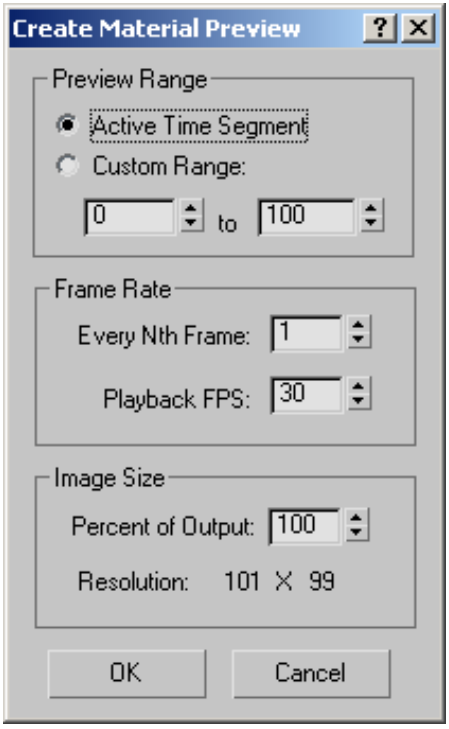
Create Material Preview Dialog



Material Editor > Make/Play/Save Preview flyout > Make Preview

The Create Material Preview dialog is displayed when you click *Make Preview (page 2–1434)* to preview an animated material. By default, the preview is saved in the *\previews* subdirectory with the name *_medit.avi*. You can use Save Preview to give the preview a different name so it won't be overwritten the next time you use Make Preview.

Interface



Preview Range group

Active Time Segment—Renders the *active time segment (page 3–904)*.

Custom Range—Renders a custom range from the start to the end frame you specify in the spinners below.

Frame Rate group

These controls specify the playback rate.

Every Nth Frame—Renders a regular sampling of the animation. For example, a value of 8 includes only every eighth frame in the preview. Default=1 (every frame).

Playback FPS—Specifies the playback rate in frames per second. Default=30 (full speed).

Image Size group

Percent of Output—Specifies the resolution of the preview. This value is a percentage; 100 percent has a resolution of 101 x 99 pixels (the size of a sample slot in the 3 x 2 array). Default=100.

Duplicate Name Dialog (Material Library)

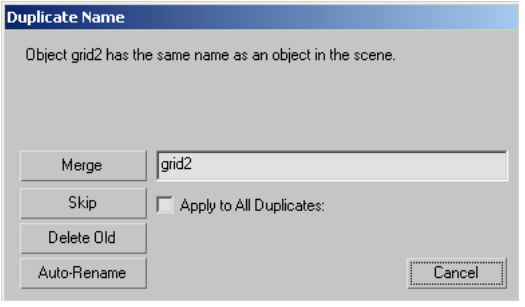
Material Editor > Get Material or Type button > Browse From group > Select Mtl Library. > File group > Open material library. > File group > Merge > Merge Material Library dialog > Open another material library or a 3ds Max scene. > Merge dialog > Select materials to merge > OK

File > XRef Objects > XRef Objects dialog > Create XRef Record from File > Choose a file. > Select objects to XRef. > Duplicate Name dialog

The Duplicate Name dialog is displayed if, after you click OK in the *Merge dialog (page 2–1453)*, one or more materials to merge have the same name as materials in the open library.

If this dialog appears while you are using *XRef Objects (page 3–394)*, in the XRef Objects dialog you can see the updated material name in the "Scene Name" column, while the original name in the source scene appears in the "Source Name" column.

Interface



Material name—At the right, the dialog displays a duplicate material name. You can edit the name to make it unique before you merge it with the open (current) library by clicking the Merge button.

Apply to All Duplicates—If you turn this on before you use the buttons, all subsequent incoming materials with duplicate names are treated the same way as the current one, and no further alert messages are displayed.

Use this option when you know that you've got several duplicate materials, and don't need constant reminders.

This check box is unavailable if you edit the material name.

Merge—Merges the material with the open, current library. This button is unavailable unless you edit the duplicate name to be a different, unique name.

Skip—Skips this material and doesn't merge it with the open, current library.

Delete Old—Deletes the "old" material in the open, current library and replaces it with the material to merge.

Auto-Rename—Click to have 3ds Max automatically rename the material by appending a sequence number to the duplicate material name.

Cancel—Cancels further merging of materials with duplicate names. If you have already merged some materials, they appear in the open, current library.

Merge Dialog (Material Library)

Material Editor > Get Material or Type button > Browse From group > Select Mtl Library. > File group > Open material library. > File group > Merge > Merge Material Library dialog > Open another material library or a 3ds Max scene.

The Merge dialog displays a list of materials to merge with the open library. The materials are from a different material library or from a 3ds Max or **VIZ Render** (*DRF (page 3–527)*) scene.

Procedures

To merge all materials:

- Click All.

To select a single material to merge:

- Click the material's name.

To select materials to merge one by one:

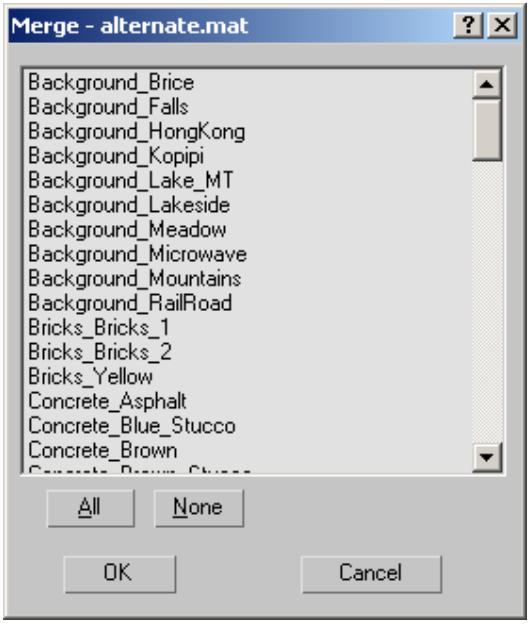
1. Click a material's name.
2. Hold down Ctrl, and click other material names.

To select a group of contiguous materials to merge:

1. Click a material's name.
2. Hold down Shift, and click another material's name.

   The dialog selects the two materials you clicked, and all materials whose names lie between the two.

Interface



Material list—Shows the names of all materials in the library or scene.

All—Selects all materials in the list.

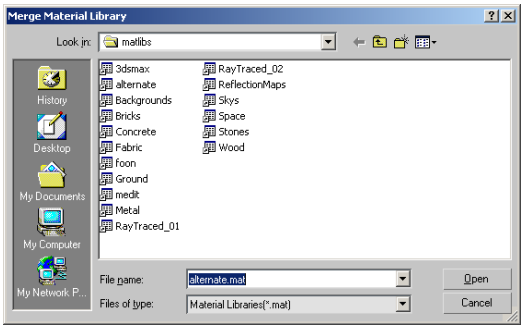
None—Deselects all materials in the list.

Merge Material Library Dialog

Material Editor > Get Material or Type button > Material/Map Browser > Browse From group > Select Mtl Library. > File group > Open material library. > File group > Merge

The Merge Material Library dialog lets you merge materials in the open material library with materials in another material library or another 3ds Max or **VIZ Render** (*DRF (page 3–527)*) scene.

Interface



This is a standard Windows file dialog. It lets you select either material library (*.mat)* files, 3ds Max (*.max*), or **VIZ Render** (*DRF (page 3–527)*) scene files.

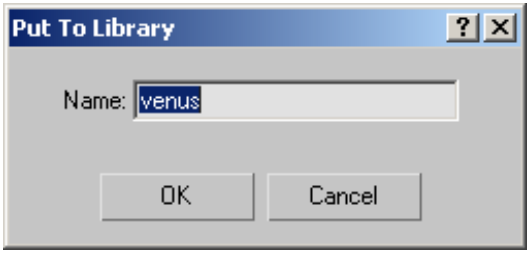
When you click OK, a further *Merge dialog (page 2–1453)* is displayed. This lets you select by name the materials you want to merge with the open library.

Put to Library Dialog

Material Editor > Put To Library

The Put To Library dialog is displayed when you want to save the material in an active *sample slot (page 2–1420)* into a material library file. It lets you change the material's name before you save it.

Interface



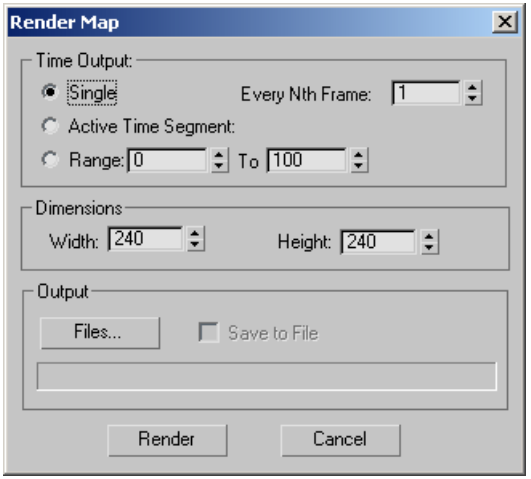
Name—Shows the name of the material to save. You can edit this name to save it under a different name.

Render Map Dialog

Material Editor > Right-click a sample slot. > Render Map > Render Map dialog

The Render Map dialog is displayed when you use Render Map on the *Sample Slots Right-Click Menu (page 2–1422)* to render the map displayed in a sample slot

Interface



Time Output group

These controls specify how many frames to render.

Single—Renders a single frame.

Every Nth Frame—Renders a regular sampling of the animation. For example, a value of 8 includes only every eighth frame in the preview. Default=1 (every frame).

Active Time Segment—Renders the *active time segment (page 3–904)*.

Range—Renders a custom range from the start to the end frame you specify in the spinners below.

Dimensions group

These controls specify the size of the rendered frames, in pixels.

Width—Specifies the frame width, in pixels.

Height—Specifies the frame height, in pixels.

Output group

These controls let you save the rendered map to a file.

Files—Click to display a file dialog that lets you specify where to save the rendered map.

When you click Save in the file dialog, you might see an additional dialog that gives you options specific to the file type you chose. These are the file formats available for a rendered map:

AVI (page 3–609)

BMP (page 3–610)

Kodak Cineon (page 3–610)

Encapsulated Postscript (page 3–612)

JPEG (page 3–620)

PNG (page 3–628)

MOV (QuickTime) (page 3–621)

SGI (page 3–633)

RLA (page 3–630)

RPF (page 3–631)

Targa (page 3–633)

TIFF (page 3–634)

Save to File—When on, the rendered map is saved to a file. When off, the map is rendered only to a *rendered frame window (page 3–5)*. This check box is unavailable unless you use Files to specify a file name, when it defaults to on.

File name field—Displays the name of the file you chose.

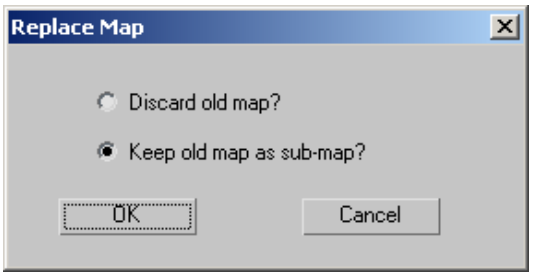
Render—Renders the map to a *rendered frame window (page 3–5)*, and to a file if you chose one.

Replace Map Dialog

Material Editor > Go to a map level or a standalone map. > Type button > Material/Map Browser > Choose a compound map.

The Replace Map dialog is displayed when you change a map type to any type of map that can have sub-maps. It gives you the choice of replacing the original ("old") map completely, or using the original map as a sub-map of the new map.

Interface



Discard old map—Discards the old map.

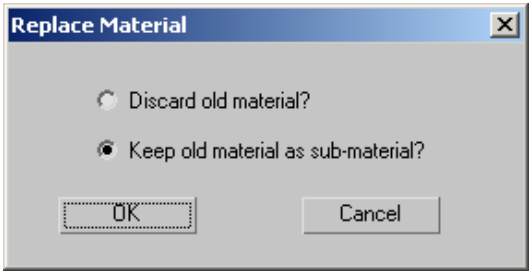
Keep old map as sub-map—Retains the old map as a sub-map.

Replace Material Dialog

Material Editor > Type button > Material/Map Browser > Choose a compound material.

The Replace Material dialog is displayed when you change a material type to one of the *compound material types (page 2–1587)*. It gives you the choice of replacing the original ("old") material completely, or using the original material as a sub-material of the new material.

Interface



Discard old material—Discards the old material.

Keep old material as sub-material—Retains the old material as a sub-material.

Update Scene Materials Dialog

Material Editor > Get Material > Material/Map Browser > Update Scene Materials from Library > Update Scene Materials dialog

The Update Scene Materials dialog lets you use a library to update materials in the scene. It appears only when you click Update Scene Materials From Library in the *Material/Map Browser (page 2–1412)*, and the scene contains materials that have the same name as materials in the library.

Procedures

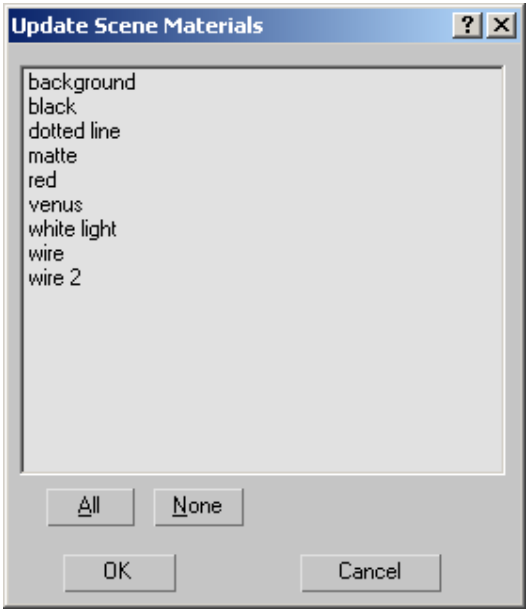
To update a material in the scene and replace it with a material from the library:

- In the Update Scene Materials dialog, select the name of the material, and then click OK.

To leave a material in the scene unchanged, do one of the following:

- In the Update Scene Materials dialog, make sure the name of the material is *not* selected, and then click OK.
- Click Cancel.

Interface



Material name list—Shows the materials that have the same name in the library and in the scene.

All—Selects all material names in the list.

None—Deselects all material names in the list.

Materials

Rendering menu > Material Editor > Type button > Material/Map Browser > In Show group, turn off Maps. > Material types are listed.

Materials create greater realism in a scene. A material describes how an object reflects or transmits light. You assign materials to individual objects or selection sets; a single scene can contain many different materials.

Different materials have different uses.

- *Standard material (page 2–1465)* is the default material. This is a versatile surface model with a large number of options.
- *Raytrace material (page 2–1512)* can create fully raytraced reflections and refractions. It also supports fog, color density, translucency, fluorescence, and other special effects.
- *Architectural material (page 2–1535)* provides a physically accurate material. It is especially intended for use with the default scanline renderer and *radiosity (page 3–51)*.
- *mental ray materials (page 2–1543)* are provided for use with the *mental ray renderer (page 3–78)*.
- *Matte/Shadow material (page 2–1584)* is specifically for making an object into a *matte object (page 3–971)* that reveals the current *environment map (page 3–934)*. A matte object is effectively invisible in the scene, but it can receive shadows cast onto it from other objects.
- *Shell material (page 2–1600)* is for storing and viewing *rendered textures (page 3–144)*.
- *Advanced Lighting Override material (page 2–1601)* is used to fine-tune the effects of a material on *radiosity solutions (page 3–51)* or the *Light Tracer (page 3–44)*. This material is not required for calculating advanced lighting, but it can help improve the result.

  Note: The Architectural material has its own controls for adjusting advanced lighting.
- *Lightscape material (page 2–1604)* helps support import and export of data from the Lightscape product.
- *Ink 'n Paint material (page 2–1605)* gives a cartoon appearance to objects.
- The *DirectX 9 Shader material (page 2–1613)* enables you to shade objects in viewports using DirectX 9 (DX9) shaders. To use this material, you must have a display driver that supports DirectX 9, and you must be using the Direct3D display driver.
-  The *XRef material (page 2–1616)* lets you externally reference a material in a different scene file.

Other material types fall into the category of *Compound materials (page 2–1587)*.

Compound Materials

Compound materials combine other materials in some way.

- *Blend (page 2–1588)* material mixes two materials on a single side of a surface.
- *Composite (page 2–1589)* material mixes up to 10 materials, using additive colors, subtractive colors, or opacity mixing.
- *Double-Sided (page 2–1591)* material lets you assign different materials to the front and back faces of an object.
- *Morpher (page 2–1592)* material uses the *Morpher modifier (page 1–729)* to manage multiple materials over time.
- *Multi/Sub-Object (page 2–1594)* material uses the sub-object level to assign multiple materials to a single object, based on material ID values.
- *Shellac (page 2–1597)* material superimposes one material on another using additive composition.
- *Top/Bottom (page 2–1599)* material lets you assign different materials to the top and bottom of faces of an object.

Procedures

To get a material:

1.  Click Get Material on the Material Editor toolbar.

   The *Material/Map Browser (page 2–1412)* is displayed.
2. Double-click a material type (not a map type) in the list, or drag the material to a sample slot.

The Material Editor replaces the original material.

To change a material type:

1. At the level of a material, click the Type button below the Material Editor toolbar.

   A modal *Material/Map Browser (page 2–1412)* is displayed. If you were at a material when you clicked Type, the Browser lists only materials (if you were at a map, it lists only maps).

2. Choose a material from the list, and then click OK.

   If you choose a compound material, a *Replace Material (page 2–1456)* dialog is displayed. This dialog lets you choose whether to keep or discard the original material.

   The Material Editor now displays controls for the new material.

SuperSampling Rollout

Material Editor > Architectural material > SuperSampling rollout
Material Editor > Raytrace material > SuperSampling rollout
Material Editor > Standard material > SuperSampling rollout

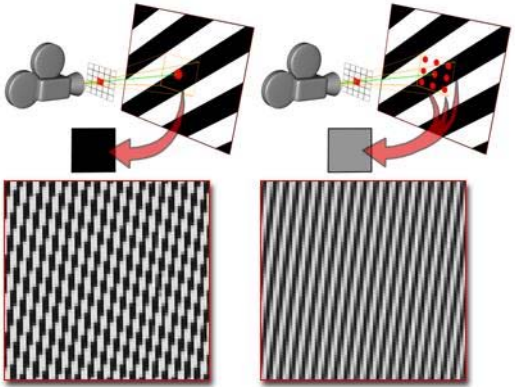
The SuperSampling rollout is used by Architectural, Raytrace, and Standard. It lets you choose a *supersampling method (page 3–1018)*. Supersampling performs an additional *antialiasing (page 3–907)* pass on the material. This requires more time but can improve image quality. Supersampling is especially helpful when you need to render very smooth specular highlights, subtle bump mapping, or high resolutions.

In 3ds Max the default is to apply a single supersampling method to all materials in the scene. This feature gives you more control over your scenes, especially larger models that make use of many materials, by letting you control the supersampling at a global level from the Rendering dialog. You can override this locally by turning off Use Global Settings. It also gives you file compatibility and workflow replication with DRF files imported from VIZ Render.

Note: SuperSampling settings are ignored by the *mental ray Renderer (page 3–78)*, which has its own sampling method.

Use supersampling when you notice artifacts in your final renderings. For example, a thin bump map might produce scintillating, jagged bumps that supersampling can correct. Supersampling requires considerably more time to render, although it does not necessarily require any additional RAM.

Note: Supersampling is *not* processed when you turn off Antialiasing in the production *renderer (page 3–27)*. You can also globally disable supersampling for all materials using the parameters for the *default scanline renderer (page 3–38)*. Globally disabling supersampling can speed up test renderings.



Supersampling uses smaller sampling points, and returns averaged values to increase the antialiasing effect.

Supersampling and Antialiasing

Supersampling is one of several antialiasing techniques. Textures, shadows, highlights, and

raytraced reflections and refractions all have their own preliminary antialiasing strategies. Supersampling is an additional step that provides a "best guess" color for each rendered pixel. The supersampler's output is then passed on to the renderer, which performs a final antialiasing pass.

A single rendered pixel represents an area of the scene's geometry. The pixel can stand in for multiple colors, especially when it appears at the edge of an object or a region of color. This is where aliasing effects occur.

When supersampling is turned off, 3ds Max simply looks at the center of the geometry covered by the pixel, and uses that for the pixel color.

When you turn on supersampling, a supersampler performs an additional antialiasing pass on the material. You can choose one of four supersamplers. As the name implies, a supersampler takes additional samples of geometry color in or near each pixel, in order to obtain a more accurate pixel color that is less prone to aliasing error.

These are the supersampling methods:

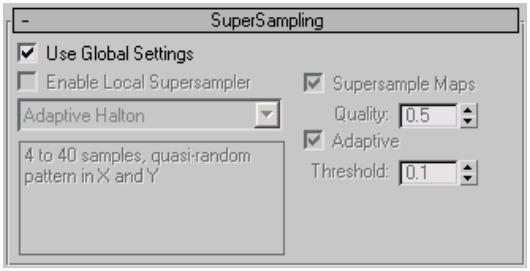
Name	Description
Adaptive Halton	Spaces samples along both X and Y axes according to a scattered, "quasi random" pattern. Depending on Quality, the number of samples can range from 4 to 40. This method is adaptive, as described below.
Adaptive Uniform	Spaces samples regularly, from a minimum quality of 4 samples to a maximum of 36. The pattern is not square, but skewed slightly to improve accuracy in the vertical and horizontal axes. This method is adaptive, as described below.
Hammersley	Spaces samples regularly along the X axis, but along the Y axis it spaces them according to a scattered, "quasi random" pattern. Depending on Quality, the number of samples can range from 4 to 40. This method is not adaptive.
MAX 2.5 Star	The sample at the center of the pixel is averaged with four samples surrounding it. The pattern is like the fives on dice. This is the supersampling method that was available in 3ds Max 2.5.

Regular sampling, as performed by the Adaptive Uniform and MAX 2.5 Star methods, is more prone to aliasing than the irregular patterns performed by the Adaptive Halton and Hammersley methods.

You can set a variable Quality value for the Adaptive Halton, Adaptive Uniform, and Hammersley methods. Quality can range from 0.0 to 1.0. A quality of 0.0 is minimal, with about four samples per pixel. A quality of 1.0 is the highest possible, with between 36 and 40 samples per pixel. High-quality supersampling is very time consuming.

Another setting for the Adaptive Halton and Adaptive Uniform supersamplers is the Adaptive toggle, which works in conjunction with the Threshold spinner. When Adaptive is on, these methods take fewer samples unless a change in color is greater than the Threshold value. In that case, they take the full number of samples specified by the Quality. To reduce the amount of time the supersampler spends, leave the Adaptive check box on.

Interface



Use Global Settings—When on, the material uses the supersampling options set in the *Default Scanline Renderer rollout (page 3–38)*. Default=on.

Enable Local Supersampling—When on, applies supersampling to the material. Default=off

Sampler drop-down list—Lets you choose which supersampling method to apply. The list is not active unless Use Global Settings is turned off. Default=Max 2.5 Star.

Supersample Maps—When on, the maps applied to the material are supersampled as well. When off, the supersampler uses pixel averages for maps. This switch is active only when Use Global Settings is turned off. Default=on.

Quality—Adjusts the quality of supersampling by controlling the number of samples used for each pixel. At the minimum value, 0.0, four samples are taken per pixel. At the maximum value, 1.0, about 40 samples are taken per pixel (the value varies depending on which shader is active). Range=0.0 to 1.0. Default=0.5.

This spinner is unavailable for the Max 2.5 Star method.

Adaptive—Visible only for the Adaptive Halton and Adaptive Uniform methods. When on, these methods take fewer samples than the Quality specifies unless samples show a change in color greater than the Threshold value. In that case, they take all the samples specified by the Quality.

Turning on Adaptive On can reduce the amount of time required to supersample. Default=on.

Threshold—Controls the Adaptive methods. Visible only for the Adaptive Halton and Adaptive Uniform methods. A change in color greater than the Threshold value causes the adaptive methods to take the full number of samples specified by the Quality. If the color does not change as much, the adaptive method takes fewer samples and does not require as much processing time. Can range from 0.0 to 1.0. Setting Threshold to 0.0 has the same effect as turning off Adaptive On. Default=0.1.

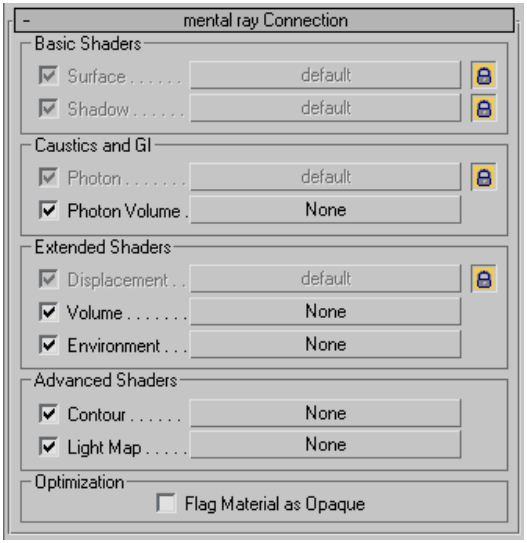
mental ray Connection Rollout

Material Editor > Click a sample slot that contains a material other than a Multi/Sub-Object or a mental ray material. > mental ray Connection rollout

The mental ray Connection rollout is available for all types of materials except the Multi/Sub-Object material and the mental ray materials themselves (for which it would be redundant). With this rollout you can add mental ray shading to conventional 3ds Max materials. These effects are visible only when you use the *mental ray renderer (page 3–78)*.

Important: The mental ray Connection rollout does not appear unless you have enabled the mental ray extensions by using the *mental ray Preferences panel (page 3–837)*. In addition, you can't assign shaders to the options in this rollout unless the mental ray renderer is the currently active renderer.

Interface



For each kind of shader on this rollout, there is a toggle and a button.

- The toggle controls whether the assigned shader is active or not.

  If no shader is assigned, the toggle has no effect.

- The button lets you assign a shader to the component type. Clicking it displays the *Material/Map Browser (page 2–1412)*: assigning a shader is just like assigning a map to a component of a standard material.

  While a shader is assigned, its name appears on the button.

-  In addition to the toggle and button, some of the shader types have a lock button to the right. When button is on, the component is inherited from the base material, and you can't assign a shader. For example, by default the Surface component is locked, and the surface is shaded using the settings of the 3ds Max material (basic parameters, maps, and so on). Turn off this button to replace the base material's settings with a mental ray shader.

  Note: Using a shader for the Surface component can result in a material whose appearance in mental ray renderings is completely different from the appearance it has in the sample slot, viewports, and scanline renderings.

Basic Shaders group

Surface—Shades the surface of objects that have this material. Default=locked to parent material.

In addition to any of the usual 3ds Max materials, the surface component can be assigned the following mental ray materials or shaders:

Shader	Library
Bump (page 2–1716)	3ds Max
DGS Material (page 2–1580)	3ds Max
Dielectric	base
Dielectric Material (page 2–1719)	3ds Max
Edge	lume
Facade	lume
Glass	lume
Glow	lume
Landscape	lume
Material to Shader (page 2–1723)	3ds Max
Metal	lume
Ocean	lume
Opacity	base
Reflect	base
Refract	base
Shader List (page 2–1723)	3ds Max
Stain	lume
Translucency	lume
Transmat	physics
Transparency	base
Two Sided	base

Shader	Library
UV Generator (page 2–1724)	3ds Max
Water Surface	lume
Wet-Dry Mixer	lume
XYZ Generator (page 2–1729)	3ds Max

Note: Unlike a standard 3ds Max material, if you assign the Surface component a bitmap with tiling turned off, the original surface color does not "show through." In renderings, you see only the untiled map, and none of the rest of the object.

Shadow—Assigns a shadow shader. Default=locked to parent material.

The shadow component can be assigned the following shaders:

Shader	Library
Edge Shadow	lume
Facade	lume
Glass	lume
Glow	lume
Material to Shader (page 2–1723)	3ds Max
Metal	lume
Shader List (page 2–1723)	3ds Max
Shadow Transparency	base
Translucency	lume
Transmat	physics
Water Surface Shadow	lume

Caustics and GI group

Photon—Assigns a photon shader. Photon shaders affect how object surfaces respond to photons; that is, they control how the surfaces behave when generating caustics and global illumination. Default=locked to parent material.

The photon component can be assigned the following shaders:

Shader	Library
DGS Material (page 2–1580)	3ds Max
Dielectric Material Photon (page 2–1719)	3ds Max
Edge	lume
Glow	lume
Material to Shader (page 2–1723)	3ds Max
Metal	lume
Photon Basic	base
Translucency	lume
Transmat	physics

Photon Volume—Assigns a photon volume shader. Photon volume shaders affect how an object's volume responds to photons; that is, they control how the volume behaves when generating caustics and global illumination.

The photon volume component can be assigned the following shaders:

Shader	Library
Material to Shader (page 2–1723)	3ds Max
Parti Volume Photon	physics
Shader List (page 2–1723)	3ds Max

Extended Shaders group

Displacement—Assigns a *displacement shader (page 3–96)*. Default=locked to parent material.

The displacement component can be assigned the following shaders:

Shader	Library
3D Displacement (page 2–1714)	3ds Max
Material to Shader (page 2–1723)	3ds Max
Ocean	lume

Volume—Assigns a *volume shader (page 3–95)*.

The volume component can be assigned the following shaders:

Shader	Library
Beam	lume
Material to Shader (page 2–1723)	3ds Max
Mist	lume
Parti Volume Photon	physics
Shader List (page 2–1723)	3ds Max
Submerge	lume

Environment—Assigns an environment shader. The environment shader provides an environment local to the material. It is visible if the material is reflective or transparent.

The environment component can be assigned the following shaders:

Shader	Library
Environment (page 2–1721)	3ds Max
Material to Shader (page 2–1723)	3ds Max
Shader List (page 2–1723)	3ds Max

Advanced Shaders group

Contour—Assigns a *contour shader (page 3–96)*.

The contour component can be assigned the following shaders:

Shader	Library
Combi	contour
Curvature	contour
Depth Fade	contour
Factor Color	contour
Layer Thinner	contour
Simple	contour
Width From Color	contour
Width From Light	contour
Width From Light Dir	contour

Light Map—Assigns a light map shader.

Warning: No light map shaders are provided with 3ds Max. This option is for users who have access to light map shaders via other shader libraries or custom shader code.

Optimization group

Flag Material as Opaque—When on, indicates that the material is fully opaque. This tells the mental ray renderer that it doesn't need to process transparency for this material, or to use the shadow shader. This can improve rendering time. Default=off.

DirectX Manager Rollout

Material Editor > DirectX Manager rollout

Lets you select a DirectX viewport shader for viewing Direct3D hardware shaders. DirectX shaders require the *Direct3D graphics driver (page 3–844)*, which uses DirectX. With DirectX shading, materials in a viewport more accurately represent how the material will appear in another application, or on other hardware such as a game engine.

DirectX viewport shaders are especially useful for previewing *texture-baked materials (page 3–144)*.

3ds Max provides two DirectX shaders:

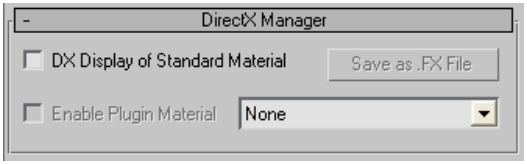
- *LightMap shader (page 2–1614)*
- *Metal Bump shader (page 2–1614)*

Note: This rollout does not appear for Multi/Sub-Object and Shell materials, which are simply containers of other materials.

See also

DirectX 9 Shader Material (page 2–1613)

Interface



DX Display of Standard Material—When on, displays the active material as a DX shader. You can save the material as an FX file by clicking Save As .FX File.

For full support of this feature, DX9 must be active.

Warning: Not all standard 3ds Max features can be represented by DX shaders, and you it is possible to create a standard material that is more complex than a DX video card can display. DX shaders do support the most commonly used material components: diffuse color (or texture), specular highlights, opacity, bump mapping, and reflections.

Save as .FX File—Click to display a Save Effect File dialog that lets you save the active material as an *FX file (page 3–946)*.

Enable Plugin Material—Turn on to use the chosen DirectX shader in shaded viewports. Default=off.

When not enabled, viewports continue to use the default *viewport (interactive) renderer (page 3–1030)* (or the *ActiveShade renderer (page 3–17)*, if that has been chosen).

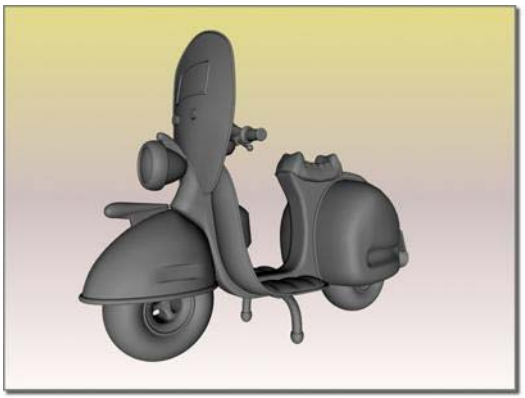
This toggle is unavailable if no shader plug-in has been chosen from the drop-down list, and when DX Display Of Standard Material is on.

Plug-in drop-down list—Use the drop-down list to choose a DirectX viewport shader.

The list is unavailable when DX Display Of Standard Material is on.

Standard Material

Material Editor > Type button > Material/Map Browser > Standard



Scooter rendered with the default standard material

Standard material is the default material in the Material Editor sample slots. There are several other kinds of *material types (page 2–1457)*.

The Standard material type provides a fairly straightforward way to model surfaces. In the real world, the appearance of a surface depends on how it reflects light. In 3ds Max, a standard material simulates a surface's reflective properties. If you don't use *maps (page 3–968)*, a standard material gives an object a single, uniform color.

This topic introduces the controls for Standard material, exclusive of mapping.

Standard Color Components

A surface of a "single" color usually reflects many colors. Standard materials typically use a four-color model to simulate this. (There are some variations depending on which *shader (page 2–1466)* you use.)

- *Ambient color (page 3–908)* is the color of the object in shadow.

- *Diffuse (page 3–929)* is the color of the object in direct, "good" lighting.
- *Specular (page 3–1014)* is the color of shiny highlights.

  Some shaders generate the specular color, rather than letting you choose it.
- *Filter (page 3–939)* is the color transmitted by light shining through the object.

  The Filter color component isn't visible unless the material's Opacity is less than 100 percent.

  Note: *Raytrace material (page 2–1512)* uses a different, six-color model to simulate surfaces. Several components are similar to those in Standard Material, but they behave differently in Raytrace.

When we describe an object's color in conversation, usually we mean its diffuse color. The choice of an ambient color depends on the kind of lighting. For moderate indoor lighting, it can be a darker shade of the diffuse color, but for bright indoor lighting and for daylight, it should be the complement of the primary (key) light source. The specular color should be either the same color as the key light source, or a high-value, low-saturation version of the diffuse color.

For more tips on choosing color components, see *Choosing Colors for Realism (page 2–1400)*.

Warning: When you change the shading type of a material, you lose the settings (including map assignments) for any parameters that the new shader does not support. If you want to experiment with different shaders for a material with the same general parameters, copy the material to a different *sample slot (page 2–1420)* before you change its shading type. That way, you can still use the original material if the new shader doesn't give you the effect you want.

Other Standard Material Components

A standard material's specular color appears in highlights. You can control the size and shape of the highlight. A polished surface has a small and strong highlight. A matte surface has a large, weak highlight, or no highlight at all.

Standard materials also have controls for making the object appear transparent, and for making it self-illuminating so that it appears to glow.

Along with the material's color components, *components* also refers to the parameters that control highlights, transparency, self-illumination, and so on.

See also

Material Components (page 2–1399)

Choosing Colors for Realism (page 2–1400)

Interface

The interface for a standard material is organized into several rollouts:

Shader Basic Parameters Rollout (page 2–1466)

Basic Parameters Rollout (Standard Material) (page 2–1470)

Extended Parameters Rollout (Standard Material) (page 2–1471)

SuperSampling Rollout (page 2–1459)

Maps Rollout (Standard Material) (page 2–1474)

Dynamics Properties Rollout (page 2–1479)

DirectX Manager Rollout (page 2–1464)

Shader Basic Parameters Rollout

Material Editor > Standard material > Shader Basic Parameters rollout > Choose shader from drop-down list.

The Shader Basic Parameters rollout lets you choose the type of shader to use with a *Standard (page 2–1465)* material. Some additional controls affect how the material appears.

Procedures

To set a material's shading type:

1. On the Shader Basic Parameters rollout, open the shader drop-down list.
2. Click the name of the shader type to use for the active material.

To use Wire mode:

- On the Shader Basic Parameters rollout, turn on Wire.

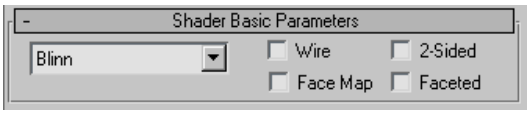
  The material is now shaded as a wireframe mesh. The wire portions of the geometry do not change; color components, shininess, and so on, remain the same.

  Tip: For a wireframe material, turn on the 2-Sided option as well.

  You have two choices for how wireframe materials are rendered. The controls for tuning wireframe shading are on the *Extended Parameters (page 2–1471)* rollout.

  - If you choose Pixels, the thickness of the wires maintains the same apparent thickness regardless of the scale of the geometry or how near or far the object is positioned. In other words, pixel wires have a constant display size as if the wires were traced over an image.
  - If you choose Units, the wires behave as if they were modeled in the geometry. They appear thinner at a distance and thicker at close range. Scaling a wireframe object does scale wire width.

Interface



shader drop-down list—Chooses a shader. The material's Basic Parameters rollout can change to show the controls for the shader you choose. Default shader=Blinn

There are seven different shaders. Some are named for what they do; others are named for their creators. These are the basic material shaders:

- *Anisotropic (page 2–1480)*: For surfaces with elliptical, "anisotropic" highlights. These highlights are good for modeling hair, glass, or brushed metal.
- *Blinn (page 2–1480)*: For rounder, softer highlights than Phong shading
- *Metal (page 2–1481)*: For metallic surfaces
- *Multi-Layer (page 2–1481)*: For surfaces with more complex highlights than Anisotropic
- *Oren-Nayar-Blinn (page 2–1482)*: For matte surfaces such as fabric or terra-cotta
- *Phong (page 2–1482)*: For surfaces with strong, circular highlights
- *Strauss (page 2–1483)*: For metallic and nonmetallic surfaces. The Strauss shader has a simpler interface than other shaders.
- *Translucent (page 2–1484)*: Similar to Blinn shading, the Translucent shader also lets you specify translucency, where light is scattered as it passes through the material.

For more information about the shaders, including illustrations, see *Understanding Shaders (page 2–1468)*.

Wire—Renders the material in *wireframe mode (page 3–1034)*. You can set the size of the wire in *Extended Parameters (page 2–1471)*.

2-Sided—Makes the material *2-sided (page 3–901)*. Applies the material to both sides of selected faces.

Face Map—Applies the material to the faces of the geometry. If the material is a mapped material, it requires no *mapping coordinates (page 3–967)*. The map is automatically applied to each facet of the object.

Faceted—Renders each face of a surface as if it were flat.

Understanding Shaders

For standard materials, a shader is an algorithm that tells 3ds Max how to calculate surface rendering. Each shader has a unique set of characteristics in order to serve a particular purpose. Some are named for what they do well, such as the Metal shader. Others are named for the person who developed them, such as the Blinn and Strauss shaders. The default shader in 3ds Max is the Blinn shader.

Note: In addition to the shaders listed below, 3ds Max supports plug-in shader types.

The following list describes the shaders supplied with the software:

- **Anisotropic:** Used for brushed metal or hair. Creates a highlight that is stretched and angled, rather than the standard circular highlight.



- **Blinn:** Has the same features as the Phong shader, but its mathematics are more accurate. This is the default shader for Standard materials.



- **Metal:** Used for making metals.



- **MultiLayer:** Two anisotropic shaders in one. Used to make two different highlights with independent controls. Simulates materials such as a metal that is covered with a shiny coat of wax.



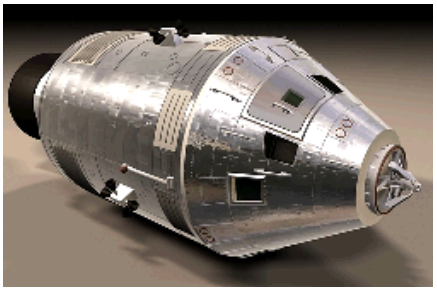
- **Oren-Nayar-Blinn:** An adaptation of the Blinn shader. It gives objects a porous, non-plastic appearance, and is suitable for surfaces like skin.



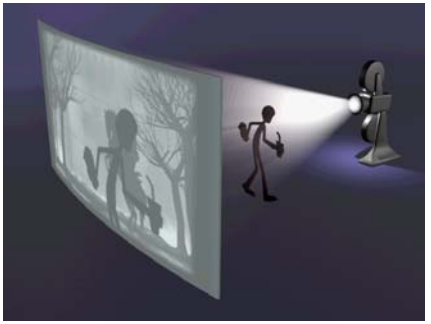
- **Phong:** A classic shading method that was the first to enable specular highlights. Suitable for plastic surfaces.



- **Strauss:** Suitable for metals. Allows you to control the degree of metallic characteristics of the material.



- **Translucent Shader:** Translucent shading is similar to Blinn shading, but it also lets you specify translucency. A translucent object allows light to pass through, and also scatters light within the object. You can use translucency to simulate frosted and etched glass.



Comparing Shader Parameters

A shader is an algorithm that tells the program how to calculate surface rendering. Each shader has a unique set of characteristics in order to serve a particular purpose.

Compare the parameters of different shader types:

1. Open the Material Editor and click an available sample slot.
2. In the list on the Shader Basic Parameters rollout, change Blinn to Anisotropic.

   The Blinn Basic Parameters rollout changes to the Anisotropic Basic Parameters rollout. Observe the differences in the available basic parameters.





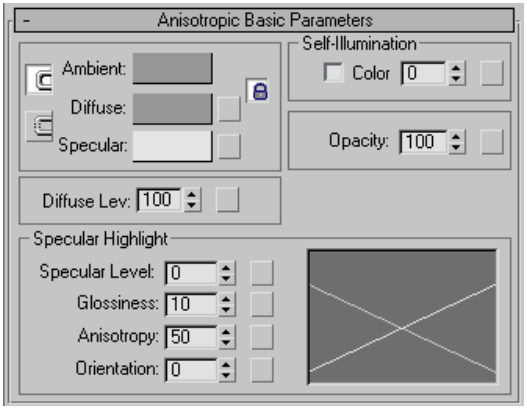
3. Select each shader type from the list and compare its parameters with the others. Some parameters are shared in common, but each shader has its own unique combination of settings.

For more information on shader types, see *Shader Basic Parameters Rollout (page 2–1466)*.

Basic Parameters Rollout (Standard Material)

Material Editor > Standard material > Basic Parameters rollout for the shader you've chosen

The Basic Parameters rollouts for Standard materials contain controls that let you set the color of your material, the shininess, the transparency, and so on, and specify *maps (page 3–968)* to use for the various components of the material.



Example: The Basic Parameters rollout for the Anisotropic shader.

Basic Parameters rollouts vary depending on which shader is chosen.

The Basic Parameters rollout changes depending on which kind of shader you choose in the *Shader Basic Parameters (page 2–1466)*.

Note: The Strauss shader's Basic Parameters rollout is simpler than those for other shaders. See *Strauss Shader (page 2–1483)* for a description.

Component Controls

The first part of the Basic Parameters rollout contains controls for overall material components. They are described in the following topics:

- *Color Controls (page 2–1485)* let you choose the material's color components, or replace them with maps.

- *Self-Illumination (page 2–1487)* makes a material appear lit from within.

  Self-illumination is not available for the *Strauss shader (page 2–1483)*.
- *Opacity (page 2–1488)* controls how opaque or transparent a material is.
- *Diffuse Level (page 2–1489)* controls the brightness of the diffuse color component.

  Diffuse Level is available only for the *Anisotropic (page 2–1480)*, *Multi-Layer (page 2–1481)*, and *Oren-Nayar-Blinn (page 2–1482)* shaders.
- *Roughness (page 2–1490)* controls how quickly the diffuse component blends into the ambient component.

  Roughness is available only for the *Multi-Layer (page 2–1481)* and *Oren-Nayar-Blinn (page 2–1482)* shaders.

Highlight Controls

The second part of the Basic Parameters rollout contains controls for specular highlights, which in some ways are the greatest difference between the various shaders. See these topics for a description:

- *Anisotropic Highlights (page 2–1480)*
- *Blinn, Oren-Nayar-Blinn, and Phong Highlights (page 2–1493)*
- *Metal Highlights (page 2–1494)*
- *Multi-Layer Highlights (page 2–1495)*

For information on highlights with the Strauss shader, see *Strauss Shader (page 2–1483)*.

Translucency Controls

For the *Translucent shader (page 2–1484)*, an additional group on the Basic Parameters rollout contains controls for *translucency (page 2–1491)*.

Extended Parameters Rollout (Standard Material)

Material Editor > Standard material > Extended Parameters rollout

The Extended Parameters rollout is the same for all shading types of Standard material. It has controls related to transparency and reflection, as well as options for Wire mode.

This topic contains tables of the Index of Refraction for some common physical materials. These can be used to create Standard materials with realistic transparency.

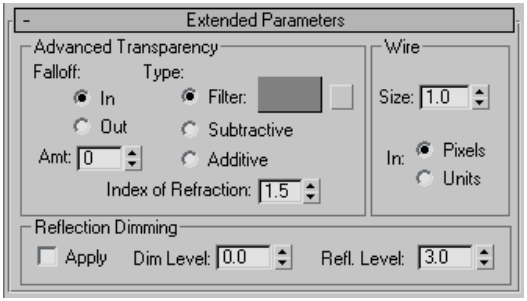
Additive Opacity and the Alpha Channel

By default, additive opacity does not generate an alpha value. In other words, the alpha value is zero, indicating no transparency. This gives correct results with backgrounds in renderings, but if you want to composite objects with additive opacity using *video post (page 3–311)* or a different compositing program, you might want to have additive opacity render with transparency. To do so, add the following line to the [Renderer] section of the *3dsmax.ini* file, and then restart 3ds Max:

AlphaOutOnAdditive=1

To revert to the default method of rendering additive opacity, in the *3dsmax.ini* file, change the value of AlphaOutOnAdditive back to 0 (zero), and then restart 3ds Max.

Interface



Advanced Transparency group

These controls affect the *opacity falloff (page 3–984)* of a transparent material.

Note: For the *Translucent shader (page 2–1484)*, these controls do not appear. They are replaced by the *Translucency controls (page 2–1491)* on the Basic Parameters rollout.

Falloff—Chooses whether falloff is in or out, and how great it is.

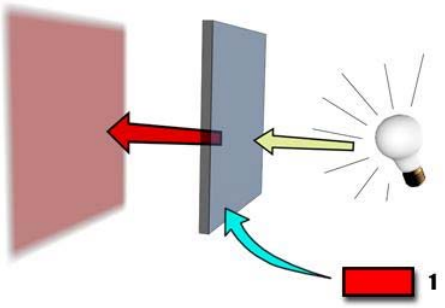
- **In—**Increases transparency toward the inside of the object, as in a glass bottle.
- **Out—**Increases transparency toward the outside of the object, as in a cloud of smoke.

Amt (Amount)—Specifies the amount of transparency at the outside or inside extreme.

Type—These controls choose how transparency is applied.

- *Filter (page 3–939)* computes a filter color that it multiplies by the color behind the transparent surface. Click the color swatch to change the filter color. Click the button to assign a map to the filter color component.

  The filter, or transmissive color, is the color transmitted through transparent or semi-transparent materials such as glass. You can use the filter color with volumetric lighting to create effects such as colored light through a stained-glass window. *Ray-traced shadows (page 3–1000)* cast by transparent objects are tinted with the filter color.



A shadow's color is changed with a red filter color.

- *Subtractive (page 3–1017)* subtracts from the color behind the transparent surface.
- *Additive (page 3–906)* adds to the color behind the transparent surface.

Index of Refraction—Sets the index of refraction (IOR) used by refraction maps and raytracing. The IOR controls how severely the material refracts transmitted light. Left at 1.0, the IOR of air, the object behind the transparent object does not distort. At 1.5 the object behind distorts greatly, like a glass marble. At an IOR slightly less than 1.0, the object reflects along its edges, like a bubble seen from under water. Default=1.0.

Common IORs (assuming the camera is in air or a vacuum) are:

Material	IOR Value
Vacuum	1.0 (exactly)
Air	1.0003
Water	1.333
Glass	1.5 (clear glass) to 1.7
Diamond	2.417

In the physical world, the IOR results from the relative speeds of light through the transparent

material and the medium the eye or the camera is in. Typically this is related to the object's density; the higher the IOR, the denser the object.

You can also use a map to control the index of refraction. IOR maps always interpolate between 1.0 (the IOR of air) and the setting in the IOR parameter. For example, if the IOR is set to 3.55 and you use a black-and-white Noise map to control IOR, the IORs rendered on the object will be set to values between 1.0 and 3.55; the object will appear denser than air. If, on the other hand, your IOR is set to 0.5, then the same map values will render between 0.5 and 1.0, as if the camera were under water and the object was less dense than the water.

Here are some more IOR values for various materials:

Material	IOR Value
Carbon Dioxide, Liquid	1.200
Ice	1.309
Acetone	1.360
Ethyl Alcohol	1.360
Sugar Solution 30%	1.380
Alcohol	1.329
Flourite	1.434
Quartz, Fused	1.460
Calspar2	1.486
Sugar Solution 80%	1.490
Glass, Zinc Crown	1.517
Glass, Crown	1.520
Sodium Chloride	1.530
Sodium Chloride (Salt) 1	1.544
Polystyrene	1.550
Quartz 2	1.553
Emerald	1.570
Glass, Light Flint	1.575
Lapis Lazuli	1.610
Topaz	1.610
Carbon Bisulfide	1.630
Quartz 1	1.644
Sodium Chloride (Salt) 2	1.644
Glass, Heavy Flint	1.650
Methylene Iodide	1.740
Ruby	1.770
Sapphire	1.770
Glass, Heaviest Flint	1.890
Crystal	2.000
Chromium Oxide	2.705
Copper Oxide	2.705
Amorphous Selenium	2.920
Iodine Crystal	3.340

Wire group

Size—Sets the size of the wire in *wireframe mode (page 3–1034)*. You can set either pixels or current units.

In—Chooses how to measure wire.

- **Pixels—**(The default.) Measures wire in pixels. With pixels, wires maintain the same apparent thickness regardless of the scale of the geometry or how near or far the object is positioned.
- **Units—**Measures wire in 3ds Max units. With units, the wires appear thinner at a distance and thicker at close range, as if they were modeled in the geometry.

Reflection Dimming group

These controls dim reflection maps that are in shadow.



Reflection dimming

Above: None

Below: 0.0 (100% dimming)

Apply—Turn on to use reflection dimming. When off, the reflection-mapped material is not affected by the presence or absence of direct light. Default=off.

Dim Level—The amount of dimming that takes place in shadow. At 0.0, the reflection map is completely dark in shadow. At 0.5, the reflection map is half dimmed. At 1.0, the reflection map is not dimmed and the material appears as if Apply were turned off. Default=0.0.

Refl. Level—Affects the intensity of the reflection that is *not* in shadow. The Reflection Level value multiplies the illumination level of the lit area of the reflection, to compensate for dimming. In most cases, the default value of 3.0 keeps the reflection in the lit area at about the same level it would appear if reflection dimming were not on.

Maps Rollout (Standard Material)

Material Editor > Standard material > Maps rollout

A material's Maps rollout lets you access and assign *maps (page 3–968)* to various components of the material.

You can choose from a large variety of map types. To find descriptions of these types, and how to set their parameters, see *Map Types (page 2–1617)*.

Assigning the Same Map to Different Parameters

Applying the same map to different parameters is useful in some cases. For example, using a pattern to map both self-illumination and opacity can make the pattern appear to glow and hover in space.

Blending Map Amounts for Opacity and Other Material Components

When you map a scalar component (such as Specular Level, Glossiness, Self-Illumination, and Opacity), the component's value in the *Basic Parameters (page 2–1470)* rollout is blended with its associated map Amount in the Maps rollout.

For example, when the Opacity spinner is set to 0, the map Amount spinner completely controls Opacity. That is, reducing the Amount value increases the transparency of the entire surface. On the other hand, when Opacity is 100, reducing the map's Amount value increases the opacity of the entire surface. You can adjust a Checker Opacity map so that the opaque areas remain opaque, while the transparent areas become semi-transparent.

Other scalar components behave in the same way. Setting the map's Amount to 100 applies all of the map. Setting the Amount to 0 is the equivalent of turning the map off. Intermediate Amount

values are blended with the value of the scalar component.

When you load old 3ds Max files or bring earlier materials from the Browser into the Materials Editor, the spinner values for Opacity, Specular Level, Glossiness, and Self-Illumination are altered, where necessary, to maintain the equivalent material effect.

Procedures

To assign a map:

1. In the Maps rollout, click a map button.

   A modal *Material/Map Browser (page 2–1412)* is displayed.

2. Use the Browse From buttons to choose where you want to look.

   If you choose Material Library and the dialog's display area is blank, you need to open a library file. Click the Open button and then choose the library to browse.

3. Use the display buttons to choose how you view maps.

   -  View List shows each map by name.
   -  View List + Icons shows a small preview and each map's name.
   -  View Small Icons shows a small preview for each map.
   -  View Large Icons shows a large preview for each map, along with the map's name.

     Tip: You can resize the Browser dialog to increase the size of the display area. This is especially useful when you view large icons.

4. Double-click the map you want.

To use the same map for different parameters:

1. In the Maps rollout, use a map button to assign a map.

   The Material Editor is now at the map level, and displays controls for the map parameters.

2.  Click Go To Parent to return to the material level, and then open the Maps rollout.

3. Drag the assigned map button to another map button.

   The *Copy (Instance) Map dialog (page 2–1451)* is displayed.

4. Choose Copy or Instance, and then click OK.

   If you choose Swap, the Material Editor swaps the two button assignments.

To view the parent material's parameters:

-  If you are currently at the map level in the Material Editor, click Go To Parent.

  The parameters for the map's parent material are displayed. Also, the Show End Result and Go To Parent buttons become unavailable.

To view a map's parameters:

- If you are currently at the material level in the Material Editor, click the button that corresponds to the map.

  The parameters for the map are displayed. Also, the *Show End Result (page 2–1446)* and *Go To Parent (page 2–1446)* buttons become available.

  In the Basic Parameters rollout, if a map has been assigned to a color component or parameter, the corresponding button displays a letter M. In the Maps rollout, if a map has been assigned, the corresponding button displays the map name.

To view a map's location:

-  Click Material/Map Navigator to view the Navigator.

  The *Material/Map Navigator (page 2–1447)* displays the hierarchy of the current material, which contains the map.

To go to a map using the Navigator:

- In the *Material/Map Navigator (page 2–1447)*, click the name of the map, or the green or red parallelogram to the left of the map's name.

  The Navigator goes to the level of the map, and the Material Editor displays the controls for the map you clicked.

  As the Navigator's map tree shows, maps for basic material components and parameters are one level below the material itself.

To preview a map in a sample slot:

1. Go to the level of the map, as described in previous procedures.

   The Material Editor displays the map's parameters.

2.  Turn off *Show End Result (page 2–1446)*.

   The sample slot shows the map instead of the material. If the map contains sub-maps, these are also visible.

   By default, the sample slot displays a map with no three-dimensional shading. You can change this in the *Material Editor Options dialog (page 2–1436)*.

To view the map interactively:

1. Select an object.
2. In the object's creation parameters, make sure that Generate Mapping Coords is on.

   If the object type does not have a Generate Mapping Coordinates toggle, you need to assign mapping coordinates by applying a *UVW Map modifier (page 1–922)*.
3. In the Material Editor, assign the mapped material to the object.
4. If you are at the material level (the top level), click the appropriate map button to go to the map level.
5.  Turn on *Show Map In Viewport (page 2–1445)*.

   The map appears on objects assigned the material in all shaded viewports. Now when you adjust the map, the viewports update to display the adjustments.

   Turning on Show Map In Viewport for one map automatically turns this button off for all other maps the material has.

   Viewports can display 2D maps such as Checker and Bitmap.

   Viewports can also display most kinds of 3D maps. The exceptions are Particle Age and Particle MBlur. Also, the appearance of the Falloff map in viewports give only a vague indication of how it will appear when rendered.

   Show Map In Viewport is unavailable if the active map type cannot display in viewports.

   Displaying mapped materials in a viewport can slow performance. If you don't need to view the texture, turn off its viewport display.

To turn off interactive texture display:

1. Go to the map level.

   If you are at the material level, click the appropriate map button to go to the map level.

2.  Turn off *Show Map in Viewport (page 2–1445)*.

   The object is shaded but the map no longer appears.

To turn a map off:

- In the Maps rollout, turn off the map's check box.

  The check box is to the left of the map's name.

To turn a map on:

- In the Maps rollout, turn on the map's check box.

  The check box is to the left of the map's name.

To change a map's strength:

- In the Maps rollout, adjust the map's Amount spinner.

  The material's sample slot reflects the change.

  Note: Adjusting a map's output (in the map's Output rollout) can also change the map's strength.

To move directly to an ancestor:

1. Below the Material Editor toolbar, click the arrow to the right of the map's *name field (page 2–1448)*.

   A drop-down list of ancestors is displayed.

2. Click a name in the Ancestor list.

   With this list, you can skip intermediate levels in the tree.

    The Ancestor drop-down list shows only part of the tree. It does not show side branches and siblings. To view these, use the *Material/Map Navigator (page 2–1447)* or the *Go Forward To Sibling (page 2–1447)* and *Go To Parent (page 2–1446)* buttons on the Material Editor toolbar.

To change a map type:

1. At the level of a map, click the button labeled Type below the Material Editor toolbar.

   A modal *Material/Map Browser (page 2–1412)* is displayed. If you were at a map, it lists only maps (if you were at a material when you clicked Type, the Browser lists only materials).

2. Choose a map type from the list, and then click OK.

   If you change a map type and the new map type can have component maps, a Replace Map dialog is displayed. This dialog gives you a choice between discarding the original map or using it as a component map.

   If the new map type does not have components, it simply replaces the original map type.

To create a standalone map tree:

1. Activate a sample slot.

2.  Click *Get Material (page 2–1439)*.

3. In the *Material/Map Browser (page 2–1412)*, make sure Browse From is set to New.

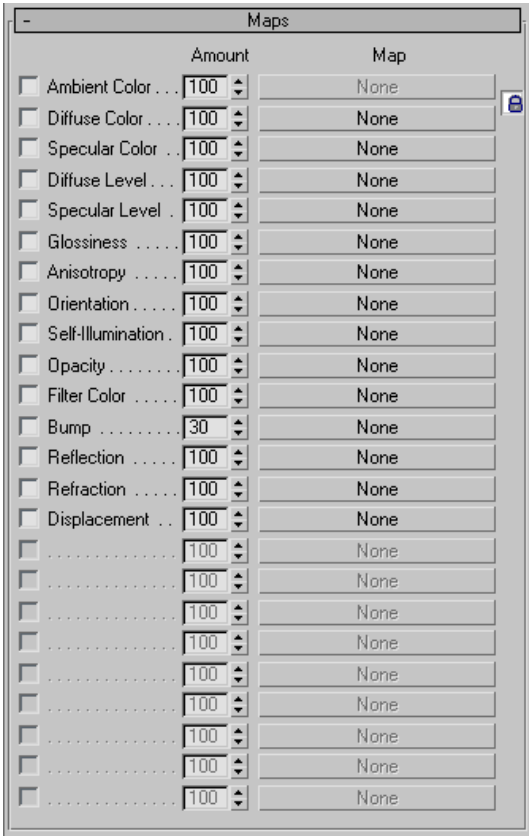
4. Double-click the name of the map type (not a material type) you want to use, or drag the map to a sample slot.

   The sample slot now contains a standalone map not associated with material parameters.

5. Use the Material Editor to modify the map as you would any other map.

   By default, the sample slot displays a map with no three-dimensional shading. You can change this in the *Material Editor Options dialog (page 2–1436)*.

Interface



The Maps rollout contains a wide button for each map type. Click this button to select a bitmap file stored on disk or to select a *procedural map type (page 3–997)*. After you select a map, its name and type appears on the button. Use the check box to the left of the button to turn the effect of the map off and on. When the check box is off, the map is not computed and has no effect in the renderer.

The Amount spinner determines the amount that the map affects the material expressed as a percentage of full intensity. For example, a diffuse map at 100% is completely opaque and covers the base material. At 50%, it is semi-transparent and the base material (the diffuse, ambient, and other colors of the material without mapping) shows through.

Ambient Color Mapping (page 2–1497)

Diffuse Color Mapping (page 2–1498)

Diffuse Level Mapping (page 2–1499)

Diffuse Roughness Mapping (page 2–1500)

Specular Color Mapping (page 2–1500)

Glossiness Mapping (page 2–1502)

Glossiness and Specular Level Settings (page 3–950)

Self-Illumination Mapping (page 2–1502)

Opacity Mapping (page 2–1503)

Filter Color Mapping (page 2–1503)

Anisotropy Mapping (page 2–1504)

Orientation Mapping (page 2–1505)

Metalness Mapping (page 2–1506)

Bump Mapping (page 2–1506)

Reflection Mapping (page 2–1508)

Refraction Mapping (page 2–1509)

Displacement Mapping (page 2–1511)

The Maps rollout can have unused, disabled control rows at the bottom. This is because the number of components that can be mapped varies depending on the current *shader (page 2–1466)*. The last four rows are always Bump, Reflection, Refraction, and Displacement, in that order.

Note: The sub-material and sub-map buttons for most materials and maps have check boxes beside each button. These turn that branch of the material or map off or on. For example, in the Top/Bottom material, the Top Material and Bottom Material buttons each have check boxes. Similarly, the Checker map has two map buttons, one for each color. Each button has check box beside it that lets you turn off that color's map.

Ambient and Diffuse Map Lock 

In the Maps rollout, the lock button to the right of the Diffuse Color map button locks ambient mapping to diffuse mapping. It is on by default. Usually it makes sense to use the same map for the ambient and diffuse components. To use different maps for ambient and diffuse, turn off the lock button. The map button for Ambient Color becomes available.

Dynamics Properties Rollout

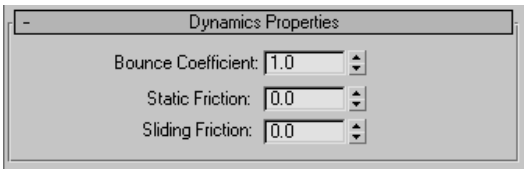
Material Editor > Standard material > Dynamics Properties rollout

The Dynamics Properties rollout lets you specify surface properties that affect the animation of an object upon collision with another object. If there are no collisions in your simulation, these settings have no effect. The dynamics properties are used by the *Dynamics utility (page 2–686)*.

Since the Dynamics Properties rollout is available at the top level of any material (including submaterials), you can specify different surface dynamic properties for each face in an object. There are also controls in the Dynamics utility that let you adjust the surface properties at the object level, but only the Materials Editor lets you alter the surface properties at the sub-object level, through use of a *Multi/Sub-Object material (page 2–1594)*.

As a default, the values in the Dynamics Properties rollout provide a surface that's similar to Teflon-coated hardened steel.

Interface



Bounce Coefficient—Sets how far an object bounces after hitting a surface. The higher the value, the greater the bounce. A value of 1 represents a "perfectly elastic collision," or a bounce in which no kinetic energy is lost. Default=1.0.

If you've seen the desktop toy with four ball bearings swinging back and forth on strings and hitting one another, you've seen an example that comes very close to a bounce coefficient of 1. Generally, hardened steel or a super ball have a bounce near 1, while lead has a bounce near 0.

Static Friction—Sets how difficult it is for the object to start moving along a surface. The higher this value, the more difficult. Default=0.0.

If something weighs ten pounds and sits on Teflon (a static friction of near 0), it takes almost no force to make it move sideways. On the other hand, if it sits on sandpaper, then the static friction might be very high, on the order of 0.5 to 0.8. A static friction near 1 is very difficult to create in the real world without adhesives or friction material.

Sliding Friction—Sets how difficult it is for the object to keep moving over a surface. The higher this value, the more difficult for the object to keep moving. Default=0.0.

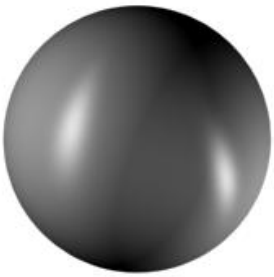
Once two objects begin to slide over one another, static friction disappears and sliding friction takes over. Generally, sliding friction is lower than static friction due to surface tension effects. For example, once steel starts sliding over brass (a value of static friction that might run from 0.05 to 0.2), the sliding friction drops to a significantly lower value, on the order of .01 to 0.1. For some materials, such as specific friction materials like brake linings, sliding friction is just as high as static friction because it is used in conjunction with a nearly frictionless material such as hardened polished steel.

Basic Material Shaders

Anisotropic Shader

Material Editor > Standard material > Shader Basic Parameters rollout > Anisotropic shader > Anisotropic Basic Parameters rollout

The Anisotropic shader creates surfaces with elliptical, "anisotropic" highlights. These highlights are good for modeling hair, glass, or brushed metal. The basic parameters are similar to those for *Blinn or Phong shading (page 2–1480)*, except for the Specular Highlight parameters, and Diffuse Level controls such as those for *Oren-Nayar-Blinn shading (page 2–1482)*.



Anisotropic highlights are elliptical, with differing U and V dimensions.

Anisotropy measures the difference between sizes of the highlight as seen from two perpendicular directions. When anisotropy is 0, there is no difference at all. The highlight is circular, as in Blinn or Phong shading. When anisotropy is 100, the difference is at its maximum. In one direction the highlight is very sharp; in the other direction it is controlled solely by Glossiness.

For more complex highlights, see the *Multi-Layer shader (page 2–1481)*.

See also

Shader Basic Parameters Rollout (page 2–1466)

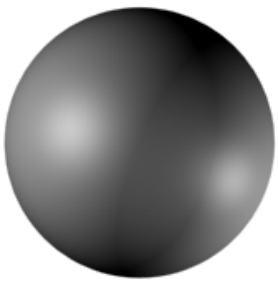
Basic Parameters Rollout (Standard Material) (page 2–1470)

Anisotropic Highlights (page 2–1492)

Blinn Shader

Material Editor > Standard material > Shader Basic Parameters rollout > Blinn shader > Blinn Basic Parameters rollout

Blinn shading is a subtle variation on Phong shading. The most noticeable difference is that highlights appear rounder. In general, you don't need to use the Soften parameter (described in *Blinn, Oren-Nayar-Blinn, and Phong Highlights (page 2–1493)*) as often as you do with Phong shading.



Blinn shading tends to have soft, round highlights.

With Blinn shading, you can obtain highlights produced by light glancing off the surface at low angles. These highlights are lost when you increase the value of Soften using Phong shading.

The Blinn and Phong shaders have the same *basic parameters (page 2–1470)*.

See also

Shader Basic Parameters Rollout (page 2–1466)

Basic Parameters Rollout (Standard Material) (page 2–1470)

Blinn, Oren-Nayar-Blinn, and Phong Highlights (page 2–1493)

Metal Shader

Material Editor > Standard material > Shader Basic Parameters rollout > Metal shader > Metal Basic Parameters rollout

Metal shading provides realistic-looking metallic surfaces and a variety of organic-looking materials.

Metal shading has a distinct curve for specular highlights. Metal surfaces also have glancing highlights. Metal materials calculate their own specular color, which can vary between the material's diffuse color and the color of the light. You can't set a metal material's specular color.



Metal shading has distinctive highlights.

Because there's no separate specular highlight, the two specular highlight spinners behave differently than the spinners for *Blinn and Phong shading (page 2–1480)*. The Specular Level spinner still controls intensity, but the Glossiness spinner affects both the intensity and size of the specular areas.

Tip: When you create a metal material, make sure the *backlight (page 2–1432)* is on in the sample slot.

See also

Shader Basic Parameters Rollout (page 2–1466)

Basic Parameters Rollout (Standard Material) (page 2–1470)

Metal Highlights (page 2–1494)

Multi-Layer Shader

Material Editor > Standard material > Shader Basic Parameters rollout > Multi-Layer shader > Multi-Layer Basic Parameters rollout



Upper left: No highlights

Upper right: Single highlight

Lower middle: Multiple highlights from the multi-layer shader

The Multi-Layer shader is similar to the *Anisotropic shader (page 2–1480)*, but it has a set of two specular highlight controls. The highlights are layered, letting you create complex highlights that are good for highly polished surfaces, special effects, and so on.

Highlights in the Multi-Layer shader can be anisotropic. Anisotropy measures the difference between sizes of the highlight as seen from two perpendicular directions. When anisotropy is

0, there is no difference at all. The highlight is circular, as in Blinn or Phong shading. When anisotropy is 100, the difference is at its maximum. In one direction the highlight is very sharp; in the other direction it is controlled solely by Glossiness.

See also

Shader Basic Parameters Rollout (page 2–1466)

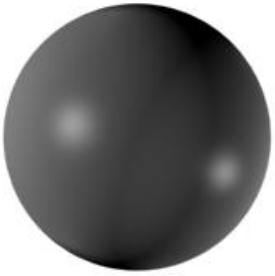
Basic Parameters Rollout (Standard Material) (page 2–1470)

Multi-Layer Highlights (page 2–1495)

Oren-Nayar-Blinn Shader

Material Editor > Standard material > Shader Basic Parameters rollout > Blinn shader > Oren-Nayar-Blinn shader > Oren-Nayar-Blinn Basic Parameters rollout

The Oren-Nayar-Blinn shader is a variant of the *Blinn shader (page 2–1480)*. It contains additional "advanced diffuse" controls, Diffuse Level and Roughness, that you can use to give the material a matte effect. This shader is good for matte surfaces such as fabric, terra-cotta, and so on.



Oren-Nayar-Blinn shading typically has a matte appearance.

See also

Shader Basic Parameters Rollout (page 2–1466)

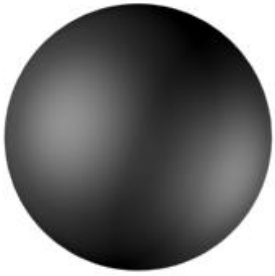
Basic Parameters Rollout (Standard Material) (page 2–1470)

Blinn, Oren-Nayar-Blinn, and Phong Highlights (page 2–1493)

Phong Shader

Material Editor > Standard material > Shader Basic Parameters rollout > Phong shader > Phong Basic Parameters rollout

Phong shading smoothes the edges between faces and renders highlights realistically for shiny, regular surfaces. This shader interpolates intensities across a face based on the averaged face normals of adjacent faces. It calculates the normal for every pixel of the face.



Phong-shaded highlights are typically less regular than Blinn highlights.

Phong shading can accurately render bump, opacity, shininess, specular, and reflection maps.

The Blinn and Phong shaders have the same *basic parameters (page 2–1470)*.

See also

Shader Basic Parameters Rollout (page 2–1466)

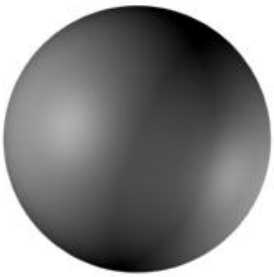
Basic Parameters Rollout (Standard Material) (page 2–1470)

Blinn, Oren-Nayar-Blinn, and Phong Highlights (page 2–1493)

Strauss Shader

Material Editor > Standard material > Shader Basic Parameters rollout > Strauss shader > Strauss Basic Parameters rollout

The Strauss shader is for modeling metallic surfaces. It uses a simpler model and has a simpler interface than the *Metal shader (page 2–1481)*.



Sample of Strauss shading

Note: The Strauss shader's Basic Parameters rollout differs a great deal from the Basic Parameters rollouts for other shaders, and is described in this topic.

See also

Shader Basic Parameters Rollout (page 2–1466)

Procedures

To change the color of a Strauss material:

1. Click the Color swatch.

   The *Color Selector (page 1–161)* is displayed.

   Right-clicking the color swatch also displays the Color Selector.

2. In the Color Selector, change the values of the color.

   As you change color values, the color also changes in the sample in the sample slot.

To reduce a material's opacity:

- Change Opacity to a value less than 100%.

  The material becomes more transparent. An object that is fully transparent (0% Opacity) is nearly invisible except for the light it reflects (the specular highlights).

  To preview transparency in the sample slots, view the sample object against a background. Click the checkered *Background (page 2–1433)* button to the right of the sample slots.

  Transparent materials render more realistically when you turn on 2-Sided in the material's *Shader Basic Parameters (page 2–1466)*.

To increase or decrease the size and intensity of highlights:

- Change the Glossiness value.

  The width of the Highlight curve and the highlights in the preview change. At 0% glossiness, the curve is at its maximum width. At 100% glossiness, the curve is extremely narrow.

  Increasing Glossiness also dims the diffuse color.

To make the material appear more metallic:

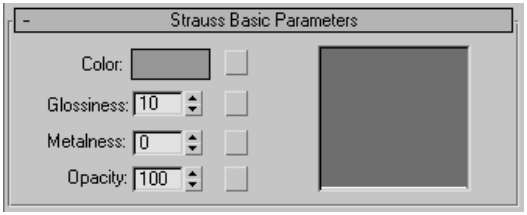
1. Increase the Glossiness value.

   The metallic effect requires visible highlights.

2. Increase the value of Metalness.

   Highlights become more focused, and the (diffuse) color component is dimmed.

Interface



Color—Controls the color of the material. This corresponds to the *diffuse color (page 3–929)* you specify for other kinds of shaders. With the Strauss shader, you control only this color. The shader calculates the ambient and specular color components.

Click the map button to assign a map to the color component. See *Diffuse Mapping (page 2–1498)*. This button is a shortcut: you can also assign color mapping in the *Maps (page 2–1474)* rollout.

Glossiness—Affects the size and intensity of the specular highlight. As you increase the value, the highlight gets smaller and the material appears shinier. Default=25.

Glossiness also controls the strength of reflection maps assigned to a Strauss material.

Click the map button to assign a map to the glossiness component. See *Glossiness Mapping (page 2–1502)*. This button is a shortcut: you can also assign glossiness mapping in the *Maps (page 2–1474)* rollout.

Metalness—Changes the metallic appearance of a material. Increasing the Metalness value increases the metallic appearance, with glancing as well as primary highlights. Because a metallic appearance principally depends on highlights, the Metalness value has little effect unless you also increase the Glossiness value. Default=0.

Tip: When you create a metal material, make sure the *backlight (page 2–1432)* is on in the sample slot.

Opacity—Sets the opacity/transparency of the material as a percentage. The effect is best previewed against a *pattern background (page 2–1433)* in the sample slot. You can control *opacity falloff (page 3–984)* in the Extended Parameters. Default=100.

Click the map button to assign a map to the opacity component. See *Opacity Mapping (page 2–1503)*. This button is a shortcut: you can also assign opacity mapping in the *Maps rollout (page 2–1474)*.

Highlight graph—This curve shows the effect of adjusting the value of Glossiness. As you decrease Glossiness, the curve grows shorter; as you increase it, the curve grows taller.

Translucent Shader

Material Editor > Standard material > Shader Basic Parameters rollout > Translucent shader > Translucent Basic Parameters rollout



The bust on the right has been made translucent.

Translucent shading is similar to Blinn shading, but it also lets you specify translucency. A translucent object allows light to pass through, and also scatters light within the object. You can use translucency to simulate frosted and etched glass.

Translucency is inherently a two-sided effect: with the translucent shader, backface illumination appears on front faces. To generate translucency,

both sides of the material receive diffuse light, though only one side is visible in renderings and shaded viewports unless you turn on 2-Sided (in the Shader Basic Parameters rollout).

If you use *radiosity (page 3–51)*, it will process light transmitted by translucency. The accuracy of this depends on the mesh: the more subdivided the faces are, the more accurate the solution will be (at a cost of processing time).

For specular highlights, you have a choice: to model materials like translucent plastic, you can choose to have highlights on both sides; to model materials like frosted glass, which is reflective on one side only, you can choose to have highlights on only one side. This is controlled by the Backside Specular toggle in the translucent highlight controls.

Tip: To simulate frosted glass, a fine-grained bump map can also help.

The translucent effect appears only in renderings. It does not appear in shaded viewports.

Note: The translucent shader does *not* simulate the scattering of light within the object. Because of this, it is better at simulating thin objects such as glass or paper, than at thick objects. For thicker objects, the light passing through might saturate excessively. To avoid this, try reducing the HSV Value of the material's Translucent Color.

Translucent materials also capture shadows cast on the backfaces of the material. However, because the translucent shader doesn't scatter light, for thicker objects the effect is not an accurate simulation of real-world translucency.

Warning: Do not use shadow maps with the translucent shader. Shadow maps result in artifacts at the edge of translucent objects.

See also

Shader Basic Parameters Rollout (page 2–1466)

Basic Parameters Rollout (Standard Material) (page 2–1470)

Translucency Setting (page 2–1491)

Translucent Highlights (page 2–1496)

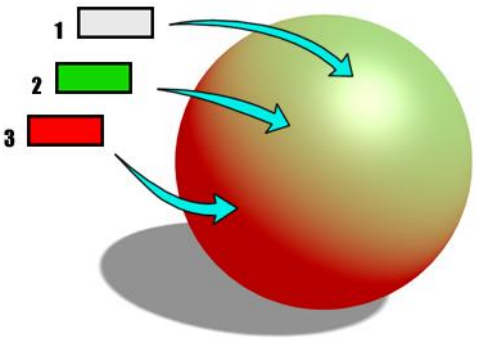
Basic Parameters for Standard Materials

Color Controls

Material Editor > Standard material > Anisotropic, Blinn, Metal, Multi-Layer, Oren-Nayar-Blinn, or Phong Basic Parameters rollout > First group in the rollout (unlabeled)

Color controls set the colors for different color components. You can set the color by clicking the color swatch to display the *Color Selector (page 1–161)*.

Note: The Metal shader does not have a Specular component, because it generates the specular color automatically. The Multi-Layer shader can have two different Specular color components, so for this material the Specular color swatches are found in the Specular Highlights group. The Strauss shader has only a single color component, which corresponds to Diffuse.



1. *Specular color*
2. *Diffuse color*
3. *Ambient color*

Procedures

To change a color component:

1. Click the color swatch next to the color component you want to change.

   The Material Editor displays a *Color Selector (page 1–161)*.

   Right-clicking the color swatch also displays the Color Selector.

2. Use the Color Selector to change the values of the color component.

   As you change color values, the color component also changes in the sample slot.

To copy one color component to another:

1. Drag the color swatch of the color you want to copy to the color swatch of the other color component.

   A *Copy or Swap Colors dialog (page 2–1452)* is displayed.

2. Click Copy to replace the second color swatch with the color you dragged. Click Swap to swap the two color components.

To lock two color components:

1.  Click the lock button between Ambient and Diffuse or between Diffuse and Specular.

   The Material Editor displays an alert that asks whether you want to lock the two color components.

2. Click Yes.

   The color above replaces the color below. In other words, Ambient replaces Diffuse and Diffuse replaces Specular.

   If two colors are locked, and you lock the other two, all three component colors are replaced by the active color.

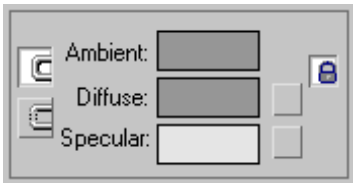
   While two colors are locked, adjustments to one color component affect the other as well.

To unlock two color components:

-  Click the lock button to turn it off.

  The two colors remain the same until you change one of them, or both.

Interface



Ambient—Controls the *ambient color (page 3–908)*. The ambient color is the color in shadow (indirect light).

Diffuse—Controls the *diffuse color (page 3–929)*. The diffuse color is the color in direct light.

Specular—Controls the *specular color (page 3–1014)*. The specular color is the color of the highlight on a shiny object. You can control the size and shape of highlights in the Specular Highlights group, described below.

Copying and Locking Color Components

For convenience in changing color components, the Material Editor lets you copy one color component to another by dragging, and to lock two color components together with the lock buttons to the left of the Ambient and Diffuse, and Diffuse and Specular color swatches.

When you drag and drop a color swatch, the *Copy or Swap Colors dialog (page 2–1452)* asks if you want to copy the color or swap the two colors.

In general, materials with two identical color components do not look realistic, and except for materials that are close to solid black, you should avoid using copied or locked color components in materials you use in a scene. Color copying and locking are best used as conveniences when you design a new basic material.

Shortcut Map Buttons

The small buttons to the right of the color swatches access the *Material/Map Browser (page 2–1447)*, where you select a map for that component. These buttons are shortcuts: you can also use the corresponding buttons in the *Maps rollout (page 2–1474)* If you have assigned a map to one of these color components, the button displays the letter M. An uppercase M means that the corresponding map is assigned and active. A lowercase m means that the map is assigned and inactive (turned off).


The lock button to the right of the Diffuse map button locks Ambient mapping to Diffuse mapping. It is on by default. Usually it makes sense to use the same map for the ambient and diffuse components. To use different maps for ambient and diffuse, turn off the lock button. A map shortcut button for Ambient appears.

Self-Illumination Setting

Material Editor > Standard material > Anisotropic, Blinn, Metal, Multi-Layer, Oren-Nayar-Blinn, or Phong Basic Parameters rollout > Self-Illumination group

These controls make the material *self illuminated (page 3–1009)*. Self-illumination creates the illusion of incandescence by replacing shadows on the surface with the diffuse color. As you increase self-illumination, the self-illumination color takes over from the ambient color. At a setting of 100, the material shows no shaded areas, although it can show specular highlights.

The self-illumination color appears in viewports. (In releases prior to 3ds Max **5,**, viewports showed the self-illumination value but not the color.)

Note: The *Strauss shader (page 2–1483)* does not have self-illumination.

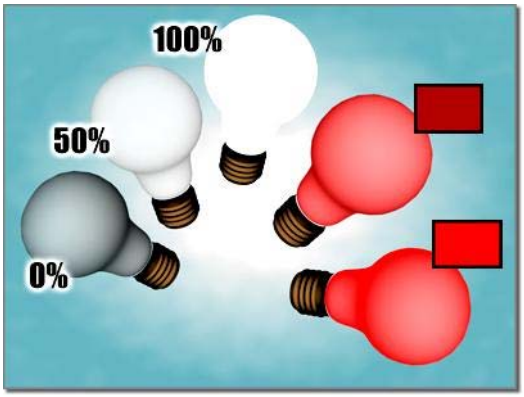
There are two ways to specify self-illumination. You can turn on the check box and use a self-illumination color, or turn off the check box and use a monochrome spinner, which is comparable to using a gray scale self-illumination color.

Self-illuminated materials do not show shadows cast onto them, and they are unaffected by the lights in the scene. The brightness (Value in the *HSV color description (page 3–1001)*) remains the same regardless of the scene's lighting.

To make a visible light source in a scene, you can combine a geometric object with a light object, and give the geometric object a self-illuminating surface. For example, you could create a lofted light bulb shape, assign it a self-illuminating white or yellowish material, and place an omni light in the same location.

To make a material both self-illuminating and transparent, use the Additive transparency type in

combination with self-illumination. See *Extended Parameters (page 2–1471)*.



An object that is self-illuminated using a percentage value and a color

Procedure

To make a material self-illuminating:

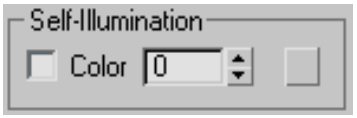
1. Click the color swatch in the Self-Illumination group.
2. In the *Color Selector (page 1–161)*, choose a color for self-illumination.
3. Use the color's Value parameter (in the HSV model) to increase or decrease the amount of *self-illumination (page 3–1009)*.

   You can also set self-illumination with a monochrome spinner. To do so, turn off the self-illumination check box and adjust the spinner.

   The self-illumination color is mixed with the material's diffuse color. The closer to black the self-illumination color, the more diffuse color is used.

   As self-illumination increases, the sample object appears flatter and more luminous.

Interface



Color check box—When on, the material uses a special self-illumination color. When off, the material uses the diffuse color for self-illumination, and displays a spinner to control the self-illumination amount. Default=off.

Color swatch—When Color is on, the color swatch shows the self-illumination color. To change the color, click the swatch and then use the *Color Selector (page 1–161)*. Adjusting the Value (in the color's *HSV description (page 3–1001)*) adjusts the amount of self-illumination. The greater the Value, the more the self-illumination color dominates both the ambient and diffuse color components.

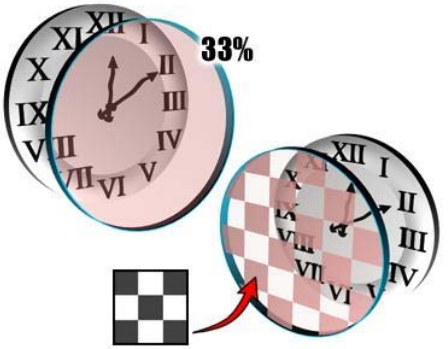
Mono spinner—When Color is off, the diffuse component is used as the self-illumination color, and this spinner lets you adjust the amount of self-illumination. At 0, there is no self-illumination. At 100, the diffuse color takes over from the ambient color.

Click the map button to assign a map to the self-illumination component. See *Self-Illumination Mapping (page 2–1502)*. This button is a shortcut: you can also assign self-illumination mapping in the *Maps rollout (page 2–1474)*.

Opacity

Material Editor > Standard material > Anisotropic, Blinn, Metal, Multi-Layer, Oren-Nayar-Blinn, or Phong Basic Parameters rollout > Opacity group (unlabeled)

Opacity controls whether a material is opaque, transparent, or translucent. (A more physically accurate way to generate translucency is to use the *Translucent shader (page 2–1484)*.)



Controlling opacity using the Opacity setting (left) or an opacity map (right).

Procedure

To reduce a material's opacity:

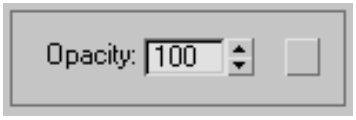
- Change Opacity to a value less than 100%.

  The material becomes more transparent. An object that is fully transparent (0% Opacity) is nearly invisible except for the light it reflects (the specular highlights).

  To preview transparency in the sample slots, view the sample object against a background. Click the checkered *Background button (page 2–1433)* to the right of the sample slots.

  Transparent materials render more realistically when you turn on 2-Sided in the material's *Shader Basic Parameters (page 2–1466)*.

Interface



Opacity—Sets the opacity/transparency of the material as a percentage. The effect is best previewed against a *pattern background (page 2–1433)* in the sample slot. You can control *opacity falloff (page 3–984)* in the Extended Parameters.

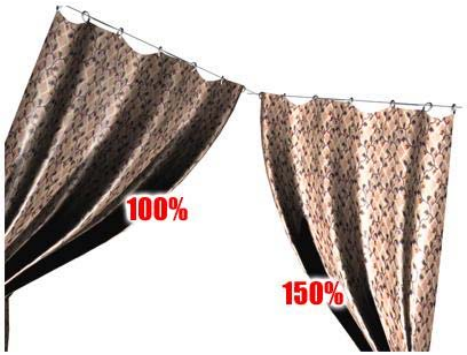
Click the map button to assign a map to the opacity component. See *Opacity Mapping (page 2–1503)*. This button is a shortcut: you can also assign opacity mapping in the *Maps rollout (page 2–1474)*.

Diffuse Level

Material Editor > Standard material > Anisotropic, Multi-Layer, or Oren-Nayar-Blinn Basic Parameters rollout > Diffuse Level group (unlabeled) or Advanced Diffuse group

Diffuse Level controls the brightness of the material's diffuse component.

Note: The Blinn, Metal, Phong, and Strauss shaders do not have Diffuse Level control.



Adjusting diffuse level

Procedure

To adjust the diffuse level:

- Change the value of Diffuse Level.

  The material grows lighter or darker. Lowering the Diffuse Level dims the material's diffuse color without affecting the specular highlight. Diffuse Level is intended primarily so you can create a *map (page 2–1499)* that makes portions of the material very dark.

Interface



Diffuse Level—Increasing this value increases diffuse brightness, and decreasing it reduces diffuse brightness without affecting the specular highlight. You can increase the diffuse level over and above the diffuse color's Value (in its *HSV description (page 3–1001)*). This parameter can range from 0 to 400. Default=100.

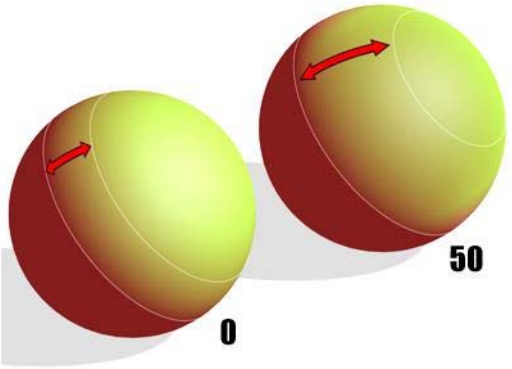
Click the map button to assign a map to the diffuse level parameter. See *Diffuse Level Mapping (page 2–1499)*. This button is a shortcut: you can also assign diffuse level mapping in the *Maps rollout (page 2–1474)*.

Roughness

Material Editor > Standard material > Multi-Layer or Oren-Nayar-Blinn Basic Parameters rollout > Advanced Diffuse group (unlabeled for Multi-Layer)

Roughness controls how quickly the diffuse component blends into the ambient component.

Note: The Roughness parameter is available only with the *Oren-Nayar-Blinn (page 2–1482)* and *Multi-Level (page 2–1481)* shaders, and with the *Arch & Design material (mental ray) (page 2–1549)*.



Increasing the blending area between diffuse and ambient with Roughness

Procedure

To adjust the roughness:

- Change the value of Roughness.

  Increasing roughness makes the material have a flatter, more matte appearance.

Interface



Roughness—As you increase this value, the matte appearance of the material increases. It also grows darker and appears more flat. At 0, the roughness is the same as it is with *Blinn shading (page 2–1480)*. Default=50.

Click the map button to assign a map to Roughness. This button is a shortcut: you can also assign Diffuse Roughness mapping in the *Maps rollout (page 2–1474)*. See *Diffuse Roughness Mapping (page 2–1500)*.

Translucency Setting

Material Editor > Standard material > Shader Basic Parameters rollout > Translucent shader > Translucent Basic Parameters rollout > Translucency group

The translucency controls are available for the *Translucent shader (page 2–1484)*.

Warning: Do not use shadow maps with the translucent shader. Shadow maps result in artifacts at the edge of translucent objects.

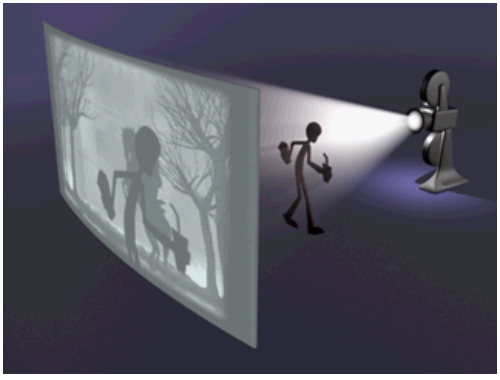
Procedure

To make a material translucent:

- Increase the *HSV Value (V) (page 3–1001)* of the Translucent Color.

  As the Value increases, the material becomes more translucent. The Hue of the Translucent Color tints the light that is scattered within the material.

  Translucent materials render more realistically when you turn on 2-Sided in the material's *Shader Basic Parameters (page 2–1466)*. The translucent effect does not appear in shaded viewports.



Example of translucency

Interface



Translucent Clr (Color)—Specifies a translucency color. This is the color of light that is scattered within the material. It does not need to be the same as the filter color, which is light transmitted by the material. The two color values are multiplied. Click the color swatch to change the translucent color. Click the button to assign a map to the translucent color component.

Filter Color—Specifies a *filter color (page 3–939)* that is multiplied by the translucent color. Click the color swatch to change the filter color. Click the button to assign a map to the filter color component.

The filter, or transmissive color, is the color transmitted through transparent or semi-transparent materials such as glass. You can use the filter color with volumetric lighting to create effects such as colored light through a stained-glass window. *Ray-traced shadows (page 3–1000)* cast by transparent objects are tinted with the filter color.

Opacity—Sets the *opacity/transparency (page 2–1488)* of the material as a percentage. The effect is best previewed against a *pattern background (page 2–1433)* in the sample slot.

Click the map button to assign a map to the opacity component. See *Opacity Mapping (page 2–1503)*. This button is a shortcut: you can also assign opacity mapping in the *Maps rollout (page 2–1474)*.

Specular Highlight Controls

Anisotropic Highlights

Material Editor > Standard material > Anisotropic Basic Parameters rollout > Specular Highlight group

Material Editor > Raytrace material > Raytrace Basic Parameters rollout > Shading: Anisotropic > Specular Highlight group

Anisotropic highlights are good for modeling hair, glass, or brushed metal.

Note: For the *Raytrace material (page 2–1512)*, the Specular Color component appears in the Specular Highlight group. Also, highlight controls that don't pertain to the current shader are labeled "N / A."

Procedures

To increase or decrease the size of a highlight:

- Change the Glossiness value.

  The width of the Highlight curves and the highlights in the preview change. At 0% glossiness, the curves are at their maximum width. At 100% glossiness, both curves are extremely narrow.

To increase or decrease the strength of a highlight:

- Change the value of Specular Level.

  The intensity of the Highlight curves and the highlights in the preview change. At 0% specular level, there is no highlight. At 100% specular level, the curves are at their maximum height with no overloading. At values greater than 100%, the curves are overloaded: they grow wider, and a wider area is at the maximum highlight intensity.

  The shape of the Highlight curves affects the blending between the specular and diffuse color regions of the material. The steeper the curve, the less blending there is and the sharper the edge of the specular highlight.

To adjust the shape (anisotropy) of the highlight:

- Change the value of Anisotropy.

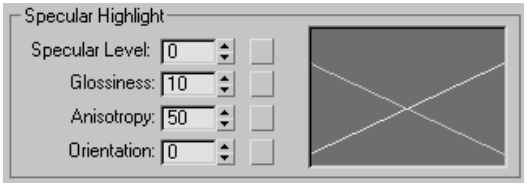
  The width of the white highlight curve and the highlights in the preview change. At 0% anisotropy, both highlight curves are the same and the highlight is circular, as in Blinn and Phong shading. At 100% anisotropy, the white highlight curve and the highlights are extremely narrow.

To adjust the orientation of the highlight:

- Change the value of Orientation.

  Highlights in the preview show the change in orientation. The display of the highlight curve does not change.

Interface



Specular Level—Affects the intensity of the specular highlight. As you increase the value, the highlight grows brighter. Default=5.

Click the map button to assign a map to the specular level component. See *Specular Level Mapping (page 2–1501)*. This button is a shortcut: you can also assign specular level mapping in the *Maps rollout (page 2–1474)*.

Glossiness—Affects the size of the specular highlight. As you increase the value, the highlight gets smaller and the material appears shinier. Default=25.

Click the map button to assign a map to the glossiness component. See *Glossiness Mapping*

(page 2–1502). This button is a shortcut: you can also assign specular level mapping in the *Maps rollout (page 2–1474)*.

Anisotropy—Controls the anisotropy, or shape, of the highlight. At 0, the highlight is round. At 100, the highlight is extremely narrow. One axis of the Highlight graph changes to show changes in this parameter. Default=50.

Orientation—Changes the orientation of the highlight. The sample slot shows changes in orientation. This is a value in degrees that can range from 0 to 9,999. Default=0.

Highlight graph—These two intersecting curves show the effect of adjusting the values of Specular Level, Glossiness, and Anisotropy. As you decrease Glossiness, the curves grow wider; as you increase Specular Level, the curves grow taller. As you adjust Anisotropy, the white curve changes to show how wide or narrow the highlight is.

Blinn, Oren-Nayar-Blinn, and Phong Highlights

Material Editor > Standard material > Blinn, Oren-Nayar-Blinn, or Phong Basic Parameters rollout > Specular Highlight group

Material Editor > Raytrace material > Raytrace Basic Parameters rollout > Shading: Blinn, Oren-Nayar-Blinn, or Phong > Specular Highlight group

The *Blinn (page 2–1480)*, *Oren-Nayar-Blinn (page 2–1482)*, and *Phong (page 2–1482)* shaders all have circular highlights and share the same highlight controls. Blinn and Oren-Nayar-Blinn highlights are somewhat softer and rounder than Phong highlights.

Note: For the *Raytrace material (page 2–1512)*, the Specular Color component appears in the Specular Highlight group. Also, highlight controls that don't pertain to the current shader are labeled "N/A."

Procedures

To increase or decrease the strength of a highlight:

- Change the value of Specular Level.

  The intensity of the Highlight curve and the highlight in the preview change. At 0% specular level, there is no highlight. At 100% specular level, the curve is at its maximum height with no overloading. At values greater than 100%, the curve is overloaded: it grows wider, and a wider area is at the maximum highlight intensity.

  The shape of the Highlight curve affects the blending between the specular and diffuse color regions of the material. The steeper the curve, the less blending there is and the sharper the edge of the specular highlight.

To increase or decrease the size of a highlight:

- Change the Glossiness value.

  The width of the Highlight curve and the highlight in the preview change. At 0% glossiness, the curve is at its maximum width. At 100% glossiness, the curve is extremely narrow.

Interface



Specular Level—Affects the intensity of the specular highlight. As you increase the value, the highlight grows brighter. Default=5.

Click the map button to assign a map to the specular level component. See *Specular Level Mapping (page 2–1502)*. This button is a shortcut: you can also assign specular level mapping in the *Maps rollout (page 2–1474)*.

Glossiness—Affects the size of the specular highlight. As you increase the value, the highlight

gets smaller and the material appears shinier. Default=25.

Click the map button to assign a map to the glossiness component. See *Glossiness Mapping (page 2–1502)*. This button is a shortcut: you can also assign specular level mapping in the *Maps rollout (page 2–1474)*.

Soften—Softens the effect of specular highlights, especially those formed by glancing light. When Specular Level is high and Glossiness is low, you can get harsh backlights on surfaces. Increase the value of Soften to mitigate this effect. At 0, there is no softening. At 1.0, the maximum amount of softening is applied. Default=0.1.

Note: The Soften control was a check box in releases prior to 3ds Max **2**. When you load a material created in an earlier version of 3ds Max, if Soften was originally off, the new Soften value is 0.0. If Soften was originally on, the new Soften value is 0.6.

Highlight graph—This curve shows the effect of adjusting the values of Specular Level and Glossiness. As you decrease Glossiness, the curve grows wider; as you increase Specular Level, the curve grows taller.

Metal Highlights

Material Editor > Standard material > Metal Basic Parameters rollout > Specular Highlight group

Material Editor > Raytrace material > Raytrace Basic Parameters rollout > Shading: Metal > Specular Highlight group

Metal-shaded materials generate their own specular color. Also, the highlight curve for the Metal shader differs in shape from the curve for *Blinn Oren-Nayar-Blinn, and Phong highlights (page 2–1493)*.

Note: For the *Raytrace material (page 2–1512)*, the Specular Color component appears in the Specular Highlight group. Also, highlight controls that don't pertain to the current shader are labeled "N/A."

Procedures

To increase or decrease the size and intensity of a highlight:

- Change the Glossiness value.

  The width of the Highlight curve and the highlights in the preview change. At 0% glossiness, the curve is at its maximum width. At 100% glossiness, the curve is extremely narrow.

To increase or decrease the strength of a highlight:

- Change the value of Specular Level.

  The intensity of the Highlight curve and the highlight in the preview change. At 0% specular level, there is no highlight. At 100% specular level, the curve is at its maximum height with no overloading. At values greater than 100%, the curve is overloaded: it grows wider, and a wider area is at the maximum highlight intensity.

  Increasing the Specular Level also dims the diffuse color.

  The shape of the Highlight curve affects the blending between the specular and diffuse color regions of the material. The steeper the curve, the less blending there is and the sharper the edge of specular and glancing highlights.

Interface



Specular Level—Affects the intensity of the specular highlight. As you increase the value, the highlight grows brighter and the diffuse color grows dimmer. Default=5.

Click the map button to assign a map to the specular level component. See *Specular Level Mapping (page 2–1501)*. This button is a shortcut: you can also assign specular level mapping in the *Maps (page 2–1474)* rollout.

Glossiness—Affects the size of the specular highlight. As you increase the value, the highlight curve grows narrower and the highlight gets smaller. Default=25.

Click the map button to assign a map to the glossiness component. See *Glossiness Mapping (page 2–1502)*. This button is a shortcut: you can also assign specular level mapping in the *Maps (page 2–1474)* rollout.

Highlight graph—This curve shows the effect of adjusting the values of Specular Level and Glossiness. As you decrease Glossiness, the curve grows wider; as you increase Specular Level, the curve grows taller.

Multi-Layer Highlights

Material Editor > Standard material > Multi-Layer Basic Parameters rollout > First Specular Layer/Second Specular Layer groups

Multi-layer highlights consist of two layers, each of them anisotropic. The highlights are transparent to each other. Where they overlap, the Multi-Layer shader blends their colors.

Procedures

To increase or decrease the size of a highlight (specular) layer:

- Change the Glossiness value.

  The width of the Highlight curves and the highlights in the preview change. At 0% glossiness, the curves are at their maximum width. At 100% glossiness, both curves are extremely narrow.

To increase or decrease the strength of a highlight (specular) layer:

- Change the value of Specular Level.

  The intensity of the Highlight curves and the highlights in the preview change. At 0% specular level, there is no highlight. At 100% specular level, the curves are at their maximum height with no overloading. At values greater than 100%, the curves are overloaded: they grow wider, and a wider area is at the maximum highlight intensity.

  The shape of the Highlight curves affects the blending between the specular and diffuse color regions of the material. The steeper the curve, the less blending there is and the sharper the edge of the specular highlight.

To adjust the shape (anisotropy) of a highlight (specular) layer:

- Change the value of Anisotropy.

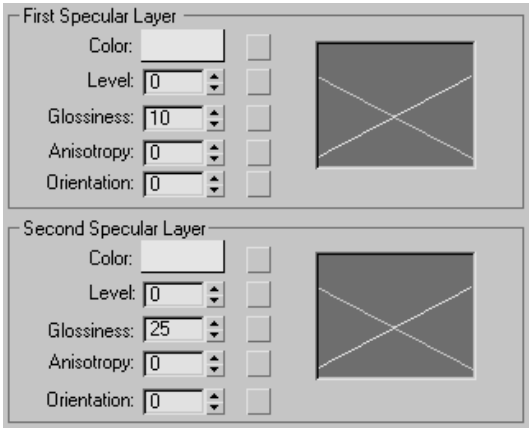
  The width of the white highlight curve and the highlights in the preview change. At 0% anisotropy, both highlight curves are the same and the highlight is circular, as in Blinn and Phong shading. At 100% anisotropy, the white highlight curve and the highlights are extremely narrow.

To adjust the orientation of a highlight (specular) layer:

- Change the value of Orientation.

  Highlights in the preview show the change in orientation. The display of the highlight curve does not change.

Interface



The First Specular Layer and Second Specular Layer groups have identical controls, which can have different settings.

Color—Controls the *specular color (page 3–1014)* of this highlight. The specular color is the color of the highlight on a shiny surface.

Level—Affects the intensity of this specular highlight. As you increase the value, the highlight grows brighter. Default: First layer=5, Second layer=0.

Click the map button to assign a map to the specular level component. See *Specular Level Mapping (page 2–1501)*. This button is a shortcut: you can also assign specular level mapping in the *Maps rollout (page 2–1474)*.

Glossiness—Affects the size of this specular highlight. As you increase the value, the highlight gets smaller and the material appears shinier. Default=25.

Click the map button to assign a map to the glossiness component. See *Glossiness Mapping (page 2–1502)*. This button is a shortcut: you can also assign specular level mapping in the *Maps rollout (page 2–1474)*.

Anisotropy—Controls the anisotropy, or shape, of this highlight. At 0, the highlight is round. At 100, the highlight is extremely narrow. One axis of the Highlight graph changes to show changes in this parameter. Default=0.

Orientation—Changes the orientation of this highlight. The sample slot shows changes in orientation. This is a value in degrees that can range from 0 to 9,999. Default=0.

Highlight graph—These two intersecting curves show the effect of adjusting the values of Level, Glossiness, and Anisotropy. As you decrease Glossiness, the curves grow wider; as you increase Specular Level, the curves grow taller. As you adjust Anisotropy, the white curve changes to show how wide or narrow the highlight is.

Translucent Highlights

Material Editor > Standard material > Translucent Basic Parameters rollout > Specular Highlight group

Like the Blinn shader, the Translucent shader has circular highlights.

Procedures

To increase or decrease the strength of a highlight:

- Change the value of Specular Level.

  The intensity of the Highlight curve and the highlight in the preview change. At 0% specular level, there is no highlight. At values greater than 100%, the curve is overloaded: it grows wider, and a wider area is at the maximum highlight intensity. At 100% specular level, the curve is at its maximum height with no overloading.

  The shape of the Highlight curve affects the blending between the specular and diffuse color regions of the material. The steeper the curve,

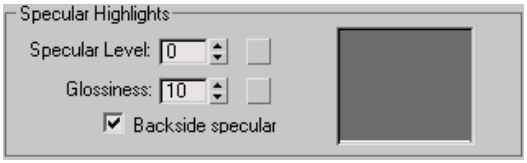
the less blending there is and the sharper the edge of the specular highlight.

To increase or decrease the size of a highlight:

- Change the Glossiness value.

  The width of the Highlight curve and the highlight in the preview change. At 0% glossiness, the curve is at its maximum width. At 100% glossiness, the curve is extremely narrow.

Interface



Specular Level—Affects the intensity of the specular highlight. As you increase the value, the highlight grows brighter. Default=0.

Click the map button next to the spinner to assign a map to the specular level component. See *Specular Level Mapping (page 2–1501)*. This button is a shortcut: you can also assign specular level mapping in the *Maps rollout (page 2–1474)*.

Glossiness—Affects the size of the specular highlight. As you increase the value, the highlight gets smaller and the material appears shinier. Default=10.

Click the map button next to the spinner to assign a map to the glossiness component. See *Glossiness Mapping (page 2–1502)*. This button is a shortcut: you can also assign specular level mapping in the *Maps rollout (page 2–1474)*.

Backside specular—When on, both sides of the material receive a specular highlight. When off, only the front side of the material receives a highlight. Default=on.

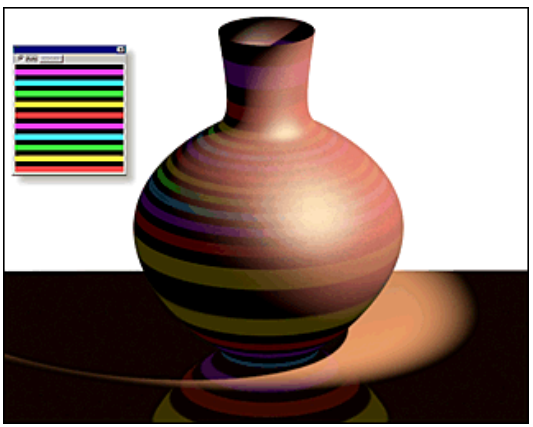
Leave Backside Specular on to model materials like translucent plastic. Turn it off to model materials like frosted glass.

Tip: When Backside Specular is turned off, the front side is always the one that receives a specular highlight. You can change this by reversing the normals of surfaces that have the translucent-shaded material.

Mapping Standard Material Components

Ambient Color Mapping

Material Editor > Standard material > Maps rollout > Ambient button



Mapping the ambient color

You can select a bitmap file or *procedural map (page 3–997)* to map an image to the material's *ambient color (page 3–908)*. The image is painted on the shaded parts of the object.

By default, diffuse mapping maps the ambient component as well, so you seldom need to use a different map for diffuse and ambient components.

If you do want to apply a separate ambient map, first click to turn off the lock button to the right of the long Map buttons in the *Maps rollout (page 2–1474)*. This unlocks ambient and diffuse mapping. The Map button for ambient color becomes available. You can then click the ambient button to select a map.

Note: Ambient color mapping is not visible in viewports or renderings unless the level of ambient light is greater than default value of black. Choose Rendering > Environment, then adjust the level of ambient light using the *Environment dialog (page 3–272)*.

Procedure

To map the ambient color:

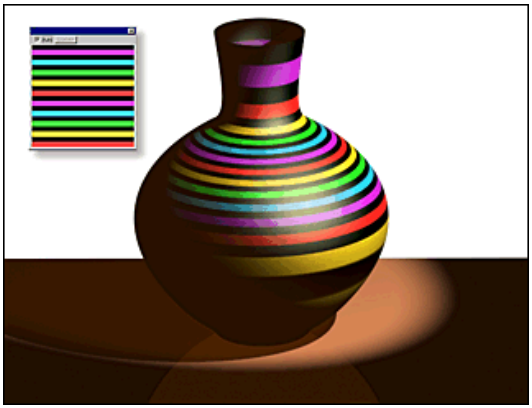
1.  Make sure the ambient and diffuse components have their maps unlocked. Click to turn off the lock button. The Map button for ambient color becomes available.
2. Click the Map button for Ambient color.

   The *Material/Map Browser (page 2–1447)* is displayed.
3. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.
4. Use the map controls to set up the map.

Diffuse Color Mapping

Material Editor > Standard material > Maps rollout > Diffuse button (or Color button for the Strauss shader)



Applying a texture by mapping the diffuse color

You can select a bitmap file or *procedural map (page 3–997)* to assign a pattern or texture to a material's *diffuse color (page 3–929)*. The colors of the map replace the material's diffuse color component. This is the most common kind of mapping.

Mapping the diffuse color is like painting an image on the surface of the object. For example, if you want a wall to be made out of brick, you can choose a map with an image of bricks, such as *Bricks (page 2–1658)*.

By default, diffuse mapping applies the same map to the *ambient color (page 3–908)* as well. You seldom need to use a different map for diffuse and ambient components.

It isn't strictly necessary to lock the ambient and diffuse maps. By turning the lock off and using a different map for each component, you can obtain interesting blend effects. But in general, the purpose of diffuse mapping is to simulate a single surface that is more complex than a basic material, and for this purpose the lock should be on.

Procedure

To map the diffuse color:

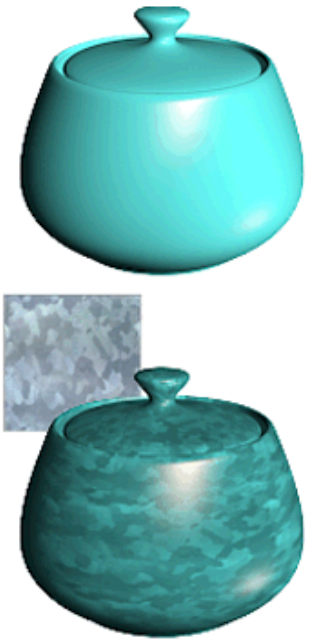
1. Make sure the ambient and diffuse components have their maps locked. This button is to the right of the map shortcut buttons for Ambient and Diffuse on the Basic Parameters rollout, and on the Maps rollout in a similar position. It is on by default. When it is on, the map button for the ambient color component is unavailable.
2. Click the Map button for Diffuse color.

   The *Material/Map Browser (page 2–1447)* is displayed.
3. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.
4. Use the map controls to set up the map.

Diffuse Level Mapping

Material Editor > Standard material > Shader Basic Parameters rollout > Anisotropic, Oren-Nayar-Blinn, or Multi-Level shader > Maps rollout > Diffuse Level button



Mapping diffuse level

Top: No mapping

Bottom: Mapping diffuse level with a bitmap

You can select a bitmap file or *procedural map (page 3–997)* to control the Diffuse Level parameter. White pixels in the map leave the diffuse level unchanged. Black pixels reduce the diffuse level to 0. Intermediate values adjust the diffuse level accordingly.

The diffuse level parameter is available with the *Anisotropic (page 2–1480)*, *Oren-Nayar-Blinn (page 2–1482)*, and *Multi-Level (page 2–1481)* shaders.

Reducing the Amount of the diffuse level map reduces the map's effect, and increases the effect of the Diffuse Level value on the Basic Parameters

rollout. When the Amount is 0 percent, the map isn't used at all.

Procedure

To map the diffuse level value:

1. Click the Map button for Diffuse Level.

   The *Material/Map Browser (page 2–1447)* is displayed.

2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Diffuse Roughness Mapping

Material Editor > Standard material > Shader Basic Parameters rollout > Oren-Nayar-Blinn or Multi-Level shader > Maps rollout > Diff. Roughness button

You can select a bitmap file or *procedural map (page 3–997)* to control the Roughness parameter on the Basic Parameters rollout. White pixels in the map increase roughness. Black pixels reduce roughness to 0. Intermediate values adjust roughness accordingly.

Note: The Roughness parameter is available only with the *Oren-Nayar-Blinn (page 2–1482)* and *Multi-Level (page 2–1481)* shaders, and with the *Arch & Design material (mental ray) (page 2–1549)*.

Reducing the Amount of the Diffuse Roughness map reduces the map's effect, and increases the effect of the Roughness value on the Basic Parameters rollout. When the Amount is 0 percent, the map isn't used at all.

Procedure

To map the roughness value:

1. Click the Map button for Diffuse Roughness (Diff. Roughness).

   The *Material/Map Browser (page 2–1447)* is displayed.

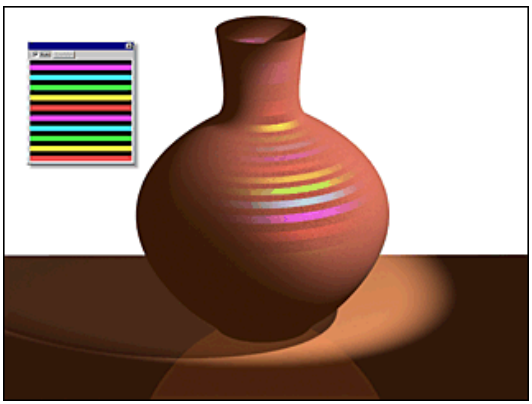
2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Specular Color Mapping

Material Editor > Standard material > Maps rollout > Specular button



Mapping the specular highlight

You can select a bitmap file or *procedural map (page 3–997)* to apply an image to the material's *specular color component (page 3–1014)*. The map's image appears only in the specular highlight areas.

When the amount spinner is at 100, all specular color is provided by the map.

Specular mapping is used primarily for special effects such as placing an image in a reflection. The important thing to remember is that, unlike

Specular Level or Glossiness mapping, which alter the *intensity and location* of specular highlights, specular mapping alters the *color* of specular highlights.

Procedure

To map the specular color:

1. Click the Map button for Specular color.

   The *Material/Map Browser (page 2–1447)* is displayed.

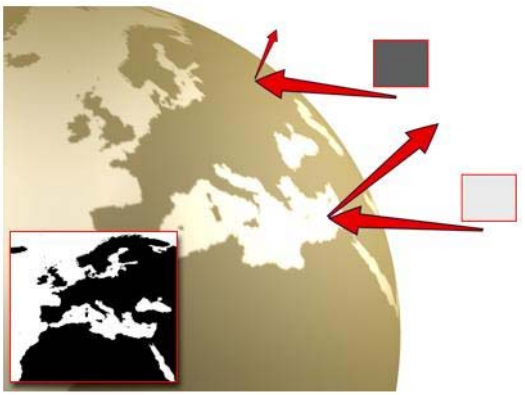
2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Specular Level Mapping

Material Editor > Standard material > Maps rollout > Specular Level button



Mapping the specular level: the sea reflects more than the land.

You can select a bitmap file or *procedural map (page 3–997)* to alter the intensity of specular highlights, based on the intensity of the bitmap. White pixels in the map produce full specular highlights. Black pixels remove the specular highlights completely, and intermediate values reduce the specular highlights accordingly.

Mapping the specular level component is different from mapping specular color. Mapping the specular level alters the *intensity* of highlights, while specular mapping alters the *color* of highlights.

Specular level mapping usually works best when you assign the same map to both Specular Level and Glossiness. (In the *Maps (page 2–1474)* rollout, you can do this by dragging from one map button to another.)

Procedure

To map the specular level value:

1. Click the Map button for the Specular Level value.

   The *Material/Map Browser (page 2–1447)* is displayed.

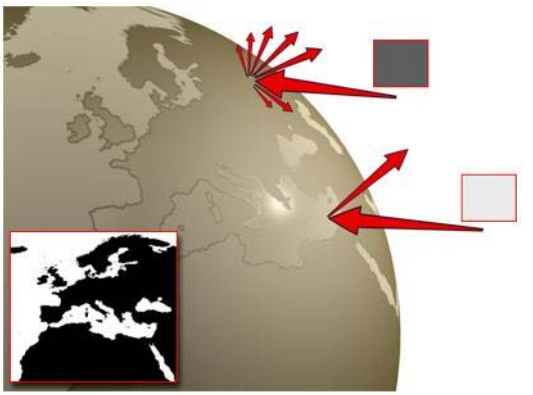
2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Glossiness Mapping

Material Editor > Standard material > Maps rollout > Glossiness button



An object with glossiness mapping. The sea appears more reflective than the land.

You can select a bitmap file or *procedural map (page 3–997)* that affects where specular highlights appear. A map assigned to glossiness determines which areas of the whole surface are more glossy and which areas are less glossy, depending on the intensity of colors in the map. Black pixels in the map produce full glossiness. White pixels remove glossiness completely, and intermediate values reduce the size of the highlight.

Mapping the glossiness component is different from mapping specular color. Mapping glossiness alters the *location* of highlights, while specular mapping alters the *color* of highlights.

Glossiness mapping usually works best when you assign the same map to both Glossiness and Specular Level. (In the Maps rollout, you can do this by dragging from one map button to the other.)

(Glossiness was known as Shininess prior to 3ds Max **3**.)

Procedure

To map the glossiness value:

1. Click the Map button for the Glossiness value.

   The *Material/Map Browser (page 2–1447)* is displayed.

2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Self-Illumination Mapping

Material Editor > Standard material > Maps rollout > Self-Illumination button



Mapping self-illumination

You can select a bitmap file or *procedural map (page 3–997)* to map the *self-illumination (page 3–1009)* value. This makes portions of an object appear to glow. White areas of the map render as fully self illuminating. Black areas render with no self-illumination. Gray areas render as partially self illuminating, depending on the grayscale value.

Self-illumination means that the glowing area is not affected by lights in the scene (its ambient

color component goes away), and does not receive shadows.

Procedure

To map the self-illumination value:

1. Click the Map button for Self-Illumination.

   The *Material/Map Browser (page 2–1447)* is displayed.

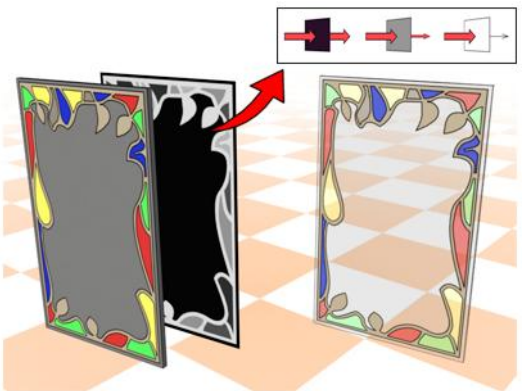
2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Opacity Mapping

Material Editor > Standard material > Maps rollout > Opacity button



The gray levels of an opacity map determine the amount of opacity.

You can select a bitmap file or *procedural map (page 3–997)* to make an object partially transparent. Lighter (higher-value) areas of the map render as opaque; darker areas render as transparent; and values in between are semi-transparent.

Setting the opacity map's Amount to 100 applies all of the map. Transparent areas are fully transparent.

Setting the Amount to 0 is the equivalent of turning the map off. Intermediate Amount values are blended with the Opacity value on the Basic Parameters rollout. Transparent areas of the map become more opaque.

Specular highlights are applied to transparent areas of the opacity map, as well as to opaque areas, creating the effect of glass. If you want the transparent areas to look like holes, map the *specular level (page 2–1502)* as well.

Procedure

To map the opacity value:

1. Click the Map button for Opacity.

   The *Material/Map Browser (page 2–1447)* is displayed.

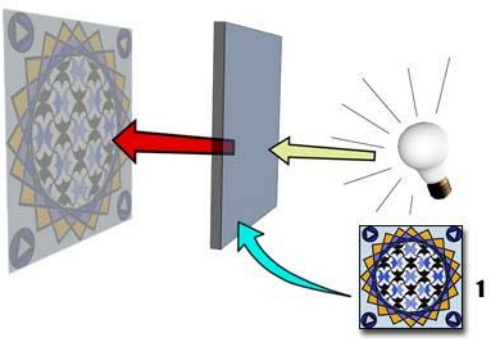
2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Filter Color Mapping

Material Editor > Standard material > Maps rollout > Filter Color button



Mapping filter color

The filter, or transmissive color, is the color transmitted through transparent or semi-transparent materials such as glass.

You can select a bitmap file or *procedural map (page 3–997)* to map the filter color component. This map applies a transparent-color effect based on the intensity of the map's pixels.

You can combine a mapped filter color with *volumetric lighting (page 3–288)* to create effects such as colored light through a stained-glass window. *Ray-traced shadows (page 3–1000)* cast by transparent objects are tinted by the filter color.

(In version 1 and 2, you specified the filter color on the *Basic Parameters rollout (page 2–1470)*. As of version 3, the filter color is specified on the *Extended Parameters rollout (page 2–1471)*.)

Procedure

To map the filter color:

1. Click the Map button for Filter color.

   The *Material/Map Browser (page 2–1447)* is displayed.

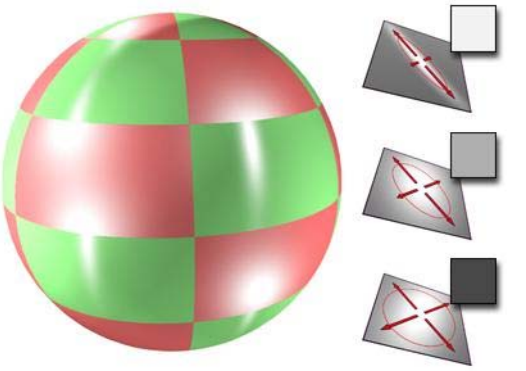
2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Anisotropy Mapping

Material Editor > Standard material > Shader Basic Parameters rollout > Anisotropic or Multi-Level shader > Maps rollout > Anisotropy button



Mapping anisotropy. The stretch of the highlight depends on the level of gray in the map.

You can select a bitmap file or *procedural map (page 3–997)* to control the Anisotropy parameter. The map controls the shape of the anisotropic highlight, roughly (but not necessarily) within the area specified by the glossiness parameter. Black and white values have little effect. Maps with a good deal of grayscale values, such as *Noise (page 2–1674)* or *Falloff (page 2–1670)*, can be very effective.

The anisotropy parameter is available with the *Anisotropic (page 2–1480)* and *Multi-Level (page 2–1481)* shaders.

The effect of mapping anisotropy is not very apparent unless the specular level is fairly high and glossiness is fairly low.

Reducing the Amount of the anisotropy map reduces the map's effect, and increases the effect of the Anisotropy value on the Basic Parameters rollout. When the Amount is 0 percent, the map isn't used at all.

Procedure

To map the anisotropy value:

1. Click the Map button for Anisotropy.

   The *Material/Map Browser (page 2–1447)* is displayed.

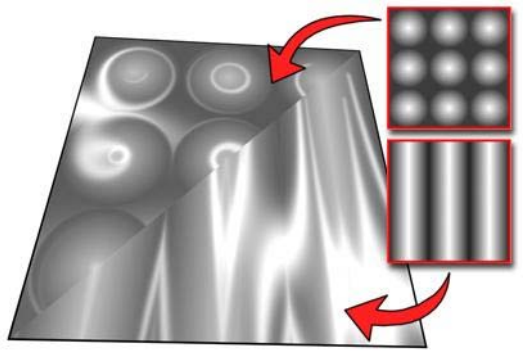
2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Orientation Mapping

Material Editor > Standard material > Shader Basic Parameters rollout > Anisotropic or Multi-Level shader > Maps rollout > Orientation button



Mapping anisotropy orientation

You can select a bitmap file or *procedural map (page 3–997)* to control the Orientation parameter. Orientation controls the position of the anisotropic highlight. Mapping orientation changes the highlight's position. Black and white values have little effect. Maps with a good deal of grayscale values, such as *Noise (page 2–1674)* or *Falloff (page 2–1670)*, can be very effective. You can also get a good effect using the same map for orientation mapping and *bump mapping (page 2–1506)*.

The orientation parameter is available with the *Anisotropic (page 2–1480)* and *Multi-Level (page 2–1481)* shaders.

Reducing the Amount of the orientation map reduces the map's effect, and increases the effect of the Orientation value on the Basic Parameters rollout. When the Amount is 0 percent, the map isn't used at all.

The effect of mapping orientation, like anisotropy, is not very apparent unless the specular level is fairly high and glossiness is fairly low.

Tip: Using an instance of the same map to control both anisotropy and orientation can give you good control over anisotropic highlights.

Procedure

To map the orientation value:

1. Click the Map button for Orientation.

   The *Material/Map Browser (page 2–1447)* is displayed.

2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Metalness Mapping

Material Editor > Standard material > Shader Basic Parameters rollout > Strauss shader > Maps rollout > Metalness button



Mapping metalness

Top: No mapping

Bottom: Mapping metalness with noise map

You can select a bitmap file or *procedural map (page 3–997)* to control the Metalness parameter. White pixels in the map increase metalness. Black pixels reduce metalness to 0. Intermediate values adjust metalness accordingly.

The metalness parameter is available with the *Strauss shader (page 2–1483)*.

Reducing the Amount of the metalness map reduces the map's effect, and increases the effect of the Metalness value on the Basic Parameters rollout. When the Amount is 0 percent, the map isn't used at all.

Procedure

To map the metalness value:

1. Click the Map button for Metalness.

   The *Material/Map Browser (page 2–1447)* is displayed.

2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

Bump Mapping

Material Editor > Standard material > Maps rollout > Bump button



An object with two different bump maps.

You can select a bitmap file or *procedural map (page 3–997)* to use for bump mapping. Bump mapping makes an object appear to have a bumpy or irregular surface. When you render an object with a bump-mapped material, lighter (whiter) areas of the map appear to be raised, and darker (blacker) areas appear to be low.

Note: The effect of a bump map is not previewed in viewports. You must render the scene to see the bump effect.

Bump mapping uses the intensity of the map to affect the surface of the material. In this case, the intensity affects the apparent bumpiness of the surface: white areas protrude, and black areas recede.

Use bump maps when you want to take the smoothness off a surface, or to create an embossed look. Keep in mind, however, that the depth effect of a bump map is limited. If you want extreme depth in a surface, you should use modeling techniques instead. For example, the *Displace modifier (page 1–629)* pushes surfaces or faces in and out based on the intensity of a bitmap image. (*Displacement mapping (page 2–1511)* is another way to do emboss a surface.)

Grayscale images can make effective bump maps. Maps that shade between white and black generally work better than maps with hard edges between the white and black areas.

The bump map Amount adjusts the degree of bumpiness. Higher values render as higher relief; low values render as low relief.

The bumps are a simulation created by perturbing face normals before the object is rendered. Because of this, bumps don't appear on the silhouette of bump-mapped objects.

Tip: If you render a bump-mapped material and notice aliasing in the highlights, try turning on *supersampling (page 2–1459)* and rendering again.

Note: Most controls on the *Output rollout (page 2–1621)* don't affect bump mapping. Only the Invert toggle is considered; it reverses the direction of the bumps.

Procedure

To assign a bump map:

1. Click the Map button labeled Bump.

   The *Material/Map Browser (page 2–1412)* is displayed.

2. Choose from the list of *map types (page 2–1617)*, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

   (If you choose Bitmap as the map type, you first see a file dialog that lets you choose the image file.)

3. Use the map controls to set up the map.

   Tip: To avoid aliasing caused by a 2D bump map, go to the bump map's Coordinates rollout. Set Blur to be in the range 0.3 to 0.6, and increase Blur Offset to be greater than 0.0. The default Blur and Blur Offset values work well for mapping other material components, but for bump mapping, lower Blur and higher Blur Offset values give better results.

To remove a bump map from a material:

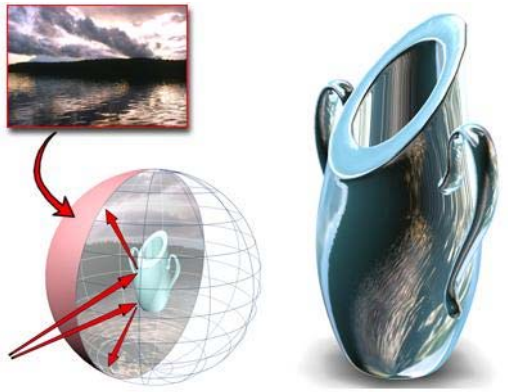
Tip: You can disable the map without removing it. Simply turn off the toggle immediately to the left of the map button on the Special Effects rollout.

1. If the Material Editor is displaying the map controls, click the *Type button (page 2–1449)* to display the Material/Map Browser. If the map controls aren't visible, click the Bump map button to display them, and then click the Type button.
2. In the Browser, choose NONE as the map type, and then click OK.

   The map is removed.

Reflection Mapping

Material Editor > Standard material > Maps rollout > Reflection button



Reflection created by mapping

You can select a bitmap file or *procedural map (page 3–997)* to use as a reflection map.

You can create three kinds of reflection: basic reflection maps, automatic reflection maps, and *flat-mirror (page 3–942)* reflection maps.

- A basic reflection map creates the illusion of chrome, glass, or metal by applying a map to the geometry so that the image looks like a reflection on the surface.
- An automatic reflection map uses no mapping at all, but looks outward from the center of the object, and maps what it sees onto the surface.

  Another way to generate reflections automatically is to assign a *Raytrace map (page 2–1698)* to be the reflection map.
- A flat-mirror reflection map is applied to a series of coplanar faces and reflects objects facing it, exactly like a real mirror.

  Reflection maps don't need *mapping coordinates (page 3–967)* because they're locked to the world, not to the geometry. The illusion of a reflection is created because the map doesn't move with the object, but with changes in the view, as do real reflections.

The most common use of reflection maps in a realistic scene is to add just a touch of reflection to an otherwise non-reflective surface. By default, reflection map strength is 100 percent, as it is for other maps. For many kinds of surfaces, however, reducing the strength gives the most realistic result. A polished table top, for example, primarily shows a wood grain; the reflections are secondary.

Reflection maps look more realistic if you increase the Glossiness and Specular Level values in the *Basic Parameters rollout (page 2–1470)*. They are also affected by the *diffuse (page 3–929)* and *ambient (page 3–908)* color values. The darker the color, the stronger the mirror effect.

Even when the Amount spinner is at 100, the reflection map is tinted by the ambient, diffuse, and *specular (page 3–1014)* colors.

In metal materials, the Diffuse color tints the reflection map. Specifically, the color from the reflection map is multiplied by the diffuse color (including a diffuse map, if one exists). The value (in the *HSV (page 3–1001)* description) of the diffuse color controls the reflection map intensity. If the diffuse color value is 255, the reflection is at full intensity; if the value is 0, the map is not visible.

In non-metal materials, the Specular color multiplies only reflection maps. The value (in the *HSV (page 3–1001)* description) of the specular color affects the reflection intensity. If the specular color value is 255, the reflection is at full intensity; if the value is 0, the map is not visible.

Procedures

To create an automatic reflection:

1. In the Maps rollout, click the Map button labeled Reflection.

2. In the *Material/Map Browser (page 2–1447)*, choose the Reflect/Refract map type, and then click OK.

   Adjusting the map's Strength slider in the parent material's Maps rollout controls how reflective the material is. At 100 percent, the material is fully reflective.

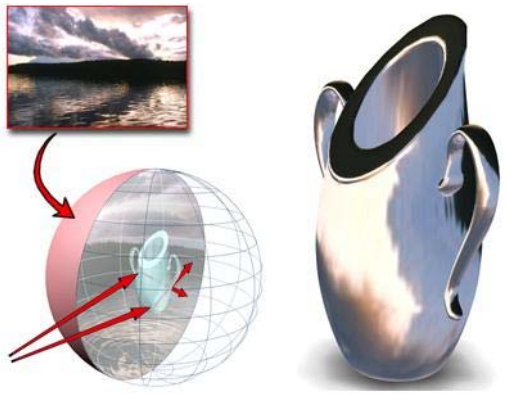
To assign a bitmap as a reflection map:

1. In the Maps rollout, click the Map button labeled Reflection. In the *Material/Map Browser (page 2–1447)*, double-click Bitmap.
2. In the Bitmap Parameters rollout, click the Bitmap button.
3. Use the file dialog to choose the bitmap file.
4. Reduce the Reflection map's Amount to get the effect you want.

   Refractions are similar to reflections. Bitmaps simulate reflections, while Reflect/Refract maps generate them based on the scene's background and geometry.

Refraction Mapping

Material Editor > Standard material > Maps rollout > Refraction button



Refractions show the scene or background through a refractive object.

You can select a bitmap file or a *procedural map (page 3–997)* such as *Reflect/Refract (page 2–1699)* to use for refraction mapping.

Refraction mapping is similar to reflection mapping. It maps the view onto the surface in such a way that the image looks like you're seeing it through the surface, rather than being reflected off it.

Like a reflection map, a refraction map's orientation is locked to the view rather than to the object. That is, as you move or rotate the object, the position of the refracted image remains fixed.

Setting the Index of Refraction

The physical properties of refractive objects often distort the image. A special parameter adjusts this distortion. It is in the parent material's *Extended Parameters rollout (page 2–1471)*.

Index of Refraction—The index of refraction (IOR) controls how severely the material refracts transmitted light. Left at 1.0, the IOR of air, the object behind the transparent object does not distort. At 1.5 the object behind distorts greatly (like a glass marble). At an IOR slightly less than 1.0, the object reflects along its edges (like a bubble seen from under water). Default=1.5 (the IOR of typical glass).

Common IORs (assuming the camera is in air or a vacuum) are:

Material	IOR Value
Vacuum	1.0 (exactly)
Air	1.0003
Water	1.333
Glass	1.5 to 1.7
Diamond	2.419

In the physical world, the IOR results from the relative speeds of light through the transparent material and the medium the eye or the camera is

in. Typically this is related to the object's density. The higher the IOR, the denser the object.

You can also use a map to control the index of refraction. IOR maps always interpolate between 1.0 (the IOR of air) and the setting in the IOR parameter. For example, if you set the IOR to 3.55 and use a black-and-white Noise map to control IOR, the IORs rendered on the object will be set to values between 1.0 and 3.55; the object will appear denser than air. If, on the other hand, you set the IOR to 0.5, then the same map values will render between 0.5 and 1.0: as if the camera is under water and the object is less dense than the water.

Here are some more IOR values for various materials:

Material	IOR Value
Carbon Dioxide, Liquid	1.200
Ice	1.309
Acetone	1.360
Ethyl Alcohol	1.360
Sugar Solution 30%	1.380
Alcohol	1.329
Flourite	1.434
Quartz, Fused	1.460
Calspar2	1.486
Sugar Solution 80%	1.490
Glass	1.500
Glass, Zinc Crown	1.517
Glass, Crown	1.520
Sodium Chloride	1.530
Sodium Chloride (Salt) 1	1.544
Polystyrene	1.550
Quartz 2	1.553
Emerald	1.570
Glass, Light Flint	1.575
Lapis Lazuli	1.610
Topaz	1.610
Carbon Bisulfide	1.630
Quartz 1	1.644
Sodium Chloride (Salt) 2	1.644
Glass, Heavy Flint	1.650
Methylene Iodide	1.740
Ruby	1.770
Sapphire	1.770
Glass, Heaviest Flint	1.890
Crystal	2.000
Diamond	2.417
Chromium Oxide	2.705
Copper Oxide	2.705
Amorphous Selenium	2.920
Iodine Crystal	3.340

Tip: The *Reflect/Refract (page 2–1699)* map type used as a Refraction map doesn't effectively model a material surrounding an object, such as a pencil in a glass of water. For this effect, use either the *Thin Wall Refraction (page 2–1703)* or the *Raytrace map type (page 2–1698)*.

Procedures

To create an automatic refraction:

1. In the Maps rollout, click the Map button labeled Refraction.
2. In the *Material/Map Browser (page 2–1447)*, choose the *Reflect/Refract (page 2–1699)* map type, and then click OK.

   At a Refraction Amount of 100 percent, the material is extremely refractive, regardless of the material's Opacity setting. At a Refraction Amount of 0 percent, the map is turned off. When the Amount is less than 100 percent, both the Reflect/Refract map and the Opacity setting control transparency.

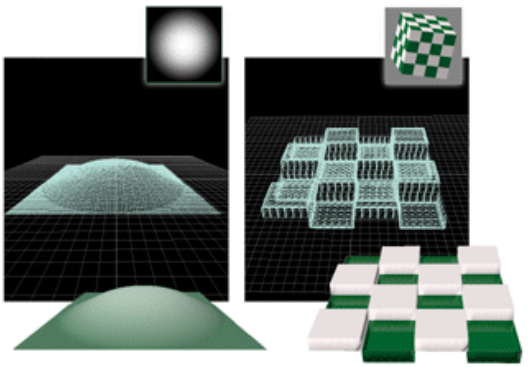
To assign a bitmap as a refraction map:

1. In the Maps rollout, click the Refraction map button.
2. Use the *Material/Map Browser (page 2–1447)* to choose the Bitmap type.
3. Use the file dialog to choose the bitmap file.

   For bitmapped refractions, you don't necessarily want to reduce the map strength.

Displacement Mapping

Material Editor > Standard material > Maps rollout > Displacement button



Using displacement mapping to alter a surface

A displacement map displaces the geometry of surfaces. The effect is similar to using the *Displace (page 1–629)* modifier. Unlike *bump mapping (page 2–1506)*, displacement mapping actually changes the geometry of the surface or patch tessellation. Displacement maps apply the gray scale of the map to generate the displacement. Lighter colors in the 2D image push outward more strongly than darker colors, resulting in a 3D displacement of the geometry.

Warning: A displacement map generates many triangular faces per surface, sometimes over 1M faces per surface. While displacement mapping can create good effects, there is a large cost in terms of time and memory.

The displacement Amount is measured as a percentage of the diagonal of the bounding box for the object that contains the patch or surface. This makes the displacement effect consistent for all surfaces in an object, and it also means that when you scale the object, the displacement is scaled with it.

You can apply a displacement map directly to the following kinds of objects:

- *Bezier patches (page 1–993)*
- *Editable meshes (page 1–996)*
- *Editable polymeshes (page 1–1022)*
- *NURBS surfaces (page 1–1078)*

For other kinds of geometry such as primitives, extended primitives, compound objects, and so on, you can't apply displacement mapping directly. To use displacement mapping with these kinds of objects, apply a *Disp Approx (page 1–628)* (Displacement Approximation) modifier. This makes the object's surface displaceable. Disp Approx works with any kind of object that you can convert to an editable mesh.

Displacement mapping isn't visible in viewports unless you apply a modifier to make it so.

- For NURBS surfaces, you can make displacement mapping visible in viewports and editable as a mesh object by using the *Displace NURBS (page 1–515)* world space modifier.
- For editable meshes and objects with Disp Approx applied to them, use the *Displace Mesh (page 1–514)* modifier to obtain the same effect.

  Note: If you apply a *UVW Map (page 1–922)* modifier to the surface, all maps obtain their coordinates from the modifier except for the displacement map, which *always* obtains its coordinates from the original surface or the Disp Approx modifier.

Under certain circumstances, such as when the underlying mesh is fairly simple, displacement

mapping of an editable mesh can cause problems because of the way the underlying mesh is tessellated. (These problems don't occur when you apply displacement mapping to a NURBS surface.) When this happens, smoothing does not work properly and you can see the underlying wireframe mesh in the surface itself. To correct this problem, use these techniques:

- Avoid applying displacement mapping to large areas of a single color. Map the diffuse color and use a small amount of variation, such as slight amount of noise, in the map you use for the diffuse color.
- Add a small amount of noise to the map you use for displacement. This can complicate the tessellation enough to ease the problem.
- Add detail to the mesh. The more initial faces, and the smoother the mesh curvature, the more even the displacement mapping will be.

Procedures

To apply a displacement map to a NURBS surface, editable mesh, or patch:

1. In a material's Maps rollout, click the map button for Displacement.

   The *Material/Map Browser (page 2–1447)* is displayed.

2. Choose from the list of map types, and then click OK.

   The Material Editor is now at the map level, and displays controls for the map parameters.

3. Use the map controls to set up the map.

To apply a displacement map to other kinds of objects:

1.  Select the object. Go to the Modify panel and choose Disp Approx from the Modifiers drop-down list.

   You can adjust the Disp Approx modifiers parameters, or you can leave them at their default settings.

2.  Go to the Material Editor.
3. In a material's Maps rollout, click the map button for Displacement.

   The *Material/Map Browser (page 2–1447)* is displayed.

4. Choose from the list of map types, and then click OK.

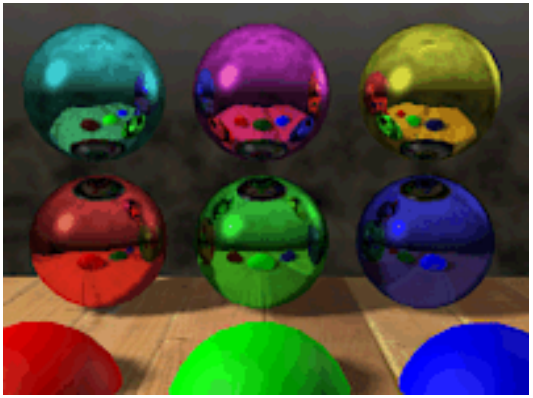
   The Material Editor is now at the map level, and displays controls for the map parameters.

5. Use the map controls to set up the map.

   For example, if you chose Bitmap as the map type, you now need to select the bitmap file to use.

Raytrace Material

Material Editor > Type button > Material/Map Browser > Raytrace



Balls using raytrace material to reflect each other

Raytrace material is an advanced surface shading material. It supports the same kinds of diffuse

surface shading that a standard material does. It can also create fully raytraced reflections and refractions. It also supports fog, color density, *translucency (page 3–1027)*, *fluorescence (page 3–942)*, and other special effects.

The reflections and refractions Raytrace material generates are more accurate than those produced by the *Reflect/Refract (page 2–1699)* map. Rendering raytraced objects can be slower than using Reflect/Refract. On the other hand, Raytrace is optimized for rendering 3ds Max scenes. You can further optimize it for your scene by excluding specific objects from raytracing.

Note: If you want accurate, raytraced reflections or refractions in a standard material you can use the *Raytrace map (page 2–1698)*, which uses the same raytracer. The Raytrace map and material share global parameter settings.

Important: Raytrace map and Raytrace material use a surface's normal to decide whether a ray is entering or exiting a surface. If you flip the normals of an object, you can get unexpected results. Making the material 2-Sided doesn't correct the problem as it often does with reflections and refractions in Standard materials.

In some cases, the colors in the Basic Parameters rollout of Raytrace material behave differently from colors in standard materials. Standard material has a diffuse shading model that does an excellent job of rendering solid, nonreflective objects such as plastic, ceramic, and so on. In effect, this model applies color to the object. The color components in Raytrace material, on the other hand, attempt to model their physical counterparts in nature.

In Raytrace material, the surface reflects its Diffuse color component without specular reflection, while the Reflect color component controls the amount of specular reflection. These two material components are layered together. The results you see depend on the layering effect. For example, if the material is not transparent and completely reflective, no diffuse color is visible. If the material is not transparent and completely nonreflective, only the diffuse color is visible.

The Dynamics Properties rollout for the Raytrace material contains the same controls as the dynamics properties for a *standard material (page 2–1465)*.

Raytrace material has a large user interface with a lot of controls. In general, if you are using Raytrace to create reflections and refractions, the controls in the Basic Parameters rollout are the only ones you need to adjust. The Extended Parameters rollout for Raytrace has controls for special effects. The Raytracer Controls rollout affects the raytracer itself. Use the Raytracer Controls to turn the raytracer on or off, and to toggle other options. Use the *Global Raytracer Settings dialog (page 2–1528)* (Rendering > Raytrace Globals) to set options globally (for all Raytrace materials and maps in the scene), including recursion depth.

Interface

Raytrace material has the following rollouts, which are described in these topics:

Raytrace Basic Parameters Rollout (page 2–1514)

Raytrace Extended Parameters Rollout (page 2–1519)

Raytracer Controls Rollout (page 2–1521)

SuperSampling Rollout (page 2–1459)

Raytrace Maps Rollout (page 2–1523)

Raytrace Dynamics Properties Rollout (page 2–1527)

The following dialogs are also part of the Raytrace material's interface:

Raytracer Global Parameters Rollout (page 2–1528)

Raytracing Acceleration Parameters Dialog (page 2–1531)

Raytrace Exclude/Include Dialog (page 2–1531)

Raytrace Antialiaser Dialog: Fast Adaptive Antialiaser (page 2–1533)

Raytrace Antialiaser Dialog: Multiresolution Adaptive Antialiaser (page 2–1534)

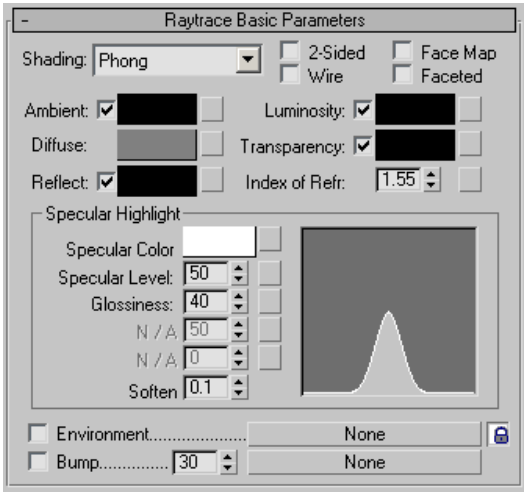
Raytrace Basic Parameters Rollout

Main toolbar > Material Editor > Type button > Material/Map Browser > Raytrace > Basic Parameters rollout

Menu bar > Tools menu > Material Editor > Type button > Material/Map Browser > Raytrace > Basic Parameters rollout

The Basic Parameters rollout for a *Raytrace material (page 2–1512)* controls the material's shading, color components, reflectivity or refractivity, and bumps.

Interface



The basic parameters in this rollout are similar to the basic parameters for standard materials, but the color components of a Raytrace material behave differently.

As with standard materials, you can use a map for Raytrace color components and various other parameters. The small buttons to the right of the color swatches and parameters take you to the *Material/Map Browser (page 2–1412)*, where you select a map of corresponding type. These are shortcuts that also have corresponding buttons in the Maps rollout. If you have assigned a map to one of these colors, the button displays the letter M. An uppercase M means that the corresponding map is assigned and active. A lowercase m means that the map is assigned and inactive (turned off).

Shading drop-down list—Chooses a shader. Depending on the shader you choose, the Specular Highlight can change to show the controls for that shader. The alternatives are:

- *Anisotropic: (page 2–1492)* For surfaces with elliptical, "anisotropic" highlights.
- *Blinn: (page 2–1493)* For rounder, softer highlights than Phong shading.
- *Metal: (page 2–1494)* For metallic highlights.
- *Oren-Nayar-Blinn: (page 2–1493)* For matte surfaces such as fabric or terra-cotta.
- *Phong: (page 2–1493)* For surfaces with strong, circular highlights. Phong is the default shading type.

2-Sided—Same as for standard materials. When on, shades and raytraces both sides of faces. By default, objects are one-sided in order to speed up rendering.

If you have a 2-sided, reflective and refractive object, and you use the *raytrace map (page 2–1698)* rather than the material, the raytracer runs until it hits the maximum recursion level. This can be time-consuming.

Wire—Same as for standard materials. When on, renders the material in *wireframe mode (page 3–1034)*. You can specify the wire size in the Extended Parameters rollout.

With pixels, wires maintains the same apparent thickness regardless of the scale of the geometry or how near or far the object is positioned. With units, the wires appear thinner at a distance and thicker at close range, as if they were modeled in the geometry.

Face Map—Applies the material to the faces of the geometry. If the material is a mapped material, it requires no *mapping coordinates (page 3–967)*. The map is automatically applied to each facet of the object.

Faceted—Renders each face of a surface as if it were flat.

Note: Raytrace material has the same *SuperSampling (page 2–1459)* options as a Standard material.

Ambient—This is *not* the same as the standard ambient color. For Raytrace material, this controls an ambient absorption factor: that is, how much the material absorbs ambient light. Setting Ambient to white is the same as locking the ambient and diffuse colors in a standard material. Default=black.

- **Ambient Color check box—**When on, the material uses an ambient color. When off, the material uses a spinner to set a grayscale value only. Default=on.
- **Color swatch—**When on, the color swatch shows the ambient color. To change the color, click the swatch and then use the *Color Selector (page 1–161)*.
- **Mono spinner—**When the check box is off, the ambient component is gray, and this spinner lets you adjust the gray value.

  Click the map button to assign a map to the ambient component. See *Ambient Mapping (page 2–1497)*. This button is a shortcut: you can also assign ambient mapping on the *Raytrace Maps rollout (page 2–1523)*.

Diffuse—Sets the diffuse color. This is the same as the standard diffuse color. It is the color that the object reflects, *without* specular reflection. Reflection and transparency effects are layered on top of the diffuse result. When Reflect is 100% (pure white), the diffuse color isn't visible. (This differs from the standard material.) Default=50% gray.

Reflect—Sets the specular reflection color. This is the color that the reflected environment (that is, the rest of the scene) is filtered through. The color's Value controls the amount of reflection. If your reflect color is saturated and the diffuse color is black, the effect is like colored chrome (for example, colored Christmas tree balls). Default=black (no reflection).

If raytracing is off (on the Raytracer Controls rollout), the object still reflects the environment, but ignores other objects in the scene. The environment can be the background color, the environment map, or the map in the Raytrace material's Environment component.

Tip: If you turn off raytraced reflections, set the Reflect color to a color other than black, and use a Reflect/Refract map for the local environment (see the Environment parameter, below), you get the same effect as a reflection map in a standard material. This can improve rendering time.

Note: Raytrace reflects and transmits the IDs in *material effects channel (page 2–1443)* (*G-buffer (page 3–946)*), so it can create glowing reflections, and so on.

- **Reflect Color check box—**When on, the material uses a reflection color. When off, the material uses a spinner to set a grayscale value only. Default=on.
- **Color swatch—**When the check box is on, the color swatch shows the reflection color. To change the color, click the swatch and then use the *Color Selector (page 1–161)*.

- **Mono spinner—**When the check box is off, the reflection color component is gray, and this spinner lets you adjust the gray value.
- **Fresnel—**Clicking the check box a second time displays this option. When active, applies a Fresnel effect to the reflection. This can add a bit of refraction to the reflecting object, depending on the viewing angle of the object.

Click the map button to assign a map to the reflect component. This button is a shortcut: you can also assign reflect mapping in the *Raytrace Maps rollout (page 2–1523)*.

Luminosity—Similar to the Standard material's self-illumination component, except that it does not depend on the diffuse color. You can have a blue diffuse object with red luminosity. Default=black.

Note: When Luminosity is off, the name of this control changes to Self-Illum (Self-Illumination).

- **Luminosity check box—**When on, the material uses a luminosity color. When off, the material uses a spinner to set a grayscale value only for self-illumination. Default=on.
- **Color swatch—**When the check box is on, the color swatch shows the luminosity color. To change the color, click the swatch and then use the *Color Selector (page 1–161)*.
- **Mono spinner—**When the check box is off, the luminosity color component is gray, and this spinner lets you adjust the gray value.

Click the map button to assign a map to the luminosity component. This button is a shortcut: you can also assign reflect mapping in the *Raytrace Maps rollout (page 2–1523)*.

Transparency—Similar to the standard material's filter color for transmitted light, combined with the standard material's opacity controls. This color filters scene elements that are behind the object with Raytrace material. Black is opaque, white is fully transparent, and any value in between filters objects behind the raytraced object. A fully saturated color in both the diffuse and transparency components gives the effect of tinted glass. If you want more of an opaque look, pick the color you want as a transparent color, copy it to the diffuse color, make the diffuse color fully saturated, and then adjust the transparency to get the effect you want. Default=black (no transparency).

If raytracing is turned off (in the Raytracer Controls rollout), the object still refracts the environment mapping, but ignores other objects in the scene.

Note: By separating the diffuse, reflect, and transparency components, Raytrace material gives you a great deal of control over how the object reacts to its environment. For example, an object might diffusely reflect red, specularly reflect green, and transmit blue. This is not a real-world effect, but it can be useful.

- **Transparency Color check box—**When on, the material uses a transparency color. When off, the material uses a spinner to set a grayscale value only. Default=on.
- **Color swatch—**When the check box is on, the color swatch shows the transparency color. To change the color, click the swatch and then use the *Color Selector (page 1–161)*.
- **Mono spinner—**When the check box is off, the transparency color component is gray, and this spinner lets you adjust the gray value.

Click the map button to assign a map to the transparency component. See *Filter Color Mapping (page 2–1503)*. This button is a shortcut: you can also assign reflect mapping in the *Raytrace Maps rollout (page 2–1523)*.

Index of Refr. (Refraction)—The index of refraction (IOR) controls how severely the material refracts transmitted light. At 1.0, the IOR of air, the object behind the transparent object does not distort. At

1.5, the object behind distorts greatly, like a glass marble. At an IOR slightly less than 1.0, the object reflects along its edges, like a bubble seen from under water. Default=1.0.

Common IORs (assuming the camera is in air or a vacuum) are:

Material	IOR Value
Vacuum	1.0 (exactly)
Air	1.0003
Water	1.333
Glass	1.5 to 1.7
Diamond	2.419

In the physical world, the IOR results from the relative speeds of light through the transparent material and the medium the eye or the camera is in. Typically this is related to the object's density, and the higher the IOR, the denser the object.

You can also use a map to control the IOR. IOR maps always interpolate between 1.0 (the IOR of air) and the setting in the IOR parameter. For example, if the IOR is set to 3.55 and you use a black-and-white Noise map to control IOR, the IORs rendered on the object will be set to values between 1.0 and 3.55. The object will appear denser than air. If, on the other hand, your IOR is set to 0.5, then the same map values will render between 0.5 and 1.0, as if the camera were under water and the object was less dense than the water.

Here are some more IOR values for various materials:

Material	IOR Value
Carbon Dioxide, Liquid	1.200
Ice	1.309
Acetone	1.360
Ethyl Alcohol	1.360
Sugar Solution 30%	1.380
Alcohol	1.329
Flourite	1.434
Quartz, Fused	1.460
Calspar2	1.486
Sugar Solution 80%	1.490
Glass	1.500
Glass, Zinc Crown	1.517
Glass, Crown	1.520
Sodium Chloride	1.530
Sodium Chloride (Salt) 1	1.544
Polystyrene	1.550
Quartz 2	1.553
Emerald	1.570
Glass, Light Flint	1.575
Lapis Lazuli	1.610
Topaz	1.610
Carbon Bisulfide	1.630
Quartz 1	1.644
Sodium Chloride (Salt) 2	1.644
Glass, Heavy Flint	1.650
Methylene Iodide	1.740
Ruby	1.770
Sapphire	1.770
Glass, Heaviest Flint	1.890
Crystal	2.000
Diamond	2.417
Chromium Oxide	2.705
Copper Oxide	2.705
Amorphous Selenium	2.920
Iodine Crystal	3.340

Specular Highlight group

The controls in this group affect the appearance of specular highlights. Specular highlights simulate

the surface of the raytraced object reflecting the lights in the scene. Changing the color or intensity of lights in the scene can change the appearance of specular highlights.

As in standard materials, as you adjust the values in this group the highlight curve at the right changes to give you an idea of the effect. The material preview in the sample slot also updates.

Specular Color—Sets the *specular color (page 3–1014)*, assuming there are white lights in the scene.

Click the color swatch to display the *Color Selector (page 1–161)* and change the highlight color.

Click the map button to assign a map to the specular color. See *Specular Mapping (page 2–1500)*. This button is a shortcut: you can also assign specular level mapping in the *Raytrace Maps rollout (page 2–1523)*.

The remaining controls in the Specular Highlight group depend on the active shader, as shown next to "Shading:" at the top of this rollout. These highlight controls are the same as for the Standard material shaders.

These are the highlight controls available to Raytrace materials:

- *Anisotropic highlights (page 2–1492)*
- *Blinn, Oren-Nayar-Blinn, and Phong highlights (page 2–1493)*
- *Metal highlights (page 2–1494)*

  Note: Highlight controls that don't pertain to the current shader are labeled "N/A."

Environment—Specifies an environment map that overrides the global environment map. Both Reflect and Transparency use the scene-wide environment map unless you use this button to specify another map. With this control, you can use different environment maps on a per-object basis, or provide an environment to specified objects when the scene as a whole has none.

This map overrides the scene-wide environment for both reflection and refraction. To override for refraction alone, see the Transparency Environment control in the Extended Parameters rollout.

Use the check box to turn this map on or off.

Tip: You can use any map as the Raytrace environment, including the *Reflect/Refract (page 2–1699)* map. Reflect/Refract map is often adequate for getting the required look, and it usually renders more quickly than raytracing the entire scene, especially if the Raytrace material is transparent. If you are using Raytrace material just to get the glass to look right on a car's headlight or on a light bulb hanging in the middle of a room, turn off the raytracer and use an environment map instead.

Lock button—Locks the Environment map to the Transparency Environment map (found on the *Extended Parameters rollout (page 2–1519)*). When on, the Transparency Environment map controls are disabled, and a map applied to the Raytrace Environment applies to the Transparency Environment as well. When off, the Transparency Environment map controls are enabled, and the Transparency Environment can have a different map assigned to it. Default=on.

Changing this button's setting here also changes it on the *Extended Parameters rollout (page 2–1519)* and the *Maps rollout (page 2–1523)*.

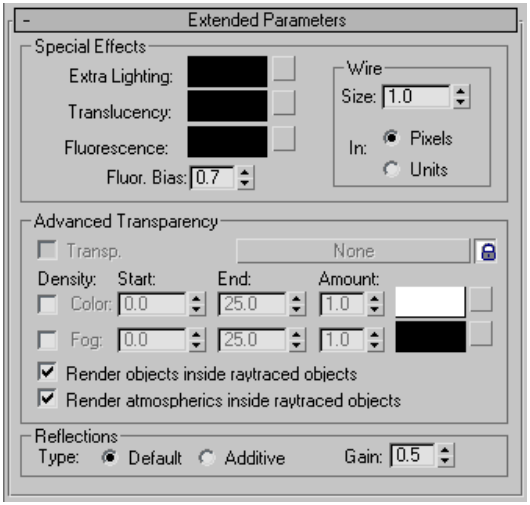
Bump—This is the same as *bump mapping (page 2–1506)* for Standard materials. Click the button to assign the map. Use the spinner to change the bump Amount. Use the check box to turn the map on or off.

Raytrace Extended Parameters Rollout

Material Editor > Type button > Material/Map Browser > Raytrace > Extended Parameters rollout

The Extended Parameters rollout for a *Raytrace material (page 2–1512)* controls the material's special effects, transparency properties, and advanced reflectivity.

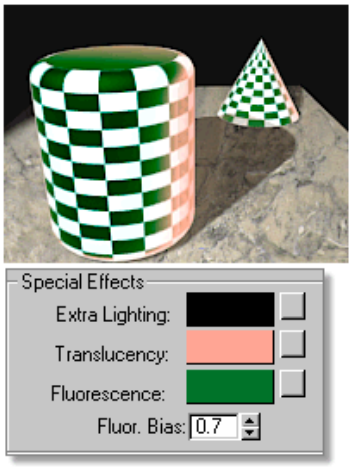
Interface



Except for the wireframe controls, the controls in the Extended Parameters rollout for Raytrace material are specific to the Raytrace material.

Special Effects group

The controls in this group are special effects. They are powerful, but you might need to experiment to use them effectively.



Object with raytrace material using translucency and fluorescence

Extra Lighting—Adds light to the surface of objects with the Raytrace material. You can view this as an ambient lighting color that you can control on a per-material basis, but don't confuse it with the ambient absorption from the Basic Parameters rollout.

By mapping this parameter, you can simulate radiosity: the ambient light that results from reflected light in a scene. One effect of radiosity is color bleeding. For example, in strong light a white shirt next to an orange wall will show a reflected orange color.

Translucency—Creates a *translucent effect (page 3–1027)*. The Translucency color is a non-directional diffuse reflection. The diffuse color on an object depends upon the angle between the surface normal and the position of the light source. By ignoring the surface normal alignment, this color component simulates translucent materials.

For thin objects, the appearance can be like shining a light on the back of a piece of rice paper. You can cast shadows onto the back of the paper and see them projected through the paper; this works well

with a projector light. On thicker objects, you can get some good wax-like effects.

Fluorescence and Fluor. Bias—Creates an effect similar to black light on a black light poster. The light from a black light is largely ultraviolet, outside the visible spectrum. Under black light, *fluorescent (page 3–942)* paints flare or glow. The fluorescence in Raytrace material takes whatever light it sees in the scene, applies the Bias to it, and then, regardless of the color of the lights in the scene, illuminates the fluorescent material as if it were lit by white light.

At 0.5, The Bias makes Fluorescence behave just like diffuse coloring. Bias values higher than 0.5 increase the fluorescent effect, making the object brighter than other objects in the scene. Bias values lower than 0.5 make the object dimmer than other objects in the scene. You can get some chromatic shifting effects with this.

Tip: Full saturation and value for the Fluorescence color help give the effect of commercial fluorescent paints.

Tip: A slight amount of Fluorescence can add to the realism of skin and eyes.

Wire group

Size—Sets the size of the wire in *wireframe mode (page 3–1034)*. You can set either pixels (the default) or current units.

In—Chooses how to measure wire. With pixels, wires maintain the same apparent thickness regardless of the scale of the geometry or how near or far the object is positioned. With units, the wires appear thinner at a distance and thicker at close range, as if they were modeled in the geometry.

Pixels—(The default.) Measures wire in pixels.

Units—Measures wire in 3ds Max units.

Advanced Transparency group

The controls in this group let you further tune transparency effects.

Transp. (Transparency Environment)—Similar to the environment map in Basic Parameters, but overrides the scene's environment map for transparency (refraction) only. Transparent objects refract this map, while reflections still reflect the scene (or the Basic Parameters Environment map, if one is chosen).

Click the button to choose the Transparency Environment map. Use the check box to toggle the effect of the map.

Lock button—Locks the Transparency Environment map to the Environment map (found on the *Basic Parameters rollout (page 2–1514)*). When on, the Transparency Environment map controls are disabled, and a map applied to the Raytrace Environment applies to the Transparency Environment as well. When off, the Transparency Environment map controls are enabled, and the Transparency Environment can have a different map assigned to it. Default=on.

Changing this button's setting here also changes it on the *Basic Parameters rollout (page 2–1514)* and the *Maps rollout (page 2–1523)*.

Density—The density controls are for transparent materials. If the material is opaque (the default), they have no effect.

Color—Sets a transmission color based on thickness. While filter (Transparency) color tints objects behind the transparent object, the density color gives the appearance of color *within* the object itself, like tinted glass.

To use, first make sure the object is transparent. Click the color swatch to display the Color Selector. Choose a color, and then turn on the check box.

The Amount controls the amount of density color. Reducing this value reduces the density color effect. Range=0 to 1.0. Default=1.0.

A thin piece of tinted glass is mainly clear, while a thick piece of the same glass has more color. The Start and End controls help you simulate this effect. They are expressed in world units. Start is the position in the object where the density color begins to appear. (Default=0.0.) End is the position in the object where the density color reaches its full Amount value. (Default=25.0) To have a lighter effect, increase the End value. To have a heavier effect, reduce the End value.

The object must be at least as thick as the Start value before the density color is visible.

You can map this color component.

Fog—Density fog is also a thickness-based effect. It fills the object with a fog that is both opaque and self illuminated. The effect is like smoke trapped in a glass, or wax at the tip of a candle. Colored fog in tubular objects can resemble neon tubes.

To use, first make sure the object is transparent. Click the color swatch to display the Color Selector. Choose a color, and then turn on the check box.

The Amount controls the amount of density fog. Reducing this value reduces the density fog effect and makes the fog translucent. Range=0 to 1.0. Default=1.0.

The Start and End controls let you adjust the fog effect based on the object's dimensions. They are expressed in world units. Start is the position in the object where the density fog begins to appear. (Default=0.0.) End is the position in the object where the density fog reaches its full Amount value. (Default=25.0) To have a lighter effect, increase the End value. To have a heavier effect, reduce the End value.

You can map this color component.

Render objects inside raytraced objects—Turns the rendering of objects inside raytraced objects on or off. Default=on.

Render atmospherics inside raytraced objects—Turns the rendering of atmospheric effects inside raytraced objects on or off. Atmospheric effects include fire, fog, volume light, and so on. Default=on.

Reflections group

Controls in this group give you finer control over reflections.

Type—When set to **Default,** reflections are layered with the Diffuse color. For example, if the material is not transparent and completely reflective, no diffuse color is visible. When set to **Additive,** reflections are added to the Diffuse color, as in *Standard materials (page 2–1465)*. The diffuse component is always visible.

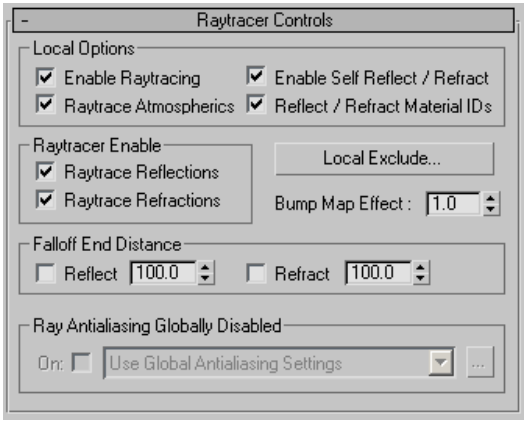
Gain—Controls reflection brightness. The lower the gain value, the brighter the reflection. At a gain of 1.0, no reflection is visible. Default=0.5.

Raytracer Controls Rollout

Material Editor > Type button > Material/Map Browser > Raytrace material > Raytracer Controls rollout

The Raytracer Controls rollout for a *Raytrace material (page 2–1512)* controls affect the operation of the raytracer itself. It can help you improve rendering performance.

Interface



Local Options group

Enable Raytracing—Turns the raytracer on or off. Default=on.

Even with raytracing off, Raytrace material and Raytrace map still reflect and refract the environment, including both the environment map for the scene, and the environment map assigned to the Raytrace material.

Raytrace Atmospherics—Turns the raytracing of atmospheric effects on or off. Atmospheric effects include fire, fog, volume light, and so on. Default=on.

Enable Self Reflect/Refract—Turns self reflection/refraction on or off. Default=on.

Can an object reflect itself? For example, a teapot's body reflects the teapot's handle, but a sphere will never reflect itself. If you don't need this effect, you can improve render time by turning off this toggle.

Tip: If you have a transparent object such as glass, and Enable Self Reflect/Refract is on, you don't have to make the object *2-sided (page 3–901)*. The raytracer sees back faces when exiting refractive objects.

Reflect/Refract Material IDs—When on, the material reflects effects assigned to material IDs in the renderer's *G-buffer (page 3–946)* on or off. Default=on.

By default, Raytrace material and Raytrace map reflect effects assigned to a material's ID, so that G-buffer effects are not lost. For example, if a raytraced object reflects a lamp made to glow with the Video Post Glow filter (Lens Effects Glow), the reflection glows as well.

Raytracer Enable group

These two check boxes turn raytracing of reflections or refractions on or off for this material. If you are using the Raytrace material to create only reflections or refractions, turn off the one you aren't using to improve rendering time.

Raytrace Reflections—Turns raytracing of reflective objects on or off. Default=on.

Raytrace Refractions—Turns raytracing of transparent objects on or off. Default=on.

Local Exclude—Displays the local *Exclude/Include dialog (page 2–1531)*.

An object that is excluded locally is excluded from this material only.

Tip: Using exclusion lists is one of the best and simplest ways to speed up the raytracer.

Bump Map Effect—Adjusts the effect of bump maps on raytraced reflections and refractions. Default=1.0.

Falloff End Distance group

Reflect—Dims reflections to black at this distance. Default=100.0.

Refract—Dims refractions to black at this distance. Default=100.0.

Note: The toggles for Reflect Falloff and Refract Falloff aren't animatable.

Raytraced Reflection and Refraction Antialiaser group

Controls in this group let you override the global antialiasing settings for raytraced maps and materials. They are unavailable if antialiasing is turned off globally. To turn on antialiasing globally, choose Rendering > Raytrace Globals to display the *Global Raytracer Settings dialog (page 2–1528)*.

On—When on, uses antialiasing. Default=unavailable unless global antialiasing is on; on if global antialiasing is turned on.

Drop-down list—Chooses which antialiasing settings to use. There are three alternatives:

Use Global Antialiasing Settings—(The default.) Uses the global antialiasing settings.

Fast Adaptive Antialiaser—Uses the Fast Adaptive antialiaser, regardless of the global setting.

Multiresolution Adaptive Antialiaser—Uses the Multiresolution Adaptive antialiaser, regardless of the global setting.

...—The button with the ellipsis to the right of the drop-down list displays another dialog to let you set antialiasing controls locally. The dialog displayed depends on which alternative you chose in the drop-down list, as follows:

Use Global Antialiasing Settings—Clicking ... displays the *Global Raytracer Settings dialog (page 2–1528)*.

Fast Adaptive Antialiaser—Clicking ... displays the *Fast Adaptive Antialiaser dialog (page 2–1533)*.

Multiresolution Adaptive Antialiaser—Clicking ... displays the *Multiresolution Adaptive Antialiaser dialog (page 2–1534)*.

When you change settings for an antialiaser locally, you don't affect the global settings for that antialiaser.

Raytrace Maps Rollout

Material Editor > Type button > Material/Map Browser > Raytrace material > Maps rollout

As with a standard material, the Maps rollout for a *Raytrace material (page 2–1512)* contains map buttons for the components of the Raytrace material that can be mapped.

You can choose from a large variety of map types. See *Map Types (page 2–1617)* to find descriptions of these types, and how to set their parameters.

Assigning the Same Map to Different Parameters

Applying the same map to different parameters is useful in some cases. For example, using a pattern as both a self-illumination map and an opacity map can make the pattern appear to glow and hover in space.

Blending Map Amounts for Opacity and Other Material Components

The Specular Level, Glossiness, Self-Illumination, and Opacity values in the four spinners in the Basic Parameters rollout are blended with their associated map Amount values in the Maps rollout.

When the Opacity spinner is set to 0, the map Amount spinner completely controls Opacity. That is, reducing the Amount value increases the transparency of the entire surface. On the other hand, when Opacity is 100, reducing the map Amount value increases the opacity of the areas where the Opacity map is less than 1. For example, you can now adjust a Checker Opacity map so that the solid areas remain solid, while the clear areas are semi-transparent.

The Specular Level, Glossiness, and Self-Illumination channels all behave in the same way. A setting of 100 applies all of the map; a setting of 0 is the equivalent of turning the map off.

When you load old 3ds Max files or bring earlier materials from the Browser into the Materials Editor, the spinner values for Opacity, Specular Level, Glossiness, and Self-illumination are altered, where necessary, to maintain the equivalent material effect.

Procedures

To assign a map:

1. Click a map button.

   A modal *Material/Map Browser (page 2–1412)* is displayed.

2. Use the Browse From buttons to choose where you want to look.

   If you choose Material Library and the dialog's display area is blank, you need to open a library file. Click the Open button and then choose the library to browse.

3. Use the display buttons to choose how you view maps.

   - View List shows each map by name.
   - View List + Icons shows a small preview and each map's name.
   - View Small Icons shows a small preview for each map.
   - View Large Icons shows a large preview for each map, along with the map's name.

     Tip: You can resize the Browser dialog to increase the size of the display area. This is especially useful when you view large icons.

4. Double-click the map you want.

To use the same map for different parameters:

1. In the Maps rollout, use a map button to assign a map.

   The Material Editor is now at the map level, and displays controls for the map parameters.

2.  Click Go To Parent to return to the material level, and then open the Maps rollout.

3. Drag the assigned map button to another map button.

   The *Copy (Instance) Map dialog (page 2–1451)* is displayed.

4. Choose Copy or Instance, and then click OK.

   If you choose Swap, the Material Editor swaps the two button assignments.

To view the parent material's parameters:

-  If you are currently at the map level in the Material Editor, click Go To Parent.

  The parameters for the map's parent material are displayed. Also, the Show End Result and Go to Parent buttons become unavailable.

To view a map's parameters:

- If you are currently at the material level in the Material Editor, click the button that corresponds to the map.

  The parameters for the map are displayed. Also, the *Show End Result (page 2–1446)* and *Go to Parent (page 2–1446)* buttons become available.

  In the Basic Parameters rollout, if a map has been assigned to a color component or parameter, the corresponding button displays a letter M. In the Maps rollout, if a map has been assigned, the corresponding button displays the map name.

To view a map's location:

-  Click Material/Map Navigator to view the Navigator.

  The *Material/Map Navigator (page 2–1447)* displays the hierarchy of the current material, which contains the map.

To go to a map using the Navigator:

- In the *Material/Map Navigator (page 2–1447)*, click the name of the map, or the green or red parallelogram to the left of the map's name.

  The Navigator goes to the level of the map, and the Material Editor displays the controls for the map you clicked.

  As the Navigator's map tree shows, maps for basic material components and parameters are one level below the material itself.

To preview a map in a sample slot:

1. Go to the level of the map, as described in previous procedures.

   The Material Editor displays the map's parameters.

2.  Turn off *Show End Result (page 2–1446)*.

   The sample slot shows the map instead of the material. If the map contains sub-maps, these are also visible.

   By default, the sample slot displays a map with no three-dimensional shading. You can change this in the *Material Editor Options dialog (page 2–1436)*.

To view the map interactively:

1. Select an object.
2. In the object's creation parameters, make sure that Generate Mapping Coords is on.

   If the object type does not have a Generate Mapping Coordinates toggle, you need to assign mapping coordinates by applying a *UVW Map modifier (page 1–922)*.
3. In the Material Editor, assign the mapped material to the object.
4. If you are at the material level (the top level), click the appropriate map button to go to the map level.
5.  Turn on *Show Map in Viewport (page 2–1445)*.

   The map appears on objects assigned the material in all shaded viewports. Now when you adjust the map, the viewports update to display the adjustments.

   Turning on Show Map In Viewport for one map automatically turns this button off for all other maps the material has.

   Viewports can display 2D maps such as Checker and Bitmap. They cannot display other kinds of maps such as 3D maps. Show Map in Viewport is unavailable if the active map type cannot display in viewports.

   Tip: Displaying mapped materials in a viewport can slow performance. If you don't need to view the texture, turn off its viewport display.

To turn off interactive texture display:

1. Go to the map level.

   If you are at the material level, click the appropriate map button to go to the map level.

2.  Turn off *Show Map in Viewport (page 2–1445)*.

   The object is shaded but the map no longer appears.

To turn a map off:

- Turn off the map's check box.

  The check box is to the left of the map's name on the Maps rollout.

To turn a map on:

- Turn on the map's check box.

  The check box is to the left of the map's name on the Maps rollout.

To change a map's strength:

- Adjust the map's Amount spinner in the Maps rollout.

  The material's sample slot reflects the change.

  Note: Adjusting a map's output (in the map's Output rollout) can also change the map's strength.

To move directly to an ancestor:

1. Click the arrow to the right of the map's *name (page 2–1448)* field.

   A drop-down list of ancestors is displayed.

2. Click a name in the Ancestor list.

   With this list, you can skip intermediate levels in the tree.

   The Ancestor drop-down list shows only part of the tree. It does not show side branches and siblings. To view these, use the *Material/Map Navigator (page 2–1447)*.

    You can also use the *Go Forward to Sibling (page 2–1447)* and *Go to Parent (page 2–1446)* buttons.

To change a map type:

1. At the level of a map, click the button labeled Type below the Material Editor toolbar.

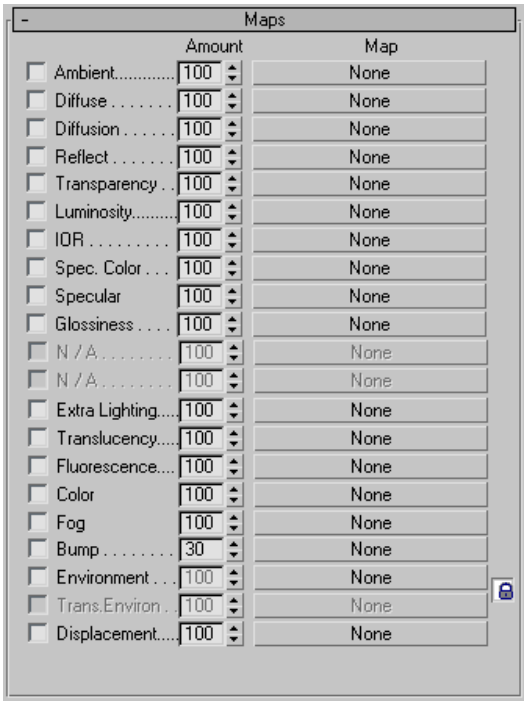
   A modal *Material/Map Browser (page 2–1412)* is displayed. If you were at a map, it lists only maps (if you were at a material when you clicked Type, the Browser lists only materials).

2. Choose a map type from the list, and then click OK.

   If you change a map type and the new map type can have component maps, a Replace Map dialog is displayed. This dialog gives you a choice between discarding the original map or using it as a component map.

   If the new map type does not have components, it simply replaces the original map type.

Interface



The Maps rollout for Raytrace contains a wide button for each map type. Click this button to select a bitmap file stored on disk or to select a *procedural-map type (page 3–997)*. After you select a map, its name and type appears on the button. Use the check box to the left of the button to turn the effect of the map off and on. When the check box is off, the map is not computed and has no effect in the renderer.

The Amount spinner determines the amount that the map affects the material, expressed as a percentage of full intensity. For example, a diffuse map at 100% is completely opaque and covers the base material. At 50%, it is semi-transparent and the base material (the diffuse, ambient, and other colors of the material without mapping) shows through.

Lock button—Locks the Environment map to the Transparency Environment map. When on, the Transparency Environment map controls are disabled, and a map applied to the Raytrace Environment applies to the Transparency Environment as well. When off, the Transparency Environment map controls are enabled, and the Transparency Environment can have a different map assigned to it. Default=on.

Changing this button's setting here also changes it on the *Basic Parameters rollout (page 2–1514)*, and the *Extended Parameters rollout (page 2–1519)*.

Diffusion Mapping

The Diffusion map component lets you apply an additional, second texture to modify the Diffuse component. Typically, you will want to reduce the Amount of this map to allow the main Diffuse map to show through.

For example, you might have a clean, bright image for a billboard. You use this image as the Diffuse map, and then use a second map as a Diffusion map to apply soot and city grime.

Tip: Animating the Amount can change the appearance of the material over time, letting the Diffusion map either conceal or reveal the underlying Diffuse map.

Note: Show Map In Viewport does not display the Diffusion map.

Raytrace Dynamics Properties Rollout

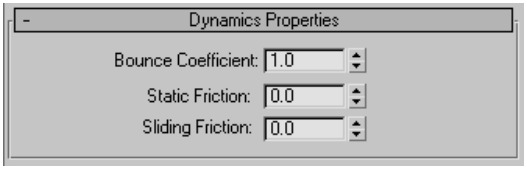
Material Editor > Type button > Material/Map Browser > Raytrace material > Dynamics Properties rollout

As with a standard material, the Dynamics Properties rollout for a *Raytrace material (page 2–1512)* lets you specify surface properties that affect the animation of an object upon collision with another object. If there are no collisions in your simulation, these settings have no effect.

Since the Dynamics Properties rollout is available at the top level of any material (including sub-materials), you can specify different surface dynamic properties for each face in an object. There are also controls in the Dynamics utility that let you adjust the surface properties at the object level, but only the Materials Editor lets you alter the surface properties at the sub-object level, through use of a *Multi/Sub-Object material (page 2–1594)*.

As a default, the values in the Dynamics Properties rollout provide a surface that's similar to Teflon-coated hardened steel.

Interface



Bounce Coefficient—Sets how far an object bounces after hitting a surface. The higher the value, the greater the bounce. A value of 1 represents a "perfectly elastic collision," or a bounce in which no kinetic energy is lost. Default=1.0.

If you've seen the desktop toy with four ball bearings swinging back and forth on strings and hitting one another, you've seen an example that comes very close to a bounce coefficient of 1. Generally, hardened steel or a super ball have a bounce near 1, while lead has a bounce near 0.

Static Friction—Sets how difficult it is for the object to start moving along a surface. The higher this value, the more difficult. Default=0.0.

If something weighs ten pounds and sits on Teflon (a static friction of near 0), it takes almost no force to make it move sideways. On the other hand, if it sits on sandpaper, then the static friction might be very high, on the order of 0.5 to 0.8. A static

friction near 1 is very difficult to create in the real world without adhesives or friction material.

Sliding Friction—Sets how difficult it is for the object to keep moving over a surface. The higher this value, the more difficult for the object to keep moving. Default=0.0.

Once two objects begin to slide over one another, static friction disappears and sliding friction takes over. Generally, sliding friction is lower than static friction due to surface tension effects. For example, once steel starts sliding over brass (a value of static friction that might run from 0.05 to 0.2), the sliding friction drops to a significantly lower value, on the order of .01 to 0.1. For some materials, such as specific friction materials like brake linings, sliding friction is just as high as static friction because it is used in conjunction with a nearly frictionless material such as hardened polished steel.

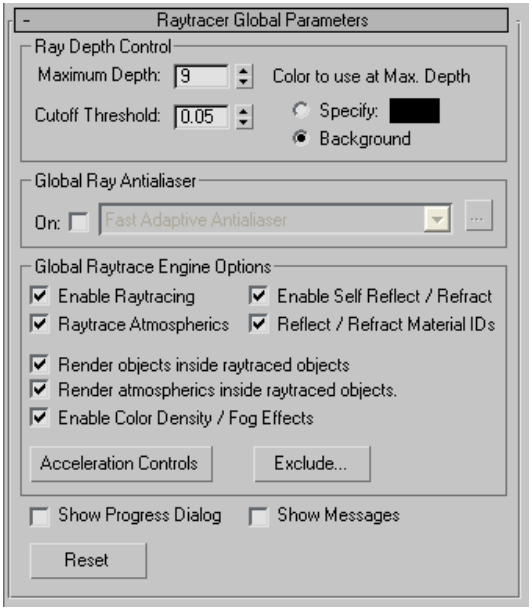
Raytracer Global Parameters Rollout

Rendering menu > Raytracer Settings
Rendering menu > Render > Render Scene dialog > Choose Default Scanline Renderer as the active production renderer. > Raytracer panel > Raytracer Global Parameters rollout
Main toolbar > Render Scene > Render Scene dialog > Choose Default Scanline Renderer as the active production renderer. > Raytracer panel > Raytracer Global Parameters rollout

Parameters in the Global Raytracer Settings dialog globally control the raytracer itself. That is, they affect *all* Raytrace materials and Raytrace maps in your scene. They also affect the generation of *Advanced Ray-traced shadows (page 3–1000)* and *Area shadows (page 3–911)*.

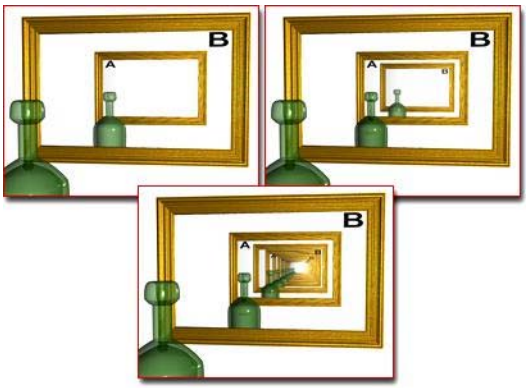
Note: These controls adjust ray-trace settings for the scanline renderer only. The settings of these controls have no impact on the mental ray renderer, which has its own ray-tracing controls.

Interface



Ray Depth Control group

Ray depth is also known as recursion depth. This is how many times a ray is allowed to bounce before it is considered lost or trapped.



Upper left: Ray depth is zero

Upper right: Ray depth of 2

Lower middle: Extremely high ray depth

Maximum Depth—Sets the maximum recursion depth. Increasing this value potentially increases the realism of your rendered scene, at a cost of rendering time. You can reduce this value to reduce rendering time. Range=0 to 100. Default=9.

Cutoff Threshold—Sets a cutoff threshold for adaptive ray levels. If the contribution of any ray to the final pixel color drops below the cutoff threshold, the ray is terminated. Default: 0.05 (5% of the final pixel color). This can speed up your rendering time considerably.

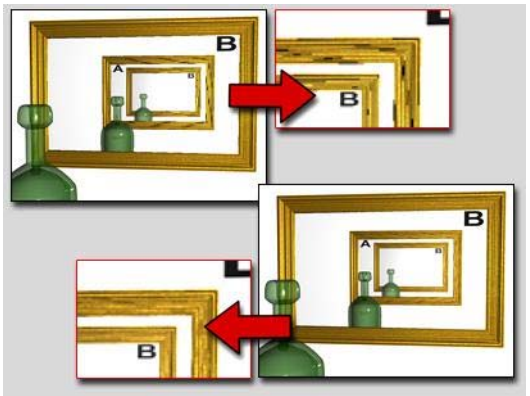
Color to use at Max Depth—As a rule, when a ray reaches the maximum depth, it is rendered the same color as the background environment. You can override the color returned at maximum depth by either selecting a color, or setting an alternative environment map. This can make the "lost" ray invisible in the scene.

Tip: If you are having trouble with getting complex objects to render, especially glass, specify the maximum recursion color to something obvious, like magenta, and your background color to something that contrasts, like cyan. The chances are that a lot of your rays are getting lost in either maximum recursion or just being shot off into the world, totally missing anything you think they should strike. Try rendering the scene again. If this is the problem, try reducing the Maximum Depth value.

- **Specify—**Specifies what color the raytracer returns when the ray is considered lost or trapped. Click the color swatch to change this color.
- **Background—**(The default.) Returns the background color when the ray is considered lost or trapped. For Raytrace material, the background color is the global environment background or the environment specified locally for the material. For Raytrace map, the background color is either the global environment background, or is set locally in the *Raytracer Parameters rollout (page 2–1704)*.

Global Ray Antialiaser group

Controls in this group let you set global antialiasing for raytraced maps and materials.



Above: No antialiasing

Below: Antialiasing of reflections

Tip: Turning on Supersample for a Raytraced material (in the *Raytrace Basic Parameters rollout (page 2–1514)*) usually provides adequate antialiasing. Use one of the raytrace antialiasers (Fast Adaptive or Multiresolution Adaptive) when you want to blur reflections or refractions.

On—When on, uses antialiasing. Default=off.

Drop-down list—Chooses which antialiasing settings to use. There are two alternatives:

- **Fast Adaptive Antialiaser—**(The default.) Uses the Fast Adaptive antialiaser.
- **Multiresolution Adaptive Antialiaser—**Uses the Multiresolution Adaptive antialiaser.

...—The button with the ellipsis to the right of the drop-down list displays another dialog to let you set antialiasing controls globally. The dialog displayed depends on which alternative you chose in the drop-down list, as follows:

- **Fast Adaptive Antialiaser—**Clicking ... displays the *Fast Adaptive Antialiaser dialog (page 2–1533)*.
- **Multiresolution Adaptive Antialiaser—**Clicking ... displays the *Multiresolution Adaptive Antialiaser dialog (page 2–1534)*.

Global Raytrace Engine Options group

These options are comparable to the local options on *Extended Parameters rollout (page 2–1519)* and the *Raytracer Controls rollout (page 2–1521)*. Their setting affects all Raytrace materials and Raytrace maps in the scene, unless you set local overrides.

Note: In versions prior to 3ds Max **5**, these controls were found on the Raytracer Options dialog, which is no longer present.

Enable Raytracing—Turns the raytracer on or off. Default=on.

Even with raytracing off, Raytrace material and Raytrace map still reflect and refract the environment, including both the environment map for the scene, and the environment map assigned to the Raytrace material.

Raytrace Atmospherics—Turns the raytracing of atmospheric effects on or off. Atmospheric effects include fire, fog, volume light, and so on. Default=on.

Enable Self Reflect/Refract—Turns self reflection/refraction on or off. Default=on.

Can an object reflect itself? For example, a teapot's body reflects the teapot's handle, but a sphere will never reflect itself. If you don't need this effect, you can improve render time by turning off this toggle.

Tip: If you have a transparent object such as glass, and Enable Self Reflect/Refract is on, you don't have to make the object *2-sided (page 3–901)*. The raytracer sees back faces when exiting refractive objects.

Reflect/Refract Material IDs—When on, the material reflects effects assigned to material IDs in the renderer's *G-buffer (page 3–946)* on or off. Default=on.

By default, Raytrace material and Raytrace map reflect effects assigned to a material's ID, so that G-buffer effects are not lost. For example, if a raytraced object reflects a lamp made to glow with the Video Post Glow filter (Lens Effects Glow), the reflection glows as well.

Render objects inside raytraced objects—Turns the rendering of objects inside raytraced objects on or off. Default=on.

Render atmospherics inside raytraced objects—Turns the rendering of atmospheric effects inside raytraced objects on or off. Atmospheric effects include fire, fog, volume light, and so on. Default=on.

Enable Color Density / Fog Effects—Turns the color density and fog features on or off.

Acceleration Controls—Click to display the *Raytracing Acceleration Parameters dialog (page 2–1531)*, which lets you adjust settings.

Exclude—Click to display the *Raytrace Exclude/Include dialog (page 2–1531)*, which lets you exclude objects from ray-tracing.

Show Progress Dialog—When on, rendering displays a window with progress bars titled Raytrace Engine Setup. Default=on.

Show Messages—When on, displays a window, Raytrace Messages, that shows status and progress messages from the raytrace engine. Default=off.

Raytracing Acceleration Parameters Dialog

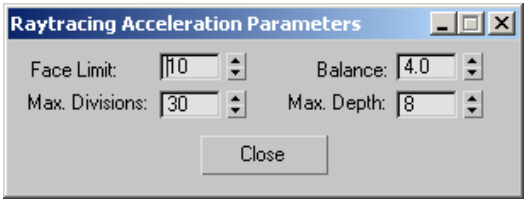
Rendering menu > Raytracer Settings > Render Scene dialog > Raytracer panel > Raytracer Global Parameters rollout > Global Raytrace Engine Options group > Acceleration Controls button

The controls in this dialog let you override the default acceleration values and specify your own requirements. Typically you don't need to use them, but if you are familiar with the requirements of your scene, they can help you optimize raytrace rendering for your specific needs and time constraints.

Raytracing subdivides the scene, organizing it into a tree for raytrace purposes. A node in this tree is known as a "voxel." Voxel trees are dynamic, and you can't explicitly specify the structure of the tree. (If you set Max. Divisions to 2, the tree is an octree, which is possibly a more familiar data structure.)

Tip: If your scene is very unbalanced (for example, it has a giant planet object and a cluster of small spacecraft objects) keep the Balance setting low.

Interface



Face Limit—Sets the maximum number of faces allowed in a lattice voxel before it is subdivided. Default=10.

Balance—Determines the sensitivity of the subdivision algorithm. Increasing this value uses more memory but can increase performance. Default=4.0.

Max. Divisions—Sets the initial lattice dimension. For example, 4 is a 4 x 4 x 4 lattice. Default=30.

Max. Depth—Sets the maximum number of lattice subdivisions. Default=8.

Raytrace Exclude/Include Dialog

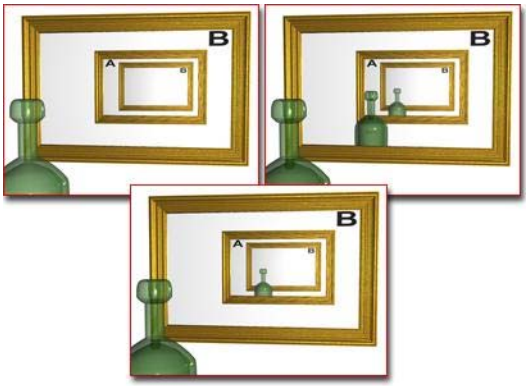
Rendering menu > Raytrace Global Include/Exclude

Material Editor > Raytrace material > Raytracer Controls rollout > Local Exclude button

Material Editor > Raytrace map > Raytracer Parameters rollout > Local Exclude button

Rendering menu > Raytracer Settings > Render Scene dialog > Raytracer panel > Raytracer Global Parameters rollout > Global Raytrace Engine Options group > Exclude button

Main toolbar > Render Scene > Render Scene dialog > Raytracer panel > Raytracer Global Parameters rollout > Global Raytrace Engine Options group > Exclude button



Upper left: Bottle excluded from both mirrors

Upper right: Bottle reflected by both mirrors

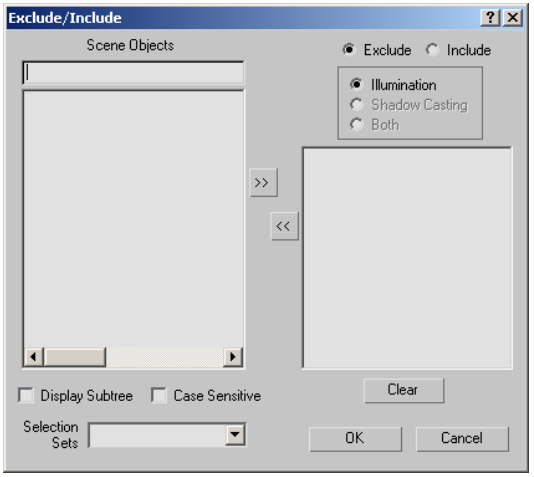
Lower middle: Bottle excluded from one mirror

These dialogs let you specify which objects the raytracer will and won't see. This is a good optimization when you are working with complex scenes and all you really need is for an object to reflect itself or some other, simpler elements in the scene.

The global dialog affects *all* Raytrace materials and Raytrace maps in your scene. The local dialog affects only the current material or map.

Tip: Another good use for exclusion can be when you are working on a scene where the world-space scale of objects is not realistic. For example, consider an animation of a space dog-fight, with a planet and its moon in the background. You would use a planet that was actually much smaller, relative to the fighters, than an actual planet. (For example, the planet might be only five times bigger than the fighters, and the moon only half as big as the fighters.) If the fighters are reflective, the reflections of the planet and moon will be incorrect and will give away the fact that you are using miniatures. In this case, raytrace the reflective fighters, and laser blasts, and so on, but exclude the planet and moon props from the raytracer. Create an environment map that represents the planet and moon at their correct scale, and make that the Raytrace material's local environment. When you render, the scale of the planet and moon geometry appears to be correct, and the reflections on the fighters will behave as your eye expects.

Interface



Both the Global and the Local Exclude/Include dialogs contain the following controls:

Exclude/Include—Choose whether raytracing will exclude or include the objects named in the list on the right.

Illumination / Shadow Casting / Both—For the raytracer, this is turned off (always set to Illumination).

Scene Objects—Select objects from the Scene Objects list on the left, then use the arrow buttons to add them to the exclusion list on the right.

The Exclude/Include dialog treats a group as an object: you can exclude or include all objects in a group by selecting the group's name in the Scene Objects list. If a group is nested within another group, it isn't visible in the Scene Objects list.

Search Field—The edit box above the Scene Objects list lets you search for object names by entering names that use wildcards.

Display Subtree—When on, indents the list according to the object hierarchy.

Case Sensitive—When on, uses case sensitivity when searching object names.

Selection Sets—Displays a list of named selection sets. Choosing a selection set from this list selects those objects in the Scene Objects list.

Clear—Clears all entries from the Exclude/Include list on the right.

OK—Closes the dialog and retains the exclude or include operations you performed.

Cancel—Closes the dialog and cancels the exclude or include operations you performed.

Raytrace Antialiaser Dialog: Fast Adaptive Antialiaser

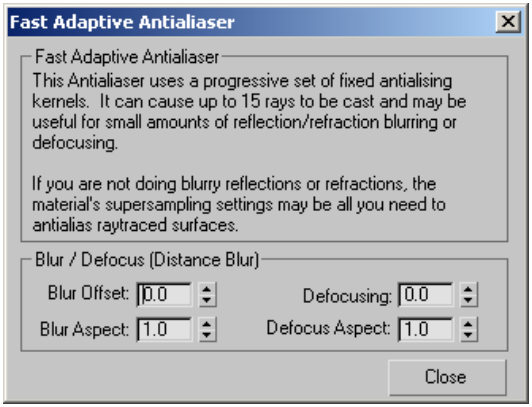
Rendering menu > Raytracer Settings > Global Raytracer Settings dialog > Global Ray Antialiaser group > Turn on global antialiasing. > Choose Fast Adaptive Antialiaser from the drop-down list. > ... button

Material Editor > Type button > Material/Map Browser > Raytrace material > Raytracer Controls rollout > Raytraced Reflection and Refraction Antialiaser group (enabled only if antialiasing is globally enabled) > Choose Fast Adaptive Antialiaser from the drop-down list. > ... button

Material Editor > Type button > Material/Map Browser > Raytrace map > Raytracer Parameters rollout > Raytraced Reflection and Refraction Antialiaser group (enabled only if antialiasing is globally enabled) > Choose Fast Adaptive Antialiaser from the drop-down list. > ... button

The Fast Adaptive Antialiaser dialog changes settings for the Raytrace material and maps Fast Adaptive antialiaser. You can use this dialog either globally, from the *Global Raytracer Settings dialog (page 2–1528)*, or locally, from the *Raytracer Controls rollout (page 2–1521)*. When you locally change settings for an antialiaser, you don't affect the global settings for that antialiaser.

Interface



Blur / Defocus (Distance Blur) group

Blur Offset is similar to blur offset for Bitmaps, while defocus is based on distance.

Blur Offset—Affects the sharpness or blurriness of the reflections or refractions without regard to distance. You can use Blur Offset to soften or defocus the details of a reflection or refraction. The value is specified in pixels. Default=0.0.

Tip: The default Blur Offset setting usually produces good results. If you see aliasing in reflections or refractions, increase its value in increments of 0.5 until the aliasing goes away.

See *Blur/Blur Offset (page 3–918)*.

Blur Aspect—This is an aspect ratio that changes the shape of the blur. Usually you will not need to change it. Default=1.0.

Tip: If you see aliasing that occurs mostly along horizontal lines, try increasing Blur Aspect to 1.5. This changes the shape of the blurred effect. The reverse is also true. If aliasing occurs mostly along vertical lines, try decreasing Blur Aspect to 0.5.

Defocusing—Defocusing is a blur based on distance. With Defocus, objects near the surface are not blurred, but objects farther away are blurred. The rays cast are spread as they leave the Raytrace material object's surface. Default=0.0.

Tip: Increasing the value of Defocusing can give a good distance blurring effect. Small adjustments are usually adequate. Try starting with a value less than 0.1, and increase or decrease it as necessary. Also try adjusting Reflect Falloff in Raytrace material or Attenuation in Raytrace map to get the best distance blurring effect.

Defocus Aspect—This is an aspect ratio that changes the shape of the defocusing. Usually you will not need to change it. Default=1.0.

Raytrace Antialiaser Dialog: Multiresolution Adaptive Antialiaser

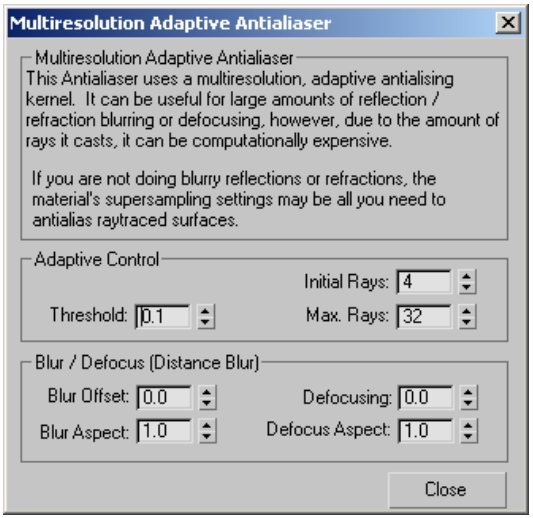
Rendering menu > Raytracer Settings > Render Scene dialog > Raytracer panel > Raytracer Global Parameters rollout > Raytracer Global Parameters > Global Ray Antialiaser group > Turn on global antialiasing. > Choose Multiresolution Adaptive Antialiaser from the drop-down list. > ... button

Material Editor > Type button > Material/Map Browser > Raytrace material > Raytracer Controls rollout > Raytraced Reflection and Refraction Antialiaser group (enabled only if antialiasing is globally enabled) > Choose Multiresolution Adaptive Antialiaser from the drop-down list. > ... button

Material Editor > Type button > Material/Map Browser > Raytrace map > Raytracer Parameters rollout > Raytraced Reflection and Refraction Antialiaser group (enabled only if antialiasing is globally enabled) > Choose Multiresolution Adaptive Antialiaser from the drop-down list. > ... button

The Multiresolution Adaptive Antialiaser dialog changes settings for the Raytrace material and maps Multiresolution Adaptive antialiaser. You can use this dialog either globally, from the *Global Raytracer Settings dialog (page 2–1528)*, or locally, from the *Raytracer Controls rollout (page 2–1521)*. When you locally change settings for an antialiaser, you don't affect the global settings for that antialiaser.

Interface



Adaptive Control group

Initial Rays—Sets the initial number of rays cast per pixel. Default=4.

Threshold—Determines the sensitivity of the adaptation algorithm. It can range from 0 to 1, where 0 always casts the maximum number of rays and 1 always casts only the minimum number of rays. Default=0.1.

Max. Rays (Maximum Rays)—Sets the maximum number of rays the algorithm will cast. Default=32.

Blur / Defocus (Distance Blur) group

Blur Offset is similar to *blur offset (page 3–918)* for bitmaps, while Defocusing is based on distance.

Blur Offset—Affects the sharpness or blurriness of the reflections or refractions without regard to distance. You can use Blur Offset to soften or defocus the details of a reflection or refraction. The value is specified in pixels. Default=0.0.

Tip: The default Blur Offset setting usually produces good results. If you see aliasing in reflections or refractions, increase its value in increments of 0.5 until the aliasing goes away.

Blur Aspect—This is an aspect ratio that changes the shape of the blur. Usually you will not need to change it. Default=1.0.

Tip: If you see aliasing that occurs mostly along horizontal lines, try increasing Blur Aspect to 1.5. This changes the shape of the blurred effect. The reverse is also true. If aliasing occurs mostly along vertical lines, try decreasing Blur Aspect to 0.5.

Defocusing—Defocusing is a blur based on distance. With Defocus, objects near the surface are not blurred, but objects farther away are blurred. The rays cast are spread as they leave the Raytrace material object's surface. Default=0.0.

Tip: Increasing the value of Defocusing can give a good distance blurring effect. Small adjustments are usually adequate. Try starting with a value less than 0.1, and increase or decrease it as necessary. Also try adjusting Reflect Falloff in Raytrace material or Attenuation in Raytrace map to get the best distance blurring effect.

Defocus Aspect—This is an aspect ratio that changes the shape of the defocusing. Usually you will not need to change it. Default=1.0.

Architectural Material

The settings for an Architectural material are physical properties, so it provides the greatest possible realism when used with *photometric lights (page 2–1301)* and *radiosity (page 3–51)*. With this combination of features, you can create lighting studies with a high degree of accuracy.

It is *not recommended* that you use the Architectural material with standard 3ds Max lights in the scene, or with the Light Tracer. The point of this material is to provide accurate modeling. Use it with photometric lights and radiosity. The mental ray renderer, on the other hand, can render the Architectural material, with some limitations described below.

Tip: If you don't need the degree of realism that the Architectural material provides, you can use a *standard material (page 2–1465)* or other material type.

Material Templates

When you create a new material, you can choose from a variety of templates. A template is simply a set of preset material parameters, which approximates the kind of material you want to create, and gives you a starting point. See *Templates Rollout (page 2–1536)*.

Rendering Architectural Materials with the mental ray Renderer

The *mental ray Renderer (page 3–78)* can render Architectural materials. There are some limitations, as follows:

- Emit Energy (Based on Luminance): This setting is ignored. The Architectural material does not contribute to the scene's lighting.
- Sampling Parameters: These settings are ignored, as the mental ray renderer uses its own sampling.

Tip: When rendering with mental ray, instead of the Architectural material, we *highly* recommend that you use the *Arch & Design material (page 2–1549)* instead. This material was designed especially for mental ray and provides superior flexibility, rendering characteristics, and speed.

Interface

Templates Rollout (page 2–1536)

Physical Qualities Rollout (page 2–1536)

Special Effects Rollout (page 2–1539)

Advanced Lighting Override Rollout (page 2–1540)

SuperSampling Rollout (page 2–1459)

mental ray Connection Rollout (page 2–1461)

Templates Rollout

Material Editor > Architectural material > Templates rollout

The Templates rollout gives you a list of material types to choose from. A template is simply a set of preset parameters for the Physical Qualities rollout, which approximates the kind of material you want to create, and gives you a starting point. Once you choose a template, you can adjust its settings and add *maps (page 2–1617)* to enhance realism and improve the material's appearance.

The templates do not affect the Diffuse Color on the Physical Qualities rollout, only the numeric settings.

Interface



Template drop-down list—Chooses the kind of material you are designing. Each template provides preset values for the various material parameters.

These are the material templates provided with 3ds Max. The purpose of most templates is clear, so the table doesn't comment on all of them.

Template	Comments
Ceramic Tile - Glazed	
Fabric	
Glass - Clear	
Glass - Translucent	
Ideal Diffuse	A neutral white material
Masonry	A good base for a diffuse map
Metal	Shiny and reflective
Metal - Brushed	Less shiny
Metal - Flat	Even less shiny
Metal - Polished	Highly shiny
Mirror	Completely shiny
Paint Flat	Another neutral white material
Paint Gloss	Also white, but shiny
Paint Semi-Gloss	Also white, only slightly shiny
Paper	
Paper - Translucent	
Plastic	
Stone	A good base for a diffuse map
Stone Polished	Has a bit of shininess; also a good base for a diffuse map
User Defined	Neutral; a good base for a diffuse map
User-Defined Metal	Somewhat shiny; also a good base for a diffuse map
Water	Completely clear and shiny
Wood Unfinished	Neutral; a good base for a map
Wood Varnished	

Physical Qualities Rollout

Material Editor > Architectural material > Physical Qualities rollout

When you create a new *Architectural material (page 2–1535)* or edit an existing one, the settings on the Physical Qualities rollout are the ones you are most likely to need to adjust.

Procedure

To match a material's luminance to a light:

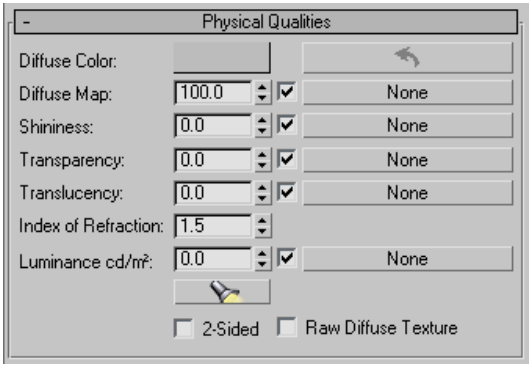
Warning: Only photometric lights give correct luminance. Also, the mental ray renderer disregards the Emit Energy setting.

- Click to turn on Set Luminance From Light (below the Luminance setting), then in a viewport, click the light.

  After you choose the light, the button turns off once more.

  Tip: If you are using a *radiosity solution (page 3–51)*, make sure to turn on Emit Energy (Based On Luminance) for any material whose luminance is greater than zero. This control is on the *Advanced Lighting Override rollout (page 2–1540)*.

Interface



Diffuse Color—Controls the *diffuse color (page 3–929)*. The diffuse color is the color this material has in direct light. Click the color swatch to display the *Color Selector (page 1–161)* and change the diffuse color.

Set color to texture average—Click to change the diffuse color to an average of the colors in the current diffuse map. (If no map is assigned, this button has no effect.)

This button is useful when you are going to reduce the diffuse map's Amount. When the diffuse map pattern appears over an average of itself, for most kinds of materials the effect is more realistic than when the pattern appears over an unrelated color.

Warning: This button will change the diffuse color even if the diffuse map is turned off.

Diffuse Map—These controls assign a map to the material's diffuse component. To assign a map, click the oblong button (labeled "None" by default). This displays the Material/Map Browser. In the Browser, choose the map type, and then click OK. If you choose Bitmap as the map type, an additional dialog prompts you to choose the particular bitmap file to use.

While a map is assigned to the material, its name appears as the map button's label.



- **Amount spinner—**The spinner at the left sets the amount of diffuse map to use. This value is a percentage: at 100.0, only the map is visible; at lower amounts, the diffuse color shows through; at 0.0, the map is not visible at all.
- **On/off—**The check box between the spinner and the map button is an on/off switch. When on, the map appears in the material. When off, the map does not appear.

Shininess—Sets the shininess of the material. This value is a percentage: at 100.0, the material is as shiny as possible; at lower values, it is less shiny; at 0.0, it is not shiny at all.

In general, the shinier a material is, the smaller its specular highlights appear. Specular highlights are reflections of the lights that illuminate the material. (The index of refraction can also affect the size of highlights.) Shininess also controls how much the material reflects other objects in the scene.

Note: Shininess alone is not sufficient to produce specular reflections and highlights on a surface. You must also consider the Index Of Refraction (IOR), described below.

The amount spinner for a shininess map scales the map. If no map is assigned, the spinner value is used by itself.

Transparency—Controls how transparent the material is. This value is a percentage: at 100.0, the material is completely transparent; at lower values, the material is partly opaque; and at 0.0, the material is completely opaque.

Tip: The effect of transparency is best previewed against a pattern background. If the material preview doesn't show a pattern of colored checks, right-click the material preview or the map preview, and choose Background from the pop-up menu.

The amount spinner for a transparency map scales the map. If no map is assigned, the spinner value is used by itself.

Translucency—Controls how translucent the material is. A translucent object transmits light, but also scatters it within the object. This value is a percentage: at 0.0, the material is completely opaque; at 100.0, the material is as translucent as possible.

The amount spinner for a translucency map scales the map. If no map is assigned, the spinner value is used by itself.

Index of Refraction—The index of refraction (IOR) controls how severely the material refracts transmitted light, and how reflective the material appears. At 1.0, the IOR of air, the object behind the transparent object does not distort. At 1.5, the object behind distorts greatly, like a glass marble. Range=1.0 to 2.5.

Common IORs (assuming the camera is in air or a vacuum) are:

Material	IOR Value
Vacuum	1.0 (exactly)
Air	1.0003
Water	1.333
Glass	1.5 to 1.7
Diamond	2.419

In the physical world, the IOR results from the relative speeds of light through the transparent material and the medium the eye or the camera is in. Typically this is related to the object's density: the higher the IOR, the denser the object.

The IOR affects how shiny a material appears; or, in the case of transparent materials such as water or glass, the amount of distortion. For nontransparent materials, the higher the IOR, the more light is reflected from the material, and the shinier the material appears.

A refractive index of 1.0 means that all light is transmitted into the material. In this case, even if the material has a high Shininess value, the surface appears perfectly diffuse, and shows no specular highlights.

Luminance cd/m2—When its luminance is greater than 0.0, the material appears to glow, and if you turn on Emit Energy (see below), it contributes energy to the *radiosity solution (page 3–51)*. Luminance is measured in candelas per meter squared.

The amount spinner for a luminance map scales the map. If no map is assigned, the spinner value is used by itself.



Set luminance from light—Obtains the material's luminance from a light in the scene. Click to turn on this button, then click a light in a viewport. The material's luminance is set to match the luminance of the light, and the button is turned off once more.

2-Sided—When on, makes the material *2-sided (page 3–901)*. Applies the material to both sides of selected faces.

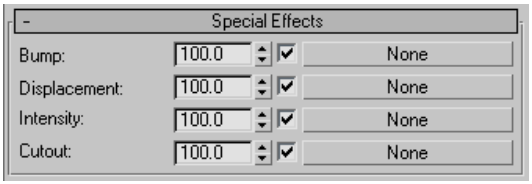
Raw Diffuse Texture—When on, excludes the material from lighting and exposure control. This makes the material render with a completely flat look, using the pure RGB values of the diffuse color or map. Default=off.

Special Effects Rollout

Material Editor > Architectural material > Special Effects rollout

When you create a new *Architectural material (page 2–1535)* or edit an existing one, the settings on the Special Effects rollout let you assign maps that create bumps or displacement, adjust light intensity, or control transparency.

Interface



Bump controls—These controls assign a *bump map (page 2–1506)* to the material. To assign a map, click the oblong button (labeled "None" by default). This displays the Material/Map Browser. In the Browser, choose the map type, and then click OK. If you choose Bitmap as the map type, an additional dialog prompts you to choose the particular bitmap file to use.

While a map is assigned to the material, its name appears as the map button's label.



- **Amount spinner—**The spinner at the left sets the amount of bump mapping to use. At 1000.0, bump mapping has its greatest effect; at lower amounts, the bumps are less pronounced; at 0.0, the bumps are not visible at all. Negative values reverse the direction of the bump effect. Range: –1000.0 to 1000.0.
- **On/off—**The check box between the spinner and the map button is an on/off switch. When on, the map is used in the material. When off, the map is not used.

Displacement controls—These controls assign a *displacement map (page 2–1511)* to the material. To assign a map, click the oblong button (labeled "None" by default). This displays the Material/Map Browser. In the Browser, choose the map type, and then click OK. If you choose Bitmap as the map type, an additional dialog prompts you to choose the particular bitmap file to use.

While a map is assigned to the material, its name appears as the map button's label.



- **Amount spinner—**The spinner at the left sets the amount of displacement mapping to use. At 1000.0, displacement mapping has its greatest effect; at lower amounts, the displacement is less pronounced; at 0.0, the displacement are not visible at all. Negative values reverse the direction of the displacement. Range: –1000.0 to 1000.0.
- **On/off—**The check box between the spinner and the map button is an on/off switch. When on, the map is used in the material. When off, the map is not used.

Intensity controls—These controls assign an intensity map to the material, modulating the material's brightness. The map is treated as a black-and-white scale of intensity values.

Tip: Applying a low-frequency *Noise map (page 2–1674)* to the intensity can help reduce the "computer-generated" look of a texture, and add a natural feel to surfaces such as bricks and carpets lit by daylight.

To assign a map, click the oblong button (labeled "None" by default). This displays the Material/Map Browser. In the Browser, choose the map type, and then click OK. If you choose Bitmap as the map type, an additional dialog prompts you to choose the particular bitmap file to use.

While a map is assigned to the material, its name appears as the map button's label.



- **Amount spinner—**The spinner at the left sets the amount of intensity mapping to use. At 100.0, intensity mapping has its greatest effect; at lower amounts, the effect is less pronounced; at 0.0, the map has no effect; and at values less than 0.0, the material is dimmed. Range: –100.0 to 100.0.
- **On/off—**The check box between the spinner and the map button is an on/off switch. When on, the map is used in the material. When off, the map is not used.

Cutout controls—These controls assign a *cutout map (page 2–1542)* to the material. To assign a map, click the oblong button (labeled "None" by default). This displays the Material/Map Browser. In the Browser, choose the map type, and then click OK. If you choose Bitmap as the map type, an additional dialog prompts you to choose the particular bitmap file to use.

While a map is assigned to the material, its name appears as the map button's label.



- **Amount spinner—**The spinner at the left sets the amount of cutout mapping to use. This value is a percentage: at 100.0, the map has its full effect; lower amounts have the effect of darkening the map, reducing its effect and increasing transparency. At 0.0, the map is completely black, making objects with this material completely transparent.
- **On/off—**The check box between the spinner and the map button is an on/off switch. When on, the map is used in the material. When off, the map is not used.

Advanced Lighting Override Rollout

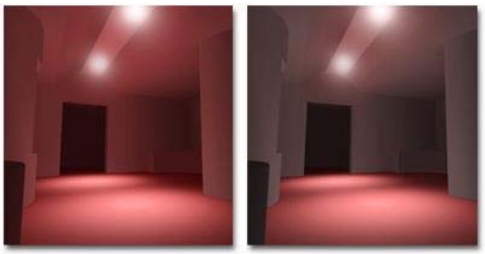
Material Editor > Architectural material > Advanced Lighting Override rollout

When you create a new *Architectural material (page 2–1535)* or edit an existing one, the settings on the Adjust Radiosity rollout let you adjust how the material behaves in a *radiosity solution (page 3–51)*.

Obtaining a Better Image

Materials with a bright diffuse color or high shininess can be highly reflective. This can lead to overexposed or washed-out radiosity solutions. In general, the best way to adjust this is to reduce the HSV Value (V) of a material's diffuse color; or, for a material with a diffuse map, reduce the map's RGB level. In some situations, the controls on this rollout can improve the appearance of the radiosity solution. Examples of situations where the material's radiosity settings can help include color bleeding and large dark areas:

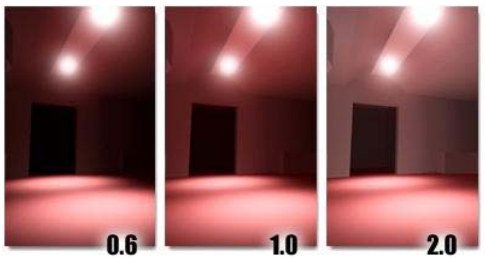
- You might want to reduce Reflectance Scale or Color Bleed Scale when a large area of color (for example, a red carpet in a room with white walls) creates excessive color bleeding. This might be physically accurate, but the eye adjusts for such effects, and the radiosity result might look better with less reflectance or less color bleeding.



Left: Excessive bleeding of the floor color onto the walls and ceiling.

Right: Reducing the floor's Reflectance Scale causes less bleeding.

- You might want to increase Reflectance Scale when the scene includes a large dark area (for example, a black floor). This can lead to a very dark radiosity result. You can maintain the floor's color but increase reflectance, giving the solution the colors you want while increasing its brightness.

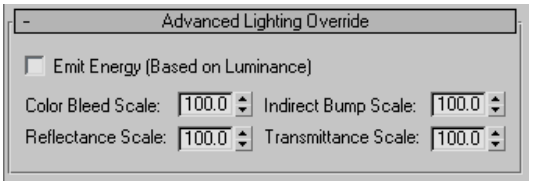


The room is lit only by spotlights pointed at the floor. Increasing reflectance of the floor brightens the entire room.

Tip: Check the *reflectance and transmittance display (page 2–1430)* to get an idea of how the current material will affect the radiosity solution.

Interface

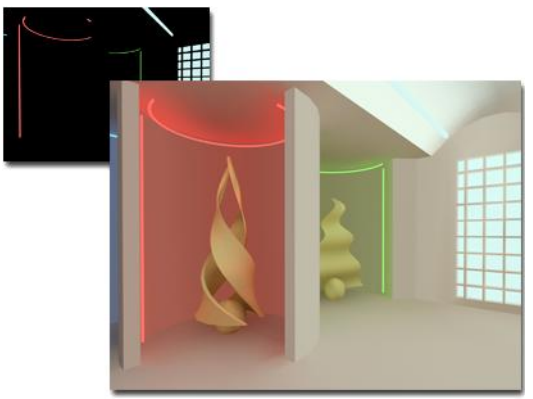
Warning: There is no problem with reducing the default scale, but increasing it for any of these parameters might cause colors to "burn out": if the scale is too great, they render as pure white, appearing overexposed.



Emit Energy (Based on Luminance)—When on, the material contributes energy to the radiosity solution, based on the material's luminance value (see above).

Note: The *mental ray renderer (page 3–78)* does not use this setting. The Architectural material does not contribute to the scene's lighting.

Increasing the Luminance (above 0.0) makes an object appear to glow in ordinary renderings, but does not contribute energy to the radiosity solution. To have radiosity processing take a self-illuminating material into account, turn on Emit Energy (Based On Luminance).



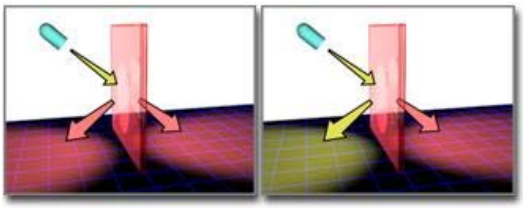
Upper left: By default, luminous neon lights do not influence the scene light.

Right: With Emit Energy on, the radiosity solution takes luminance into account.

Tip: When you increase luminance to achieve a special effect in the rendering (for example, to make the globe surrounding a lamp appear to be glowing), probably you shouldn't turn on Emit Energy (in the example, both the globe and lamp would then add light to the scene). When you

increase luminance because the object really glows (for example a neon light tube), then you should turn on Emit Energy, so that the object contributes light to the scene.

Color Bleed Scale—Increases or decreases the saturation of reflected color. Range=0.0 to 100.0. Default=100.0.

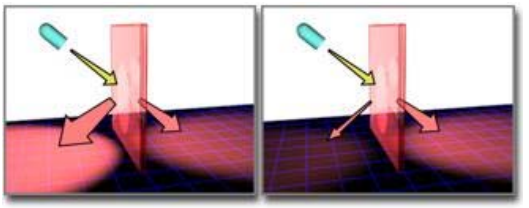


Color Bleed increases or decreases the saturation of reflected color.

Indirect Bump Scale—Scales the effect of the base material's *bump mapping (page 2–1506)* in areas lit by indirect light. When this value is zero, no bump mapping is done for indirect light. Increasing Indirect Light Bump Scale increases the bump effect under indirect lighting. This value does not affect the Bump amount in areas where the base material is lit directly. Range=–999.0 to 999.0. Default=100.0.

Tip: This parameter is useful because indirect bump mapping is simulated and not always accurate. Indirect Light Bump Scale lets you adjust the effect by hand.

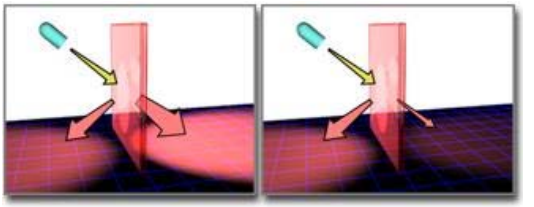
Reflectance Scale—Increases or decreases the amount of energy the material reflects. Range=0.0 to 100.0. Default=100.0.



Reflectance Scale increases or decreases the energy of reflected rays.

Tip: Don't use this control to increase self-illumination. Use the material's Luminance instead. The Luminance control is on the *Physical Qualities rollout (page 2–1536)*.

Transmittance Scale—Increases or decreases the amount of energy the material transmits. Range=0.1 to 5.0. Default=1.0.



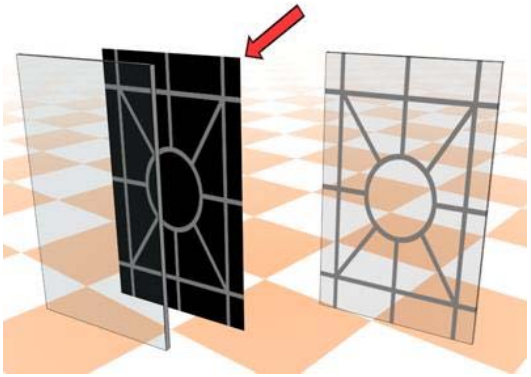
Transmittance Scale increases or decreases the energy of transmitted rays.

Cutout Mapping

Material Editor > Architectural material > Special Effects rollout > Cutout map button

Assigning a *bitmap (page 2–1631)* or *procedural map (page 3–997)* to the Cutout component of an *Architectural material (page 2–1535)* makes the material partially transparent. Lighter (higher-value) areas of the map render as opaque; darker areas render as transparent; and values in between are semi-transparent.

Setting the cutout map's Amount to 100 applies all of the map. Transparent areas are fully transparent. Setting the Amount to 0 is the equivalent of turning the map off. Intermediate Amount values are blended with the Transparency value on the Physical Qualities rollout. Transparent areas of the map become more opaque.



The gray levels of a cutout map determine the amount of transparency.

Procedure

To use a cutout map:

1. Click the Cutout map button.

   The *Material/Map Browser (page 2–1412)* is displayed.

2. Choose from the list of *map types (page 2–1617)*, and then click OK.

   The Material is now at the map level, and displays map controls.

   (If you choose Bitmap as the map type, you first see a file dialog that lets you choose the image file.)

3. Use the map controls to set up the map.

To remove a cutout map from a material:

Tip: You can disable the map without removing it. Simply turn off the toggle immediately to the left of the map button on the Special Effects rollout.

1. If the Material Editor is displaying the map controls, click the *Type button (page 2–1449)* to display the Material/Map Browser. If the map controls aren't visible, click the Cutout map button to display them, and then click the Type button.
2. In the Browser, choose NONE as the map type, and then click OK.

   The map is removed.

mental ray Materials

3ds Max comes with several materials created specifically for use with the *mental ray renderer (page 3–78)*. These materials are visible in the *Material/Map Browser (page 2–1412)* when mental ray is the active renderer.

- *mental ray (page 2–1544)*

  The mental ray material has components for the surface shader, and for the other nine optional shaders that make up a material in mental ray.

- *Arch & Design (page 2–1562)*

  The Arch & Design material is a monolithic material shader designed to support most materials used in architectural and product-design renderings.

- *Car Paint (page 2–1576)*

  Car Paint, available as both a material and a shader, has components for a paint layer with embedded metal flakes, a clear-coat layer, and a Lambertian dirt layer.

- *DGS (page 2–1580)*

  DGS stands for Diffuse, Glossy, and Specular. This material behaves in a physically realistic way.

- *Glass (page 2–1582)*

  The Glass material simulates both the surface properties and the light-transmitting (photon) properties of glass.

- *Subsurface Scattering Materials (page 2–1583)*

The subsurface scattering materials, supported by a shader library from mental images, can model skin and similar organic materials.

Note: When you wire the parameters of an object with a mental ray material assigned, names of material parameters might differ from those in the Material Editor interface. Also, parameters not supported by 3ds Max might appear as blanks in the wiring menu.

mental ray Material

Material Editor > Type button > Material/Map Browser > mental ray

Note: The mental ray material appears in the Browser only if the mental ray renderer is the currently active renderer.

The mental ray material lets you create a material exclusively for use by the *mental ray renderer (page 3–78)*. A mental ray material consists, at the top level, of from one to 10 kinds of shaders, or shading components.

Important: You must assign a shader to the material's Surface component. Otherwise, the mental ray material will not be visible when you render.

Interface

The interface to the mental ray material consists of two rollouts:

- *Material Shaders Rollout (mental ray Material) (page 2–1544)*
- *Advanced Shaders Rollout (mental ray Material) (page 2–1548)*

Material Shaders Rollout (mental ray Material)

Material Editor > Type button > Material/Map Browser > mental ray > Material Shaders rollout

Note: The mental ray material appears in the Browser only if the mental ray renderer is the currently active renderer.

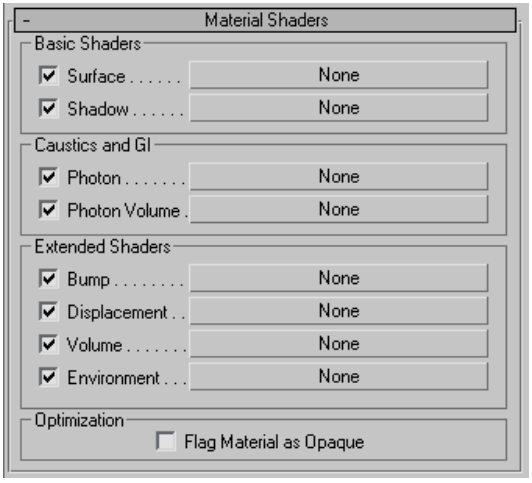
The mental ray material lets you create a material exclusively for use by the *mental ray renderer (page 3–78)*. The Material Shaders rollout provides controls for the main kinds of component shaders you are likely to assign.

Important: You must assign a shader to the material's Surface component. Otherwise, the mental ray material will not be visible when you render.

See also

Advanced Shaders Rollout (mental ray Material) (page 2–1548)

Interface



Each shader component has a toggle at the left of its name. When the toggle is on, the shader is used in rendering. When the toggle is off, the shader is not used, even if it has been assigned. Clicking the button to the right of the component name displays

the *Material/Map Browser (page 2–1412)* so you can assign a particular shader to the component.

Basic Shaders group

Surface—Shades the surface of objects that have this material.

In addition to any of the usual 3ds Max materials, the surface component can be assigned the following mental ray materials or shaders:

Shader	Library
Ambient/Reflective Occlusion	base[1] (see note, below)
Bump (page 2–1716)	3ds Max
DGS Material (page 2–1580)	3ds Max
Dielectric	base
Dielectric Material (page 2–1719)	3ds Max
Edge	lume
Facade	lume
Glass	lume
Glow	lume
Landscape	lume
Material to Shader (page 2–1723)	3ds Max
Metal	lume
Ocean	lume
Opacity	base
Reflect	base
Refract	base
Shader List (page 2–1723)	3ds Max
SSS Physical Material	subsurface scattering
Stain	lume
Texture Wave	base
Translucency	lume
Transmat	physics
Transparency	base
Two Sided	base
UV Generator (page 2–1724)	3ds Max
Water Surface	lume
Wet-Dry Mixer	lume
XYZ Generator (page 2–1729)	3ds Max

Note: As of the current version of 3ds Max, the Ambient/Reflective Occlusion shader has been updated to support certain capabilities for texture baking (see *this note*). If you load a file containing a material that uses the older version of the shader, that same version is still used in the scene, and the shader is renamed "Ambient/Reflective Occlusion (base) (old)". The old version of the shader continues to be used in the scene until you reapply it in the Material Editor.

Note: Unlike a standard 3ds Max material, if you assign the Surface component a bitmap with tiling turned off, the original surface color does not "show through." In renderings, you see only the untiled map, and none of the rest of the object.

Shadow—Assigns a shadow shader.

The shadow component can be assigned the following shaders:

Shader	Library
Edge Shadow	lume
Facade	lume
Glass	lume
Glow	lume
Material to Shader (page 2–1723)	3ds Max
Metal	lume
Shader List (page 2–1723)	3ds Max
Shadow Transparency	base
Translucency	lume
Transmat	physics
Water Surface Shadow	lume

Caustics and GI group

Photon—Assigns a photon shader. Photon shaders modify the appearance of caustics and global illumination. They modify light energy (luminous flux) rather than color (radiance).

The photon component can be assigned the following shaders:

Shader	Library
Shader	**Library**
DGS Material (page 2–1580)	3ds Max
Dielectric Material Photon (page 2–1719)	3ds Max
Edge	lume
Glow	lume
Material to Shader (page 2–1723)	3ds Max
Metal	lume
Photon Basic	base
SSS Physical Material	subsurface scattering
Translucency	lume
Transmat	physics

Photon Volume—Assigns a photon volume shader. Like a photon shader, a photon volume shader modifies caustics and global illumination, but it affects photons that pass through the inside of the object, rather than photons that collide with its surface.

The photon volume component can be assigned the following shaders:

Shader	Library
Material to Shader (page 2–1723)	3ds Max
Parti Volume Photon	physics
Shader List (page 2–1723)	3ds Max

Extended Shaders group

Bump—Assigns a bump shader. Bump shading for mental ray materials is similar to *bump mapping (page 2–1506)* for standard materials.

The bump component can be assigned the following shaders:

Shader	Library
Bump (page 2–1716)	3ds Max
Ocean	lume
Shader List (page 2–1723)	3ds Max

Displacement—Assigns a *displacement shader (page 3–96)*.

The displacement component can be assigned the following shaders:

Shader	Library
3D Displacement (page 2–1714)	3ds Max
Material to Shader (page 2–1723)	3ds Max
Ocean	lume

Volume—Assigns a *volume shader (page 3–95)*.

The volume component can be assigned the following shaders:

Shader	Library
Beam	lume
Material to Shader (page 2–1723)	3ds Max
Mist	lume
Parti Volume Photon	physics
Shader List (page 2–1723)	3ds Max
Submerge	lume

Environment—Assigns an environment shader. Like an environment you assign using the Render Scene dialog, the environment shader changes the scene background.

The environment component can be assigned the following shaders:

Shader	Library
Environment (page 2–1721)	3ds Max

Shader	Library
Material to Shader (page 2–1723)	3ds Max
Shader List (page 2–1723)	3ds Max

Optimization group

Flag Material as Opaque—When on, indicates that the material is fully opaque. This tells the mental ray renderer that it doesn't need to process transparency for this material, or to use the shadow shader (if one has been assigned). This can improve rendering time. Default=off.

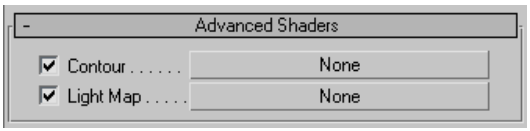
Advanced Shaders Rollout (mental ray Material)

Material Editor > Type button > Material/Map Browser > mental ray > Advanced Shaders rollout

Note: The mental ray material appears in the Browser only if the mental ray renderer is the currently active renderer.

The mental ray material lets you create a material exclusively for use by the *mental ray renderer (page 3–78)*. The Advanced Shaders rollout provides controls for two component shaders that aren't always used.

Interface



Each shader component has a toggle at the left of its name. When the toggle is on, the shader is used in rendering. When the toggle is off, the shader is not used, even if it has been assigned. Clicking the button to the right of the component name displays the *Material/Map Browser (page 2–1412)* so you can assign a particular shader to the component.

Contour—Assigns a *contour shader (page 3–96)* to the material.

The contour component can be assigned the following shaders:

Shader	Library
Combi	contour
Curvature	contour
Depth Fade	contour
Factor Color	contour
Layer Thinner	contour
Simple	contour
Width From Color	contour
Width From Light	contour
Width From Light Dir	contour

Note: Contours don't render unless you have also enabled them on the Render Scene dialog > Renderer panel > *Camera Effects rollout (page 3–101)*.

Light Map—Assigns a light map shader to the material.

Warning: No light map shaders are provided with 3ds Max. This option is for users who have access to light map shaders via other shader libraries or custom shader code.

Arch & Design Material (mental ray)

Material Editor > Type button > Material/Map Browser > Arch & Design Material

Material Editor > any material > Click a Map button. > Material/Map Browser > Car Paint Shader

Note: The Car Paint material and shader appear in the Browser only if the mental ray renderer is the currently active renderer.

The mental ray Arch (architectural) & Design material improves the image quality of architectural renderings and improves workflow and performance in general and for glossy surfaces such as floors in particular. Special features include advanced options for reflectivity and transparency, ambient occlusion settings, and the ability to round off sharp corners and edges as a rendering effect.

See also

Arch & Design Material (mental ray) – Overview (page 2–1562)

Arch & Design Material (mental ray) – Tips & Tricks (page 2–1569)

Interface

Main material parameters rollout (page 2–1549)

BRDF rollout (page 2–1554)

Special Effects rollout (page 2–1555)

Advanced Rendering Options rollout (page 2–1557)

Fast Glossy Interpolation rollout (page 2–1559)

Special Purpose Maps rollout (page 2–1561)

General Maps rollout (page 2–1562)

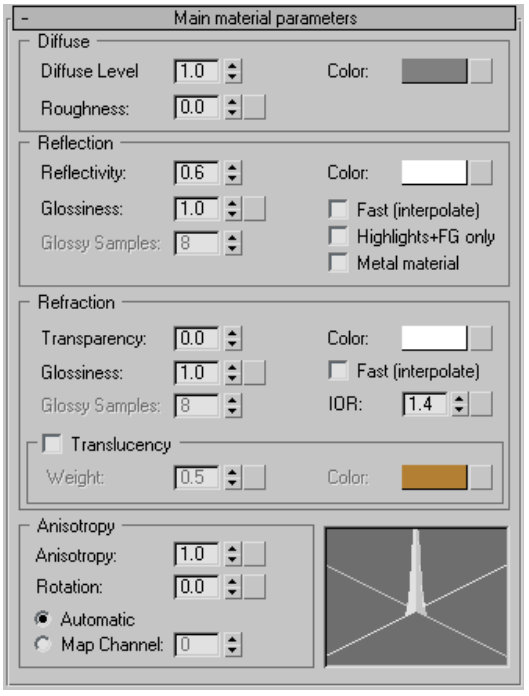
Templates rollout



Provides access to Arch & Design material presets for quickly creating different types of materials such as wood, glass, and metal. You can also use these as starting points for generating customized materials. Choose a template from the drop-down list; a description of the material then appears in the left-hand pane.

Note: The Arch & Design material works only with the mental ray renderer, so in order to see it represented accurately in the sample slots, you must first set mental ray to render in the Material Editor. For details, see *Assign Renderer Rollout (page 3–35)*.

Main material parameters rollout



Diffuse group

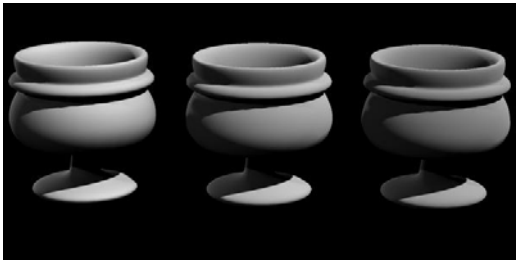
Diffuse Level—*Diffuse Level (page 2–1489)* controls the brightness of the diffuse color component. Range=0.0 to 1.0. Default=1.0.

Note: Because the material is energy conserving, the actual diffuse level used depends on the reflectivity and transparency, as discussed in the introductory section, above.

Color—Controls the *diffuse color (page 3–929)*. The diffuse color is the color in direct light. Default=50% gray.

Roughness—*Roughness (page 2–1490)* controls how quickly the diffuse component blends into the ambient component. Range=0.0 to 1.0. Default=0.0.

The diffuse component uses the Oren-Nayar shading model. When the Roughness value is 0.0 this is identical to classical Lambertian shading, but with higher values the surface gets a more "powdery" look, as shown in the following illustration.



Left: Roughness=0.0; Center: 0.5; Right: 1.0

Reflection group

Reflectivity—The overall level of reflectivity. Range=0.0 to 1.0. Default=0.6.

The Reflectivity and Color values combine to define the level of reflections as well as the intensity of the traditional highlight, also known as the *specular highlight*.

This value is the maximum value; the actual value also depends on the angle of the surface and comes from the BRDF curve. This curve (see *BRDF rollout (page 2–1554)*) lets you define 0–degree reflectivity for surfaces facing the view and 90–degree reflectivity for surfaces perpendicular to the view.



Left: No reflectivity, with a purely diffuse material

Center: Angle-dependent reflectivity, with 0–degree reflectivity of 0.1 and a 90–degree reflectivity of 1.0

Right: Constant reflectivity, with both 0–degree reflectivity and 90–degree reflectivity of 0.9

Note: In the preceding illustration, the high reflectivity automatically "subtracts" from the white diffuse color. If this didn't happen, the material would become unrealistically overbright, and would break the laws of physics.

Color—The overall color of reflected light. Default=white.

Glossiness—Defines the surface "glossiness," ranging from 1.0 (a perfect mirror) to 0.0 (a diffusely reflective surface). Default=1.0.



Left: Glossiness=1.0; Center, 0.5; Right, 0.25

Glossy Samples—Defines the maximum number of samples (rays) that mental ray shoots to create glossy reflections. Higher values cause slow rendering but create a smoother result. Lower values render faster but create a grainier result. Generally 32 is enough for most cases.

Available only when Glossiness does not equal 1.0. Because a Glossiness value of 1.0 creates a "perfect mirror," it is meaningless to shoot multiple rays for this case, hence only one reflection ray is shot.

Note: If you set Glossy Samples to 0, the reflections take the form of a "perfect mirror" and only one ray is shot, regardless of the actual value of Glossiness . You can use this to boost performance for surfaces with very weak reflections. The highlight still respects the Glossiness value.

Glossy reflections need to trace multiple rays to yield a smooth result, which can affect performance. For this reason, the material includes the following two special features designed to enhance performance:

Fast (interpolate)—When on, a smoothing algorithm allows rays to be reused and smoothed. The result is faster and smoother glossy reflections at the expense of accuracy. Interpolation is explained in greater detail in the section on the *Fast Glossy Interpolation rollout (page 2–1559)*.

Note: This method works best on flat surfaces.

Highlights+FG only—When on, mental ray traces no actual reflection rays. Instead, only the highlights are shown, as well as soft reflections emulated with the help of using Final Gathering.

The Highlights+FG Only mode takes no additional rendering time compared to a non-glossy (diffuse) surface, yet can yield surprisingly convincing results. While it might not be completely convincing for "hero" objects in a scene, it can work very well for less-essential scene elements. It tends to work best on materials with weak reflections or extremely glossy (blurred) reflections, as shown in the following illustration:



The two cups on the left use real reflections, while those on the right use Highlights+FG Only.

Metal material—Metallic objects actually influence the color of their reflection, whereas other materials do not. For example, a gold bar will have gold colored reflections, but a red glass orb does not have red reflections. This is supported through the Metal Material option:

- When off, the Reflection Color parameter defines the color, and the Reflectivity parameter together with the BRDF settings defines the intensity and colors of reflections.
- When on, the Diffuse Color parameter defines the color of reflections, and Reflectivity parameter sets the "weight" between diffuse reflections and glossy (metallic) reflections.



Left: Non-metallic reflections (Metal Material is off). Reflections clearly contain the color of the objects they reflect and are not influenced by the color of the materials.

Center: Metallic reflections (Metal Material is on). Now the color of reflections are influenced by the color of the object.

Right: A variant of this with Reflectivity=0.5, creating a 50:50 mix between colored reflections and diffuse reflections

Refraction group

Transparency—Defines the level of refraction. Range=0.0 to 1.0. Default=0.0.

Due to the material's energy-conserving nature, the value set in the Transparency parameter is the maximum value; the actual value depends on the reflectivity as well as the BRDF curve.

Color—Defines the color of refraction. While this color can be used to create "colored glass," a slightly more accurate method to do this is described in the *Colored Glass section (page 2–1570)* of the Tips & Tricks topic.

Glossiness—Defines the sharpness of the refraction/transparency, ranging from 1.0 (completely clear transparency) to 0.0 (extremely diffuse or blurry transparency). Default=1.0.



Left: Refraction Glossiness=1.0; Center: 0.5; Right: 0.25

Glossy refraction needs to trace multiple rays to yield a smooth result, which can affect performance. For this reason, the material includes the following special feature designed to enhance performance:

Fast (interpolate)—When on, a smoothing algorithm allows rays to be reused and smoothed. The result is faster and smoother glossy refraction at the expense of accuracy. Interpolation is explained in greater detail in the section on the *Fast Glossy Interpolation rollout (page 2–1559)*.

Note: This method works best on flat surfaces.

Glossy Samples—Defines the maximum number of samples (rays) that mental ray shoots to create glossy refraction. Higher values cause slow rendering but create a smoother result. Lower values render faster but create a grainier result, like frosted glass. Generally 32 is enough for most cases.

Available only when Glossiness does not equal 1.0. Because a Glossiness value of 1.0 creates a perfectly clear (non-blurry) transparency, it is meaningless to shoot multiple rays for this case, hence only one refraction ray is shot.

Note: If you set Glossy Samples to 0, the refraction takes the form of a "perfect lens" and only one ray is shot, regardless of the actual value of Glossiness. You can use this to boost performance for draft renderings.

IOR—The Index of Refraction, which is a measurement of how much a ray of light bends when entering a material.

The direction in which light bends depends on whether it is entering or exiting the object. The Arch & Design material use the direction of the surface normal as the primary cue for figuring out whether it is entering or exiting. It is therefore important to model transparent, refractive objects with the surface normals pointing in the proper direction.

The IOR can also be used to define the BRDF curve, which is what happens in the class of transparent materials known as "dielectric" materials, and is illustrated here:



Left: IOR=1.0; Center: 1.2; Right: 1.5

The leftmost cup looks completely unrealistic and is almost invisible. Because an IOR of 1.0, which equals that of air, is impossible in solid matter, we get no change in reflectivity across the material and hence perceive no edges or changes of any kind. On the other hand, the center and rightmost cups have realistic changes in reflectivity guided by the IOR.

Instead of basing the reflectivity on the IOR, you can instead use the BRDF mode to set it manually:



Different types of transparency

As in the previous illustration, the leftmost cup acquires its curve from the index of refraction. The center cup has a manually defined curve, which has been set to a 90 degree reflectivity of 1.0 and a 0 degree reflectivity of 0.2; this looks a bit more like metallized glass. The rightmost cup uses the same BRDF curve, but instead is set to *thin-walled transparency (page 2–1558)*. Clearly, this method is better for making non-refractive objects than simply setting IOR to 1.0, as we tried above.

Translucency group

Translucency is handled as a special case of transparency; in order to use translucency there must first exist some level of transparency. The implementation of translucency in the Arch & Design material is a simplification concerned solely with the transport of light from the back of an object to its front faces and is not true SSS (subsurface scattering) effect. You can create an SSS-like effect by using glossy transparency coupled with translucency, but this is neither as fast nor as powerful as the dedicated *SSS shaders (page 2–1583)*.

Translucency—When on, the remaining Translucency become available and take effect when rendering.

Weight—Determines how much of the existing transparency is used as translucency. For example, if Weight=0.0, all of the transparency is used as transparency. If Weight=0.3, 30 percent of the transparency is used as translucency.



All: Transparency=0.75. Left: Weight=0.0; Center: Weight=0.5; Right: Weight=1.0

Translucency is intended for use primarily in *thin-walled mode*, as in the example above) to model things like curtains, rice paper, and similar effects. In thin-walled mode it simply allows the shading of the reverse side of the object to bleed through. The shader also operates in *solid mode*, but, as explained above, the SSS shaders are better suited for such purposes.



Solid translucency with: Left: Weight=0.0; Center: Weight=0.5; Right: Weight=1.0

Color—The translucency color.

Anisotropy group

Anisotropy—Controls the anisotropy, or shape, of the highlight. At 1.0, the highlight is round; that is, no anisotropy. At 0.01, the highlight is elongated. One axis of the highlight graph changes to show changes in this parameter. Default=1.0.



Left: Anisotropy=1.0; Center: Anisotropy=4.0; Right: Anisotropy=8.0

Rotation—Changes the orientation of the highlight. The sample slot shows changes in orientation. This value can range from 0.0 to 1.0, with 1.0=360 degrees. So, for example, 0.25=90 degrees and 0.5=180 degrees. Default=0.0.



Left: Anisotropy Rotation=0.0; Center: Anisotropy Rotation=0.25; Right: Anisotropy Rotation=[texture map]

Tip: When using texture-mapped Anisotropy Rotation, make sure the texture is not antialiased (filtered). You can achieve this by setting the map's Blur parameter to 0.0. Otherwise the antialiased pixels cause local vortices in the anisotropy that appear as seam artifacts.

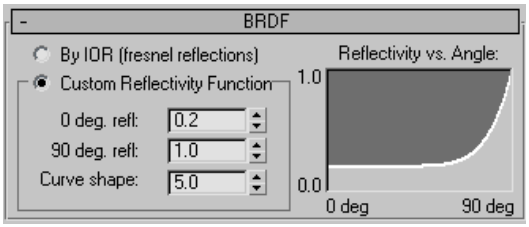
Automatic/Map Channel—Lets you optionally apply anisotropy to a specific map channel.

If the Map Channel setting is Automatic, the base rotation follows the object's local coordinate system. If it is any other value (in other words, a specific map channel), the space that defines the stretch directions of the highlights is derived from that channel's texture space.

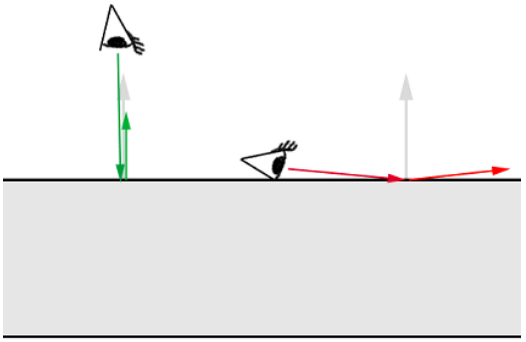
Warning: Deriving the anisotropy from the texture space creates only one space per triangle and can cause visible seams between triangles.

Also see *Brushed Metal (page 2–1575)*.

BRDF rollout



BRDF stands for *bidirectional reflectance distribution function*. As explained in the *introduction (page 2–1566)*, this property lets the material's reflectivity be ultimately guided by the angle from which the object surface is viewed.



0 degree (green) and 90 degree (red) view angles

[BRDF method]—Lets you choose how the BRDF curve is defined:

- **By IOR (fresnel reflections)—**How the reflectivity depends on the angle is guided solely by the

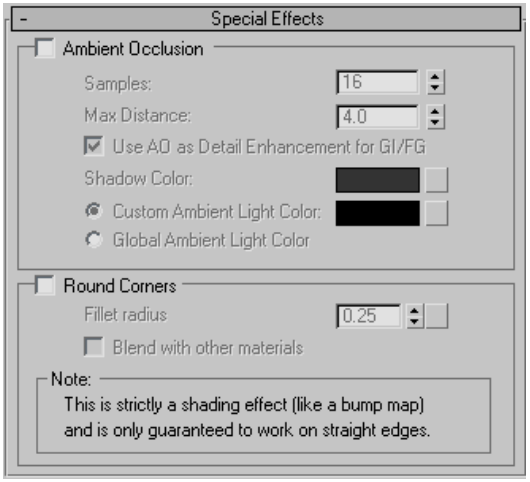
material's index of refraction. This is known as *Fresnel reflections* and follows the behavior of most dielectric materials such as water and glass.

- **Custom Reflectivity Function—**When chosen, the following settings determine reflectivity based on angle of view.
  - **0 deg. refl.—**Defines the reflectivity for surfaces directly facing the viewer (or incident ray).
  - **90 deg. refl.—**Defines the reflectivity of surfaces perpendicular to the viewer.
  - **Curve shape—**Defines the falloff of the BRDF curve.

This mode is used for most hybrid materials or for metals. Most materials exhibit strong reflections at grazing angles; hence the 90 degree reflectivity parameter can generally be kept at 1.0, using the Reflectivity parameter to guide the overall reflectivity instead. Metals tend to be fairly uniformly reflective and the 0 degree reflectivity value is high (0.8–1.0), but many other layered materials, such as linoleum and lacquered wood have lower 0 degree reflectivity values, in the range 0.1–0.3). For further information, see *Quick Guide to Some Common Materials (page 2–1569)*.

Reflectivity vs. Angle graph—Depicts the combined Custom Reflectivity Function settings.

Special Effects rollout



Provides settings for *ambient occlusion (page 2–1567)* (AO) and round corners/edges.

Ambient Occlusion group

Ambient Occlusion—When on, enables AO and makes the remaining group controls available.

Samples—The number of samples (rays) shot for creating AO. Higher values yield smoother results but render more slowly, while lower values render faster but look grainier. Values in the range 16–64 cover most situations.

Max Distance—Defines the radius within which mental ray looks for occluding objects. Smaller values restrict the AO effect to small crevices only but are much faster to render. Larger values cover larger areas but render more slowly.

The following illustrations show the raw AO contribution with two different distances:



Left: Higher Max Distance value; Left: Lower Max Distance value

Use AO as Detail Enhancement for GI/FG—Applies short-distance AO, multiplying it with the existing indirect illumination (Final Gathering or GI/photons) to bring out small details.

The following illustration depicts a model helicopter that is lit almost exclusively by indirect light:



Left: Without AO; Right: With AO

Note how the helicopter does not feel "grounded" in the left image and the shadows under the landing skids are far too vague. The right image uses AO to "punch out" the details and the contact shadows.

Shadow Color—Sets the darkness of the AO shadows. It is used as the multiplier value for completely occluded surfaces. In practice, a black color makes the AO effect very dark; a middle-gray color makes the effect less noticeable (brighter), and so on.

Custom/Global Ambient Light Color—You can specify a color for the ambient light used in AO, or use the global color specified on the *Environment panel (page 3–272)* > Common Parameters rollout.

This parameter is used for doing more traditional AO; that is, supplying an imagined "ever-present ambient light" that is then attenuated by the AO effect to create shadows.

While traditional AO is generally used when rendering without other indirect light, you can also combine it with existing indirect light. Bear in mind that this "ever-present ambient light" is inherently non-physical, but can possibly help lighten some troublesome dark corners.

Round Corners group

This effect rounds off corners and straight edges as a rendering effect only; it has no effect on geometry.

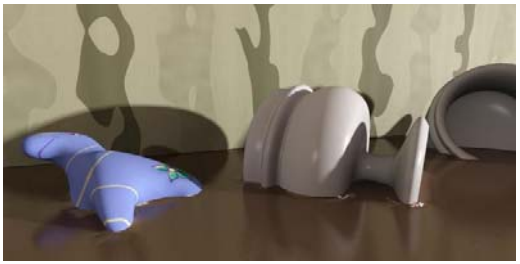
The rounding effect happens to convex corners and surfaces that actually intersect. Concave corners that merely touch will not display the effect. To get the effect to work in concave corners the objects must be pushed into each other a little. The effect is intended for straight edges and is not guaranteed to work properly for highly curved, complex intersections.

Round Corners—When on, rounds off corners and straight edges at render time.

Fillet Radius—Specifies the radius of the filleted corners and/or edges.

Blend with other materials—By default, the rounding effect happens only between surfaces of the same material, but if you turn this on the filleting is performed against any material.

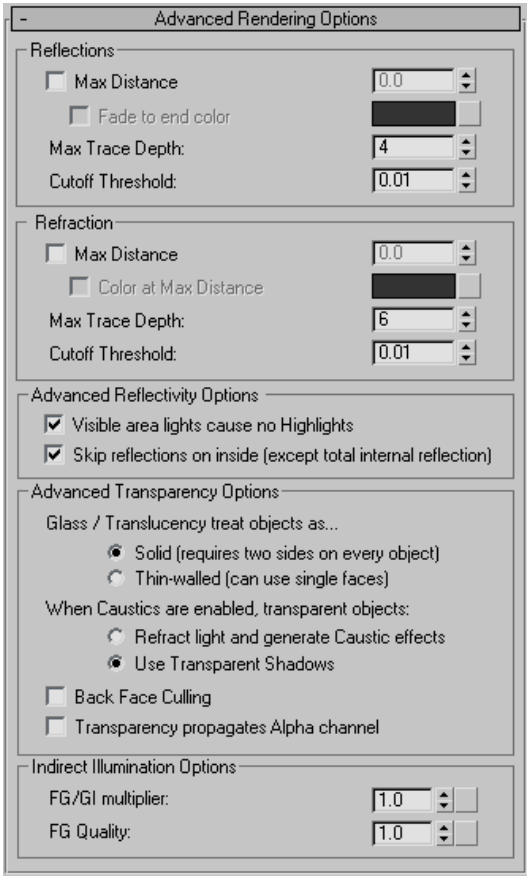
In the following image, the molten chocolate is rounded off against the submerged objects even though they use different materials. In actuality, the molten chocolate is a completely flat plane.



Objects in molten chocolate

Advanced Rendering Options rollout

These parameters define performance-boosting options.



Reflections group

Max Distance—Allows limiting reflections to a certain distance, which both speeds up rendering and avoids pulling distant objects into extremely glossy reflections.

Fade to end color—When on, reflections fade to this color. When off, reflections fade to the environment color. The former tends to be more useful for indoor scenes; the latter, for outdoor scenes.

Available only when Max Distance is on.



Left: Full reflections (Max Distance=off); Center: 100mm; Right: 25mm

Max Trace Depth—When this trace depth is reached, the material behaves as if the Highlights+FG Only switch is on; that is, it shows only highlights and "emulated" reflections created with the help of Final Gathering.

Cutoff Threshold—The level at which reflections are rejected; that is, not traced. It's a relative value: For example, the default setting of 0.01 means that rays that contribute less than 1 percent to the final pixel are ignored. A setting of 0.25 means that mental ray discards rays that contribute less than a quarter of the value of the final pixel.

Refraction group

The optimization settings for refraction (transparency) are nearly identical to those for reflections. The exception is that of Color At Max Distance, which behaves differently.

Max Distance—Allows limiting refraction to a certain distance.

Color at Max Distance—When on, the material simulates physically correct absorption. At the distance specified by Max Distance, the refracted image has the color given by Color At Max Distance, but the rays are not limited in reach. At twice the distance, the influence of Color At Max Distance is double, at half the distance half, etc.

When off, transparency rays simply fade to black. This is like smoked glass and other highly absorbent materials. Transparency just stops at the specified distance. This has the same performance advantage as using the Max Distance for reflections: Tracing shorter rays is much faster.

Available only when Max Distance is on.



Left: No limit (Max Distance=off); Center: Fade to black; Right: Fade to blue

Max Trace Depth—When this trace depth is reached, the material refracts black.

Cutoff Threshold—The level at which refraction is rejected; that is, not traced. It's a relative value: For example, the default setting of 0.01 means that rays that contribute less than 1 percent to the final pixel are ignored. A setting of 0.25 means that mental ray discards rays that contribute less than a quarter of the value of the final pixel.

Advanced Reflectivity group

Visible area lights cause no Highlights—When on, mental ray area lights (*Omni (page 2–1298)* and *Spotlight (page 2–1299)*) with the Show Icon In Renderer property on create no specular highlights. Default=on.

The Show Icon In Renderer check box is found on the light's Area Light Parameters rollout. When on, the light is visible and reflects in any glossy, reflective objects. If both the reflection of the visible area light and the highlight is rendered, the light is added twice, causing an unrealistic brightening effect. When on, this switch causes visible area lights to lose their highlights and instead only appear as reflections. Note that this does not apply to the *Highlights+FG Only* mode, which doesn't actually reflect anything.

Skip reflections on inside (except total internal reflection)—Most reflections *inside* transparent objects are very faint, except in the special case that occurs at certain angles known as *total internal reflection* (TIR). When on, this option saves rendering time by ignoring the weak reflections completely but retaining the TIRs. Default=on.

Advanced Transparency Options group

The options give you control over some of the deepest details of the Arch & Design material.

Glass/Translucency treat objects as...—

- **Solid—**The object behaves as if it is made of a solid, transparent substance.
- **Thin-walled—**The object behaves as if made of wafer-thin sheets of a transparent material.

For more information, see *Solid vs. Thin-Walled (page 2–1567)*.



Left: Solid; Right: Thin-walled

When Caustics are enabled, transparent objects:—When not rendering caustics, the Arch & Design material uses a shadow shader to create transparent shadows. For objects such as window panes this is perfectly adequate, and actually creates a better result than using caustics, because the direct light is allowed to pass more or less undisturbed through the glass into a space such as a room.

Traditionally, enabling caustics in mental ray causes all materials to stop casting transparent shadows and instead start to generate refractive caustics. In most architectural scenes this is undesirable; you might want a glass decoration on a table to generate caustic effect, but still want the windows to the room to let in normal direct light. This switch makes this possible at the material level.

- **Refract light and generate Caustic effects—**The material refracts light and generates caustics.
- **Use Transparent Shadows—**No caustics are produced; the material and object simply transmit the light, with greater shadowing in thicker areas.

In the following illustration, the left side shows the result with Use Transparent Shadows chosen, and the right side shows the result with Refract Light And Generate Caustic Effects chosen. You can freely mix the two modes in the same rendering. Photons are automatically treated accordingly by the built-in photon shader, shooting straight through as direct light in the former case, and being refracted as caustics in the latter.



Left: Using transparent shadows; Right: Using refractive caustics.

Back Face Culling—When on, enables a special mode that makes surfaces invisible to the camera when seen from the reverse side. You can use this to create "magic walls" in a room. If all walls are planes with the normals facing inwards, the Back Face Culling switch allows the room to be rendered from "outside." The camera can see into the room, but the walls will still exist and cast shadows, bounce photons, etc. while being invisible when the camera goes outside.



Left: Back Face Culling=off; Right: Back Face Culling=on

Transparency propagates Alpha channel—Defines how transparent objects treats any alpha channel information in the background. When on, refraction and other transparency effects propagate the alpha of the background "through" the transparent object. When off, transparent objects have an opaque alpha.

Indirect Illumination Options

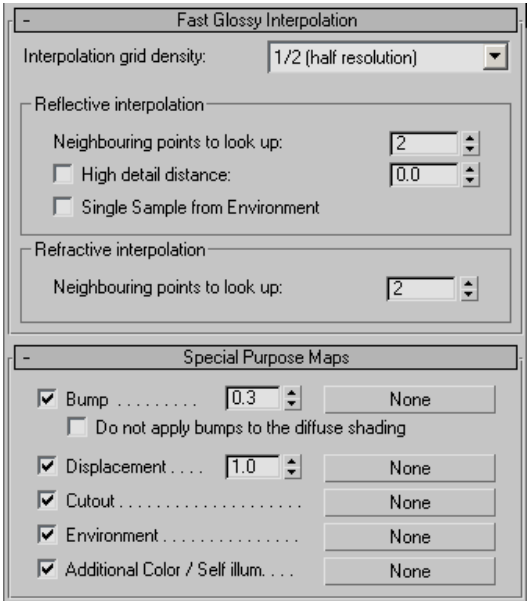
FG/GI multiplier—Allows tweaking of how strongly the material responds to indirect light.

FG Quality—A local multiplier for the number of final gather rays shot by the material.

Fast Glossy Interpolation rollout

Glossy reflections and refraction can be *interpolated*, which causes them to render faster and look smoother. Interpolation works by precalculating glossy reflection in a grid across the image. The number of samples (rays) taken at each point is govern by the *Reflection > Glossy Samples* or *Refraction > Glossy Samples* parameters, as in the non-interpolated case.

Note that interpolation can cause *artifacts*. Because it is done on a low-resolution grid, it can lose details. And because it blends neighbors of this low-resolution grid, it can cause oversmoothing. For this reason it is useful primarily with flat surfaces. Interpolation does not work well with wavy, highly detailed surfaces or surfaces that use bump maps.



Interpolation grid density—The resolution of the grid used for interpolating glossy reflections and refraction. Choose a setting from the drop-down list.

Within the grid, data is stored and shared across the points. Using a lower grid resolution is faster, but causes greater loss of detail information.

Reflective interpolation group

Neighboring points to look up—Defines how many stored grid points (in an N by N group around the currently rendered point) is looked up to smooth out reflective glossiness. The default is 2, and higher values will "smear" the glossiness more, but are hence prone to more oversmoothing artifacts.

In the following illustration, the reflection of the left cup in the floor does not use interpolation, and some grain is evident (here intentionally exaggerated). The floor tiles under the other two cups use a half-resolution interpolation with point lookup set to 2 (center) and 4 (right), respectively.



Left: No interpolation; Center: Looking up two points; Right: Looking up four points.

The preceding image also illustrates one of the consequences of using interpolation: The foot of the left cup, which is near the floor, is reflected quite sharply, and only the parts of the cup far from the floor are blurry. However, the interpolated reflections of the right cups have a base level of blurriness, due to the smoothing of interpolation, which makes even the closest parts somewhat blurry. In most scenes with weak glossy reflections this discrepancy will never be noticed, but in other cases this can make things like legs of tables and chairs feel "unconnected" with a glossy floor, if the reflectivity is high. To resolve this, you can use the High Detail Distance setting (see following).

High detail distance—Allows tracing of a second set of detail rays to create a "clearer" version of objects within the specified radius.

In the following illustration, all three floor tiles use interpolation but the two on the right use different distances for the detail distance.



Left: No detail distance; Center: 25mm detail distance; Right: 150mm detail distance

This also allows an interesting trick: Set Reflection > Glossy Samples to 0, which renders reflections as if they were mirror-perfect, but also use interpolation to introduce blur into this reflection, and perhaps use High Detail Distance to make nearby parts less blurry. This is a fast way to obtain a glossy reflection.

The floor tiles in the following illustration are rendered with mirror reflections, and the blurriness comes solely from the interpolation. This renders as fast as or faster than pure mirror reflections, yet gives a satisfying illusion of true glossy reflections, especially when utilizing the High Detail Distance option, as on the right.



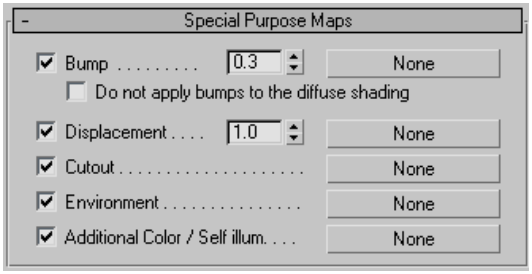
Left: No detail distance; Right: With detail distance

Single Sample from Environment—Creating realistically blurry glossy reflections normally requires taking multiple samples from the environment, which can result in grainy, slow-rendering environment reflections. With this check box on, mental ray instead takes only *one* sample, thus preventing the grain. This also prevents blurring the environment, so it is best used together with a local, "pre-blurred" environment map. You can do the pre-blurring in an image-processing program or with the Material Editor > *Coordinates rollout (page 2–1625)* > Blur and Blur Offset settings.

Refractive interpolation group

Neighboring points to look up—Defines how many stored grid points (in an N by N group around the currently rendered point) are looked up to smooth out refractive glossiness. The default is 2; higher values tend to "smear" the glossiness more, but are hence prone to more oversmoothing artifacts.

Special Purpose Maps rollout



Lets you apply bump, displacement, and other maps. Each left-justified setting has a check box for enabling and disabling the map, and a button for defining the map.

Bump—Lets you apply a bump map and multiplier.

Do not apply bumps to the diffuse shading—When off, the bumps apply to all shading components: diffuse, highlights, reflections, refractions, etc. When on, bumps are applied to all components except the diffuse. This means bumps are seen in reflections, highlights, etc. but the diffuse shading shows no bumps. It is as if the material's diffuse surface is smooth, but is covered by a bumpy lacquer coating.



Left: Do Not Apply Bumps ...=off; Right: Do Not Apply Bumps ...=on

Displacement—Lets you apply a displacement map and multiplier.

Cutout—Lets you apply an opacity map to completely remove parts of objects. A classic example is to map an image of a tree to a flat plane and use opacity to cut away the parts of the tree that are not there.



Left: Mapped transparency; Right: Using Cutout

Environment—Lets you apply an environment map and shader.

Additional Color/Self illum.—Lets you apply any shader. The output of this shader is added on top of the shading done by the Arch & Design material and can be used for self-illumination-type effects, as well as adding any additional shading you want.

General Maps rollout

This rollout enables application of maps or shaders to any of the Arch & Design material parameters. Of course, you can apply a shader to a parameter at its standard location in the user interface by clicking its Map button, so the principal value of this rollout is that it also lets you toggle a parameter's shader, using the check box, without removing the map.

Arch & Design Material (mental ray) – Overview

This topic serves as an introduction to the Arch & Design material for mental ray. For a detailed reference to the material interface, see *Arch & Design Material (mental ray) (page 2–1549)*. For a variety of suggestions for using the material to create specific effects, see *Arch & Design Material (mental ray) – Tips & Tricks (page 2–1569)*.



A range of material effects available with the Arch & Design material

What is the Arch & Design material?

The mental ray Arch & Design material is a monolithic material shader designed to support most materials used in architectural and product-design renderings. It supports most hard-surface materials such as metal, wood and glass. It is especially tuned for fast glossy reflections and refractions (replacing the *DGS material (page 2–1580)*) and high-quality glass (replacing the *dielectric material (page 2–1719)*).

The major features are:

- **Easy to use, yet flexible** - controls are arranged logically in a most-used-first fashion.
- **Templates** - allow fast access to settings combinations for common materials.
- **Physically accurate** - the material is energy conserving, making it impossible to create shaders that break the laws of physics.
- **Glossy performance** - advanced performance boosts including interpolation, emulated glossiness, and importance sampling.
- **Tweakable BRDF** (bidirectional reflectance distribution function) - the user can define how reflectivity depends on angle.
- **Transparency** - "Solid" or "thin" materials: transparent objects such as glass can be treated as either solid (refracting, built out of multiple faces) or thin (nonrefracting, can use single faces).
- **Round corners** - simulate fillets to allow sharp edges to still catch the light in a realistic fashion.
- **Indirect Illumination control** - set the final gather accuracy or indirect illumination level on a per-material basis.
- **Oren-Nayar diffuse** - allows "powdery" surfaces such as clay.
- **Built-in Ambient Occlusion** - for contact shadows and enhancing small details.
- **All-in-one shader** - photon and shadow shader built in.
- **Waxed floors, frosted glass and brushed metals** - all fast and easy to set up.

Physics and the display

The Arch & Design material attempts to be physically accurate, hence its output has a high dynamic range. How visually pleasing the material looks depends on how colors inside the renderer are mapped to colors displayed on the screen.

When working with the Arch & Design material it is highly recommended that you operate through a tone mapper/exposure control such as the *Logarithmic Exposure Control (page 3–297)*, or at the very least use gamma correction.

A note on gamma

Describing all the details of gamma correction is beyond the scope of this topic; this is just a brief overview.

The color space of a normal, off-the-shelf computer screen is not linear. The color with RGB value 200 200 200 is not twice as bright as a color with RGB value 100 100 100, as one might expect.

This is not a bug because, due to the fact that our eyes see light in a non-linear way, the former color is actually perceived to be about twice as bright as the latter. This makes the color space of a normal computer screen roughly perceptually uniform. This is a good thing, and is actually the main reason 24-bit color (with only 8 bits or 256 discrete levels for each of the red, green and blue components) looks as good as it does to our eyes.

The problem is that physically correct computer graphics operates in a true linear color space where a value represents actual light energy. If one simply maps the range of colors output to the renderer naively to the 0-255 range of each RGB color component it is incorrect.

The solution is to introduce a mapping of some sort. One of these methods is called *gamma correction*.

Most computer screens have a gamma of about 2.2 (known as the sRGB color space), but most software defaults to a gamma of 1.0, which makes everything look too dark (especially midtones), and light does not "add up" correctly.

Using a gamma of 2.2 is the theoretically correct value, making the physically linear light inside

the renderer appear in a correct linear manner on screen.

However, because the response of photographic film isn't linear either, users find that this theoretically correct value looks too bright and washed out. A common compromise is to render to a gamma of 1.8, making things look more photographic; that is, as if the image had been shot on photographic film and then developed.

Tone mapping

Another method for mapping the physical energies inside the renderer to visually pleasing pixel values is known as *tone mapping*. This can be accomplished either by rendering to a floating-point file format and using external software, or with a plug-in that allows the renderer to do it on the fly. In 3ds Max such plug-ins are known as exposure controls and are accessed from the Environment dialog.

Note: When using the Logarithmic Exposure Control, the need for gamma correction is largely replaced. You should use either gamma correction or the Logarithmic Exposure Control but not both.

Use Final Gathering and Global Illumination

The Arch & Design material is designed to be used in a realistic lighting environment; one that incorporates full direct and indirect illumination.

mental ray provides two basic methods for generating indirect light: Final Gathering and Global Illumination. For best results, be sure to use at least one of these methods.

At the very least, enable Final Gathering, or use Final Gathering combined with Global Illumination (photons) for quality results. Performance tips for using Final Gather and Global Illumination can be found *here (page 2–1569)*.

If you use an environment for your reflections, make sure the same environment (or a blurred copy of it) is used to light the scene through Final Gathering. In 3ds Max this means you should include a *Skylight (page 2–1296)* in your scene set to Use Scene Environment, or use *Daylight system (page 1–418)* with Skylight set to mr Sky.

Use physically correct lights

Traditional computer-graphics light sources live in a cartoon universe where the intensity of the light doesn't change with the distance. The real world doesn't agree with that simplification. Light decays when leaving a light source due to the fact that light rays diverge from their source and the intensity of the light changes over distance. This decay of a point light source is $1/d^2$; in other words, light intensity is proportional to the inverse of the square of the distance to the source.

One of the reasons for this traditional oversimplification is the fact that in the early days of computer graphics, tone mapping was not used and problems of colors blowing out to white in the most undesirable ways was rampant. (Raw clipping in sRGB color space is displeasing to the eye, especially if one color channel clips earlier than the others. Tone mapping generally solves this by "soft clipping" in a more suitable color space than sRGB.)

However, as long as only Final Gathering (FG) is used as indirect illumination method, such traditional simplifications still work. Even light sources with no decay still create reasonable renderings. This is because FG is concerned only with the transport of light from one surface to the next, not with the transport of light from the light source to the surface.

It's when working with Global Illumination (GI) (that is, with photons) the troubles arise.

When GI is enabled, light sources shoot photons. For the Arch & Design material (or any other

mental ray material) to be able to work properly, it is imperative that the energy of these photons to match the direct light cast by that same light. And since photons model light in a physical manner, decay is built in.

Hence, when using GI:

- Light sources must emit photons at the correct energy.
- The direct light must decay in a physically correct way to match the decay of the photons.

Therefore it is important to make sure the *light shader* and the *photon emission shader* of the lights work well together.

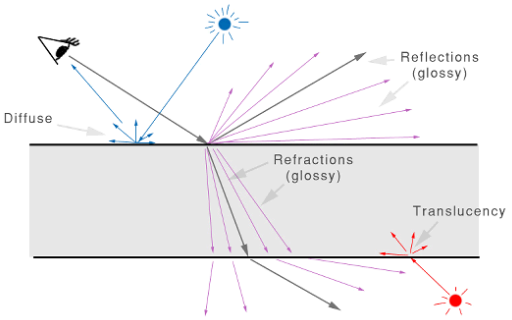
In 3ds Max this is most easily solved by using the *photometric lights (page 2–1301)*. All of these lights are guaranteed to have their photon energy in sync with their direct light. It is built in and automatic and one does not need to worry about it.

Features

The Shading Model

From a usage perspective, the shading model consists of three components:

- **Diffuse** - diffuse channel /including Oren Nayar "roughness").
- **Reflections** - glossy anisotropic reflections (and highlights).
- **Refraction** - glossy anisotropic transparency (and translucency).



The Arch & Design material shading model

Direct and indirect light from the scene cause diffuse reflections as well as translucency effects. Direct light sources also create specular highlights.

Ray tracing is used to create reflective and refractive effects, and advanced importance-driven multi-sampling is used to create glossy reflections and refraction.

The rendering speed of the glossy reflections/refraction can further be enhanced by interpolation as well as "emulated" reflections with the help of Final Gathering.

Conservation of Energy

One of the most important features of the material is that it is automatically energy conserving. This means that it makes sure that diffuse + reflection + refraction <= 1. In other words, no energy is magically created and the incoming light energy is properly distributed to the diffuse, reflection and refraction components in a way that maintains the first law of thermodynamics.

In practice, this means, for example, that when adding reflectivity, the energy must be taken from somewhere, and hence the diffuse level and the transparency will be automatically reduced accordingly. Similarly, adding transparency happens at the cost of the diffuse level.

The rules are as follows:

- Transparency takes energy from diffuse; that is, at 100% transparency, there is no diffuse at all.
- Reflectivity takes energy from both diffuse and transparency; that is, at 100% reflectivity there is neither diffuse nor transparency.
- Translucency is a type of transparency, and the Translucency Weight parameter defines the percentage of transparency vs. translucency.



From left to right: Reflectivity=0.0, 0.4, 0.8, and 1.0



From left to right: Transparency=0.0, 0.4, 0.8, and 1.0

Conservation of energy also means that the level of highlights is linked to the glossiness of a surface. A high Reflection Glossiness value causes a narrow, intense highlight, while a lower value causes a wider, less intense highlight. This is because the energy is now spread out and dissipated over a larger area.

BRDF - how reflectivity depends on angle

In the real world, the reflectivity of a surface is often view-angle dependent. A fancy term for this is *bidirectional reflectance distribution function* (BRDF); that is, a way to define how much a material reflects when seen from various angles.



The reflectivity of the wood floor depends on the view angle.

Many materials exhibit this behavior. The most obvious examples are glass, water, and other dielectric materials with Fresnel effects (where the angular dependency is guided strictly by the index of refraction), but other layered materials such as lacquered wood and plastic display similar characteristics.

The Arch & Design material allows this effect to be defined by the index of refraction, and also allows an explicit setting for the two reflectivity values for:

- 0 degree faces (surfaces directly facing the camera)
- 90 degree faces (surfaces 90 degrees to the camera)

For more information, see *BRDF rollout (page 2–1554)*.

Reflectivity features

The final surface reflectivity is in reality caused by the sum of three components:

- The diffuse effect
- The actual reflections
- Specular highlights that simulate the reflection of light sources



Diffuse, reflections, and highlights

In the real world, highlights are just glossy reflections of the light sources. In computer graphics it's more efficient to treat these separately. However, to maintain physical accuracy the material automatically keeps highlight intensity, glossiness, anisotropy, etc. in sync with the intensity, glossiness and anisotropy of reflections. Thus, there are no separate controls for these as both are driven by the reflectivity settings.

Transparency features

The material supports full glossy anisotropic transparency and includes a translucent component, described in detail *here (page 2–1553)*.



Translucency

Solid vs. Thin-Walled

The transparency/translucency property can treat objects as either *solid* or *thin-walled*.

If all objects were treated as solids at all times, every window pane in an architectural model would have to be modeled as *two* faces: an entry surface that refracts the light slightly in one direction, and immediately following it an exit surface, where light is refracted back into the original direction.

Not only does this entail additional modeling work, it is a waste of rendering power to simulate refraction that has very little net effect on the image. Hence the material allows modeling the entire window pane as a single flat plane, foregoing any actual refraction of light.



Solid vs. thin-walled transparency and translucency

In the preceding illustration the helicopter canopy, the window pane, the translucent curtain, and the right-hand sphere all use thin-walled transparency or translucency, whereas the glass goblet, the plastic horse, and the left-hand sphere all use solid transparency or translucency.

Cutout opacity

Beyond the "physical" transparency, which models an actual property of the material, the material provides a completely separate, non-physical "cutout opacity" channel to allow "billboard" objects such as trees, or to cut out objects such as a chainlink fence with an opacity mask.

Special Effects

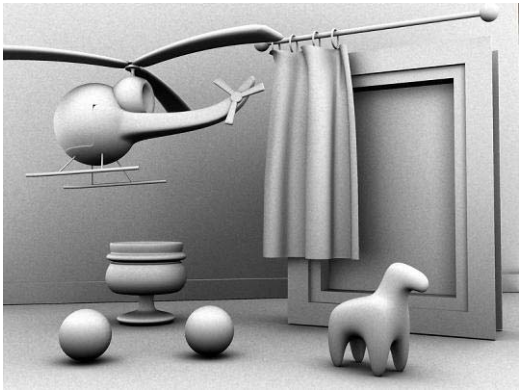
Built-in Ambient Occlusion

Ambient Occlusion (AO) is a method spearheaded by the film industry for emulating the look of true global illumination by using shaders that calculate the extent to which an area is occluded, or prevented from receiving incoming light.

Used alone, an AO shader, such as the separate mental ray Ambient/Reflective Occlusion shader, creates a grayscale output that is dark in areas light cannot reach and bright in areas where it can:

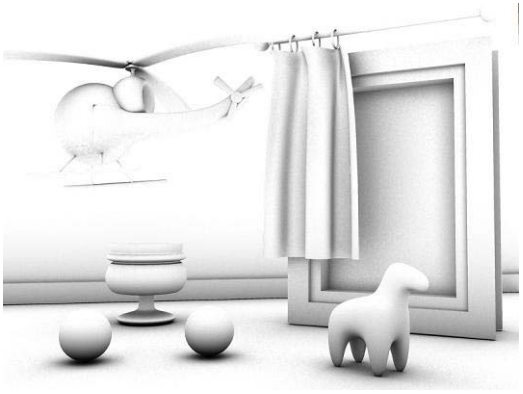
The following image illustrates the main results of AO: dark crevices and areas where light is blocked

by other surfaces, and bright areas that are exposed to the environment.



An example of AO applied to a scene

One important aspect of AO is that the user can tune the distance it looks in for occluding geometry.



AO looked up within a shorter radius

Using a radius creates a localized AO effect; only surfaces within the given radius are considered as occluders (which is also faster to render). The practical result is that the AO provides nice "contact shadow" effects and makes small crevices visible.

The Arch & Design material gives you two ways to utilize its built-in AO:

- Traditional AO for adding an omnipresent ambient light that is then attenuated by the AO to create details.
- Use AO for *detail enhancement* together with existing indirect lighting methods such as Final Gathering or photons.

The latter method is especially interesting when using a highly smoothed indirect illumination solution, such as a high photon radius or an extremely low final gather density, which could otherwise lose small details. By applying the AO with short rays these details can be brought back.

Round corners

Computer-generated imagery tends to look unrealistic, partly because edges of objects are geometrically sharp, whereas most edges in the real world are slightly rounded, chamfered, worn, or filleted in some manner. This rounded edge tends to "catch the light" and create highlights that make edges more visually appealing.

The Arch & Design material can create the illusion of rounded edges at render time. This feature is intended primarily to speed up modeling, so that you need not explicitly fillet or chamfer edges of objects such as a tabletop.



Left: No round corners; Right: Round corners

The function is not a displacement; it is merely a shading effect, like bump mapping, and is best suited for straight edges and simple geometry, not advanced, highly curved geometry.

Performance Features

Finally, the Arch & Design material contains a large set of built-in functions for optimal performance, including but not limited to:

- Advanced importance sampling with ray rejection thresholds
- Adaptive glossy sample count
- Interpolated glossy reflection/refraction with detail enhancements
- Ultra-fast emulated glossy reflections (Highlights+FG Only mode)
- The option to ignore internal reflections for glass objects
- The choice between traditional transparent shadows, suitable for objects such as a window pane, and refractive caustics, suitable for solid glass objects, on a per-material basis.

Arch & Design Material (mental ray) – Tips & Tricks

This topic contains information to help you more effectively use the *Arch & Design material (page 2–1549)* for mental ray.

See also

Arch & Design Material (mental ray) – Overview (page 2–1562)

Arch & Design Material (mental ray) (page 2–1549)

Final Gather Performance

The final gather algorithm in mental ray 3.5 is vastly improved from earlier versions, especially in its adaptiveness. This means you can often use much lower ray counts and much lower densities than in previous versions of mental ray.

In many cases, you can render still images with such extreme settings as 50 rays and a density of 0.1. If this causes "oversmoothing" artifacts, you can use the *built-in ambient occlusion (page 2–1555)* to solve those problems.

When using final gather together with GI (photons), make sure the photon solution is fairly smooth by first rendering with Final Gather disabled first. If the photon solution is noisy, increase the photon search radius until it "calms down," and then enable Final Gather.

Quick Guide to Some Common Materials

Following are some quick rules of thumb for creating various materials. Each assumes the basic default settings as a starting point.

General Rules of Thumb for Glossy Wood, Flooring, etc.

These are the kind of "hybrid" materials you might require for architectural renderings; lacquered wood, linoleum, etc.

For these materials, set BRDF to Custom Reflectivity Function; that is, you'll define a custom BRDF curve. Start out with 0 degree reflectivity of 0.2, 90 degree reflectivity of 1.0, and apply a suitable texture map to the Diffuse Color. Set Reflectivity between 0.6 and 1.0.

How glossy is the material? Are reflections clear or blurry? Are they strong or weak?

- For clear, fairly strong reflections, keep Reflection Glossiness at 1.0.
- For slightly blurry but strong reflections, set a lower Reflection Glossiness value. If performance becomes an issue try turning on Fast (interpolate).
- For slightly blurry but also very weak reflections, you can "cheat" by applying a lower Reflection Glossiness value for broader highlights while setting Reflection Glossy Samples to 0. This shoots only one mirror ray

for reflections, but if they are very weak, the viewer can often not really tell.

- For moderately blurry surfaces, set an even lower Reflection Glossiness value and maybe increase the Reflection Glossy Samples value. Again, for improved performance turn on Fast (interpolate).
- For extremely blurry surfaces or surfaces with very weak reflections, try turning on Highlights+FG Only.

A typical wooden floor could use Reflection Glossiness of 0.5, Reflection Glossy Samples of 16, Reflectivity of 0.75, a nice wood texture for Diffuse Color, perhaps a slight bump map. If bumpiness should appear only in the lacquer layer, turn on Special Purpose Maps rollout > Do Not Apply Bumps To The Diffuse Shading.

Linoleum flooring could use the same settings but with a different texture and bump map, and probably with slightly lower Reflectivity and Reflection Glossiness values.

Ceramics

Ceramic materials are glazed; that is, they're covered by a thin layer of transparent material. They follow rules similar to the general materials mentioned above, but choose *By IOR (fresnel reflections)* and set IOR to about 1.4 and Reflectivity to 1.0.

Set the Diffuse Color to a suitable texture or color, such as white for white bathroom tiles.

Stone Materials

A stone object usually has a fairly matte finish, or has reflections that are so blurry they are nearly diffuse. You can simulate the "powdery" character of stone with the Diffuse Roughness parameter; try 0.5 as a starting point. Porous stone such as brick would have a higher value.

Stone would have a very low Reflection Glossiness (lower than 0.25) and one can most likely use Highlights+FG Only to good effect for very good performance. Use a nice stone texture for Diffuse Color, some kind of bump map, and perhaps a map that varies the Reflection Glossiness value.

The Reflectivity would be around 0.5-0.6 with By IOR (fresnel reflections) off and 0 degree reflectivity at 0.2 and 90 degree reflectivity at 1.0

Glass

Glass is a dielectric, so By IOR (fresnel reflections) should definitely be on. The IOR of standard glass is 1.5. Set Diffuse Level to 0.0, Reflectivity to 1.0 and Transparency to 1.0. This is enough to create basic, completely clear refractive glass.

If this glass is for a window pane, turn Thin-walled on. If this is a solid glass block, turn Thin-walled off and consider if caustics are necessary or not, and set Refractive Caustics accordingly.

If the glass is frosted, set Refraction Glossiness to a suitable value. Tune the Refraction Samples for good quality or turn on Fast (interpolate) for performance.

Colored glass

For clear glass, use the tips in the preceding section. Colored glass, however, is a different story.

Many shaders set the transparency at the surface of the glass. And indeed this is what happens if one simply sets a Refractive Color to some value, such as blue. For glass with Thin-walled turned on this works perfectly. But for solid glass objects this is not an accurate representation of reality.

The scene in following illustration contains two glass blocks of different sizes, a sphere with a spherical hole inside it, and a glass horse.

Note: The spherical hole was created by inserting a second sphere with its normals flipped inside the

outer sphere. Don't forget to flip the normals of such surfaces or they will not render correctly.



With a blue refraction color: Glass with color changes at the surface

The problems are evident:

- The two glass blocks are of different thicknesses, yet they are exactly the same level of blue.
- The inner sphere is darker than the outer one.

Why does this happen?

Consider a light ray that enters a glass object. If the color is located at the surface, the ray is colored somewhat as it enters the object, retains this color through the object, and receives a second coloration (attenuation) when it exits the object:

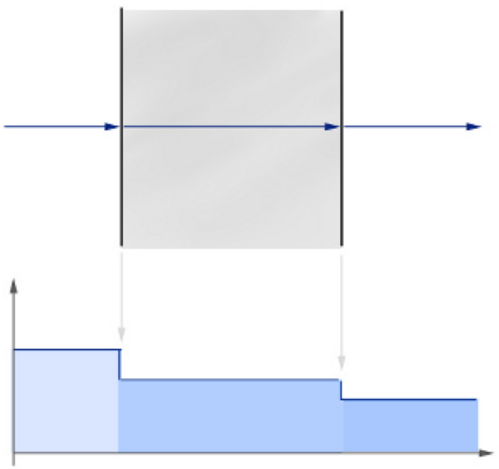


Diagram for glass with color changes at the surface

In the above illustration the ray enters from the left, and at the entry surface it drops in level and gets slightly darker (the graph illustrates the level schematically). It retains this color throughout its travel through the medium and then drops in level again at the exit surface.

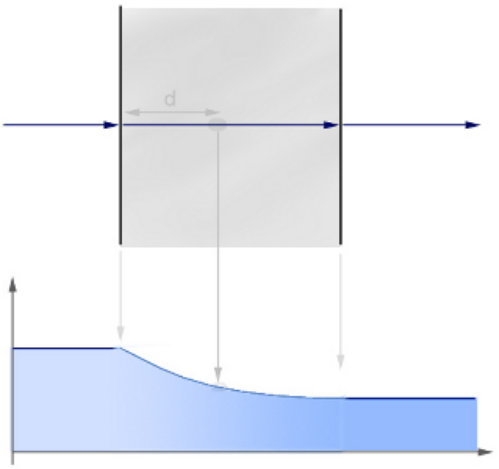
For simple glass objects this is quite sufficient. For any glass using *Thin-walled* it is by definition the correct thing to do, but for any complex solid it is not. It is especially wrong for negative spaces inside the glass (like the sphere in our example) because the light rays have to travel through four surfaces instead of two, getting two extra steps of attenuation at the surface.

In real colored glass, light travels through the medium and is attenuated as it goes. In the Arch & Design material this is accomplished by turning on Advanced Rendering Options > Refraction > Max Distance, setting the Color At Max Distance, and setting the Refraction Color to white. This is the result:



Glass with color changes within the medium

The result is clearly much more satisfactory: The thick glass block is a deeper blue than the thin one, and the hollow sphere now looks correct. In diagram form it looks as follows:



d=Max Distance where attenuation is Color at Max Distance

The ray enters the medium and is attenuated throughout its travel. The strength of the attenuation is such that precisely the Max Distance (d in the figure) the attenuation matches that of Color At Max Distance. In other words, at this depth the attenuation is the same as was received immediately at the surface with the previous scene. The falloff is exponential, so at double the Max Distance value the effect is that of Color At Max Distance squared, and so on.

There is one minor tradeoff:

To render the shadows of a material correctly using this method, you must either use caustics or make sure mental ray is rendering shadows in Segments mode (see *Shadows and Displacement Rollout (mental ray Renderer) (page 3–114)*).

Using caustics naturally gives the most correct-looking shadows (the above image was rendered without caustics), but requires that the scene has caustic photons enabled and contains a physical light source that shoots caustic photons.

On the other hand, the mental ray Segments shadows have a slightly lower performance than the more widely used Simple shadow mode. But if it is not used, the shadow intensity will not take the attenuation through the media into account properly. However, the image might still look pleasing.

Water and Liquids

Water, like glass, is a *dielectric* with an IOR of 1.33. Hence, the same principles as for glass (above) apply to bodies of water, which truly need to refract their environment. An example is water running from a tap. Colored liquids use the same principles as colored glass.



Water into wine

To create a liquid in a container, as in the preceding image, it is important to understand how the Arch & Design material handles refraction through multiple surfaces vs. the real-world behavior of light in such circumstances.

What is important for refraction is the transition from one medium to another with a different IOR. Such a transition is known as an *interface*.

For lemonade in a glass, imagine a ray of light travelling through the air (IOR=1.0). When it enters the glass, it is refracted by the IOR of the glass (1.5). The ray then leaves the glass and enters the liquid; that is, it passes through an interface from a medium of IOR 1.5 to another medium of IOR 1.33.

One way to model this in computer graphics is to make the glass one separate closed surface, with the normals pointing outward from the surface of

the glass and an IOR of 1.5, and a second, closed surface for the liquid, with the normals pointing outward and an IOR of 1.33, leaving a small air gap between the container and the liquid.

This approach works, but can cause a problem: When light goes from a higher IOR to a lower there is a chance of an effect known as *total internal reflection* (TIR). This is the effect you see when diving into a swimming pool and then looking up: You can see the objects above the surface only in a small circle straight above. Anything below a certain angle shows only a reflection of the pool and things below the surface. The larger the difference in the IOR of the two media, the greater the chance of TIR.

So in our example, as the ray goes from glass (IOR=1.5) to air, there is a large chance of TIR. But in reality the ray would move from a medium of IOR=1.5 to one of IOR=1.33, which is a much smaller step with a much smaller chance of TIR. This looks different:



Left: Correct refraction; Right: the "air gap" method

The result on the left is the correct one, but how it is obtained?

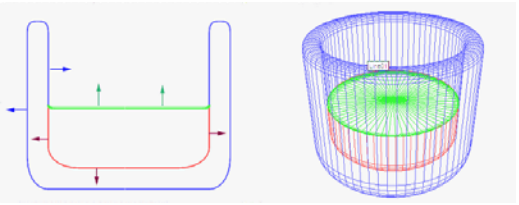
The solution is to rethink the modeling, and not to think in terms of media, but in terms of interfaces. In our example, we have three different interfaces, where we can consider the IOR as the ratio between the IORs of the outside and inside media:

- air-glass interface (IOR=1.5/1.0=1.5)
- air-liquid interface (IOR=1.33/1.0=1.33)
- glass-liquid interface (IOR=1.33/1.5=0.8)

In the most common case of an interface with air, the IOR to use is the IOR of the media (because the IOR of air is 1.0), whereas in an interface between two different media, the situation is different.

To correctly model this scenario, then, we need three surfaces, each with a different Arch & Design material applied:

- the air-glass surface (blue), with normals pointing out of the glass, covering the area where air directly touches the glass, having an IOR of 1.5
- the air-liquid surface (green), with normals pointing out of the liquid, covering the area where air directly touches the liquid, having an IOR of 1.33
- the glass-liquid surface (red), with normals pointing out of the liquid, covering the area where the glass touches the liquid, having an IOR of 0.8



The three interfaces for a liquid in a glass

By setting suitable Max Distance and Color At Max Distance values for the two liquid materials (to get a colored liquid), we obtain the glass on the left in the preceding rendered image.

Ocean and Water Surfaces

A water surface is a slightly different matter than a visibly transparent liquid.

The ocean isn't blue; it is *reflective*. Not much of the light that penetrates the surface of the ocean gets anywhere of interest. A small amount of light is scattered back up again, doing a bit of literal subsurface scattering.

To make an ocean surface with the Arch & Design material, follow these steps:

1. Set Diffuse Level to 0.0, Reflectivity to 1.0, and Transparency to 0.0. That's right: No refraction is necessary.
2. Set IOR to 1.33 and turn on By IOR (fresnel reflections). Apply an interesting wobbly shader to Bump (Ocean (lume) works well here) and your ocean is basically done.

This ocean has reflections guided *only* by the IOR. But this might work fine; try it. Just make sure there is something there for it to reflect. Add a sky map, objects, or a just a blue gradient background. There must be something or it will be completely black.



The ocean isn't blue; the sky is.

For a more tropical look, try setting Diffuse Color to a slightly blue-green color, set the Diffuse Level to a fairly low number such as 0.1, and turn on Do Not Apply Bumps To The Diffuse Shading.

Now you have a base color in the water that emulates the small amount of scattering that occurs in the top level of the ocean.



Enjoy the tropics.

Metal

Metals are reflective, which means they need something to reflect. The best looking metals come from having a true HDRI environment, either from a spherically mapped HDRI photo, or something like the mental ray physical sky.

To create classic chrome, turn off By IOR (fresnel reflections), set Reflectivity to 1.0, 0 degree reflectivity to 0.9 and 90 degree reflectivity to 1.0. Set Diffuse Color to white, and turn on Metal Reflections.

This creates an almost completely reflective material. Tweak the Reflection Glossiness parameter for various levels of blurry reflections. Also consider using the *Round Corners effect (page 2–1556)*, which tends to work very well with metallic objects.

Metals also influence the color of their reflections. Because you turned on Metal Reflections, this is already happening; try setting the Diffuse Color to a golden color to create gold.

Try various levels of Reflection Glossiness (with the help of Fast (interpolate) for performance, when necessary).

You can also change the Reflectivity value. This has a slightly different meaning when Metal Material is on; it blends between the reflections (colored by the Diffuse Color) and normal diffuse shading. This allows a blend between the glossy reflections

and the diffuse shading, both driven by the same color. For example, an aluminum material would need a bit of diffuse blended in, whereas chrome would not.



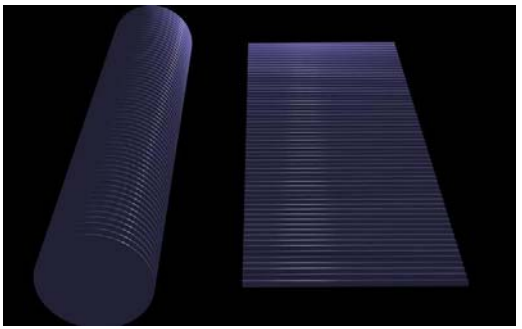
Gold, silver, and copper

Brushed Metal

Brushed metal is an interesting special case. In some cases, creating a brushed metal requires only turning down the Reflection Glossiness to a level where you obtain a very blurred reflection. This is sufficient when the brushing direction is random or the brushes are too small to be visible even as an aggregate effect.

For materials that have a clear brushing direction and/or where the actual brush strokes are visible, creating a convincing look is slightly more involved.

The tiny grooves in a brushed metal surface all work together to cause anisotropic reflections. This can be illustrated by the following schematic, which simulates the brush grooves by modeling many tiny adjacent cylinders, shaded with a simple Phong shader:



Many small adjacent cylinders

As you can see, the specular highlights in the cylinders work together to create an aggregate effect which is the anisotropic highlight.

Also note that the highlight isn't continuous; it is actually broken up into small, adjacent segments. So the primary visual cues that a material is brushed metal are:

- Anisotropic highlights that stretch out in a direction perpendicular to the brushing direction
- A discontinuous highlight with breaks in the brushing direction

Many attempts to simulate brushed metals have looked only at the first effect: the anisotropy. Another common mistake is to think that the highlight stretches in the brushing direction. Neither is true.

Hence, to portray brushed metals, it is necessary to simulate these two visual cues. The first is simple: Use Anisotropy and Anisotropy Rotation to make anisotropic highlights. The second can be done in several ways:

- with a bump map
- with a map that varies the Anisotropy or Reflection Glossiness values
- with a map that varies the Reflection Color

Each has advantages and disadvantages, but the one we suggest here is the last one. The reason for

choosing this method is that it works well with interpolation.

1. Create a map for the brush streaks. The possible ways to do this include painting a map in a paint program, or using a Noise map that has been stretched heavily in one direction. The map should vary between middle-gray and white.
2. Apply this map to the Reflection Color in a scale suitable for the brushing.
3. Set Diffuse Color to white (or the color of the metal) but set Diffuse Level to 0.0 (or a small value).
4. Make sure Metal Material is on.
5. Set Reflection Glossiness to 0.75.
6. Set Anisotropy to 0.1 or a similar value. Use Anisotropy Rotation to align the highlight properly with the map. If necessary use Anisotropy Channel to base it on the same texture space as the map.



Brushed metal

Car Paint Material/Shader (mental ray)

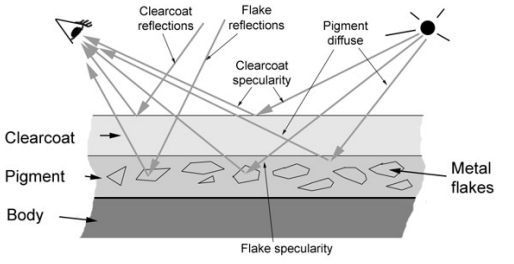
Material Editor > Type button > Material/Map Browser > Car Paint Material

Material Editor > any material > Click a Map button. > Material/Map Browser > Car Paint Shader

Note: The Car Paint material and shader appear in the Browser only if the mental ray renderer is the currently active renderer.

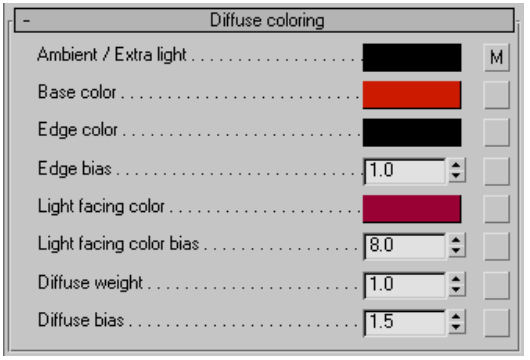
Car Paint is available as both a mental ray material and shader; both have identical parameters, and support the following unique characteristics of real-world car paint:

- The lowest surface, applied directly to the car body, is a thin layer of pigment. The properties of this layer are such that the perceived color changes depending on the viewing angle as well as the incident angle of the incoming light.
- Tiny metal flakes are suspended within this layer. The flakes reflect light and can be seen glittering on a sunny day, due to individual flakes reflecting sunlight directly at the observer.
- On top of this is a clear-coat layer, which can be more or less reflective and more or less glossy, depending on the quality of the layer and any added wax coating. Most notably, this layer tends to exhibit a pronounced Fresnel effect, reflecting more light at glancing angles.
- Includes a Lambertian dirt layer for that "unwashed" look.



Interface

Diffuse coloring rollout



Ambient/Extra light—The ambient light component.

Note: This parameter is treated differently from the ambient/ambience parameter pair of many other base shaders in that it is influenced by the additional Diffuse Coloring parameters, and hence represents incoming light, rather than the object's "ambient color."

Base color—The base diffuse color of the material.

Edge color—The color seen at glancing angles (that is, edges), which tends to appear much darker. For deep metallic paints seen on sports cars it tends to be almost black.

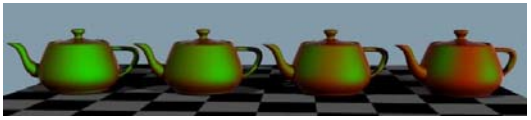
Edge bias—The falloff rate of the color towards the edge. Higher values make the edge region narrower; lower values make it wider. The useful range is 0.0 to approximately 10.0, where the value 0.0 turns the effect off.



Color shift due to view angle, shifting between a red base color and a blue edge color (atypical colors chosen for demonstration purposes) with varying Edge Bias values

Light facing color—The color of the area facing the light source.

Light facing color bias—The falloff rate of the color towards the light. Higher values make the colored region facing the light smaller/narrower; lower values make it larger/wider. The useful range is 0.0 to approximately 10.0, where the value 0.0 turns the effect off.

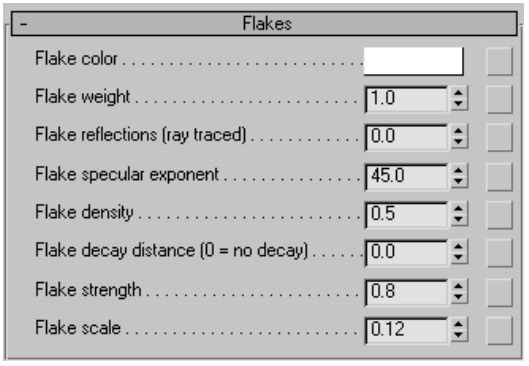


Color shift due to view angle, shifting between a red base color and a green light facing color (atypical colors chosen for demonstration purposes) with varying Light Facing Color Bias values

Diffuse weight—Controls the overall level of the Diffuse Coloring parameters.

Diffuse bias—Modifies the falloff of the diffuse shading. Higher values push the diffuse peak towards the light source, and lower values flatten the diffuse peak. The useful range is approximately 0.5 to 2.0, where 1.0 represents standard Lambertian shading.

Flakes rollout



Flake color—The color (reflectivity) of the flakes, which is generally white.

Flake weight—A scalar multiplier for the flake color.

Flake reflections (ray traced)—The amount of ray-traced reflection in the flakes, which allows

glittery reflections of, for example, an HDRI environment. The value 0.0 turns the effect off.

This effect should generally be very subtle; a value of 0.1 is often enough. The final intensity of reflections also depends on the Flake Color and Flake Weight values.

Flake specular exponent—The Phong specular exponent for the flakes.

Flake density—The density of the flakes. The useful range is from 0.1 to approximately 10.0, where lower values indicate sparser flakes and higher values indicates denser flakes.

Flake decay distance—The distance at which the influence of the flakes fades out. A value of 0.0 disables fading. Any positive value causes the Flake Weight value to be modulated so that it reaches zero at this distance.

Because flakes are relatively small, using can introduce rendering artifacts if their visual density becomes significantly smaller than a pixel. If the oversampling of the rendering is set high, small flakes can also potentially trigger massive oversampling and hence overlong rendering times needlessly, because the averaging caused by the oversampling will essentially cancel out the flake effect. If you experience these issues, use Flake Decay Distance to counteract them.



Flakes at different distances with no flake decay. The farthest flakes might cause flicker in animations, or trigger unnecessary oversampling and long render times (rendered here with low oversampling for illustrative purposes).

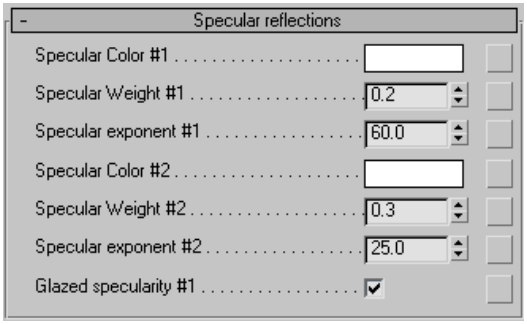


Using flake decay. The flake strength diminishes with distance. The same intentionally low oversampling as in the previous image has been used.

Flake strength—The difference between the orientation of the flakes. The useful range is 0.0 to 1.0 where 0.0 means that all flakes are parallel to the surface, while higher values vary the orientation of flakes increasingly.

Flake scale—The size of the flakes. The procedural texture is calculated in object space, and will hence follow the object. Thus, the scale is influenced by any scale transformation on the object.

Specular reflections rollout



Specular Color #1—The color of the primary specular highlight.

Specular Weight #1—A scalar multiplier applied to Specular Color #1.

Specular exponent #1—The Phong exponent of Specular Color #1.

Specular Color #2—The color of the secondary specular highlight.

Specular Weight #2—A scalar multiplier applied to Specular Color #2.

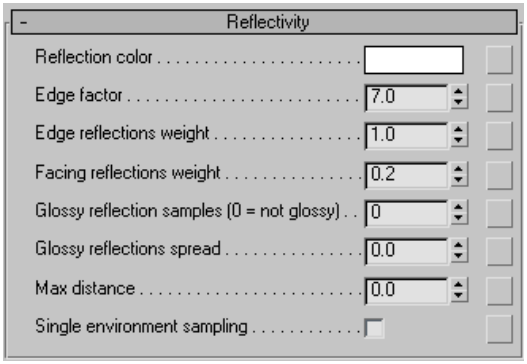
Specular exponent #2—The Phong exponent of Specular Color #2.

Glazed specularity #1—Enables a special mode on the primary specular highlight called *glazing*. By applying a threshold to the specular highlight, it makes the surface appear more polished and shiny. For a new sports car with a lot of wax, turn this on. For a beat-up car in the junkyard, turn it off.



Left to right: Flake specularity only; standard specularity; "glazed" mode enabled; "glazed" mode specularity with flakes

Reflectivity rollout



Reflection color—The color of the reflections in the clear-coat layer. This is generally white.

Edge factor—Clear coat tends to reflect more at glancing angles (edges). This parameter defines the "narrowness" of this edge.

Edge reflections weight—The reflective strength at the edge (generally 1.0).

Facing reflections weight—The reflective strength at facing angles (generally low: 0.1 - 0.3).

Glossy reflection samples—Enables a glossy clear coat. This parameter sets the number of glossy reflection rays traced. A value of 0 disables glossiness.

Glossy reflections spread—Sets the amount of glossiness. Cars are generally near-mirrors so this value should be kept small.

Max distance—Limits the reach of reflective rays.

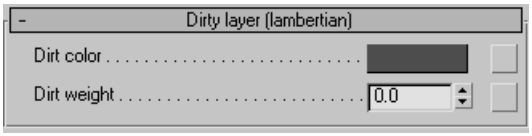
Single environment sampling—Optimizes lookup of environment maps.

Dirty layer (lambertian) rollout



Real cars are rarely clean. This shows the dirt layer (hand-painted dirt-placement map), including a bump map applied in the dirty regions.

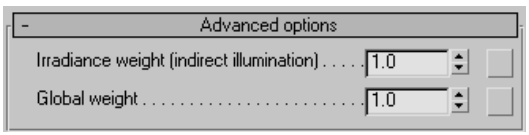
A simple Lambertian dirt layer covers the underlaying paint and clear-coat layers.



Dirt color—The color of the dirt.

Dirt weight—The amount of dirt in the layer. This is typically connected to a texture shader to obtain variations in the dirt across the surface. If the value is 0.0, no dirt is added.

Advanced options rollout



Irradiance weight (indirect illumination)—The influence of indirect light (photons and final gathering) on the surface. It is internally divided by pi (3.14159); for example, a value of 1.0 means the standard 1.0/pi weight.

Global weight—A global tuning parameter that affects the entire diffuse, flake, and specular subsystems. It does not affect reflections or dirt.

Shaders rollout

This rollout enables application of maps or shaders to any of the Car Paint parameters. Of course, you can apply a shader to a parameter by clicking its Map button, so the principal value of this rollout is that it also lets you toggle a parameter's shader, using the check box, without removing the map.

DGS Material (mental ray)

Material Editor > Type button > Material/Map Browser > DGS Material (physics_phen)

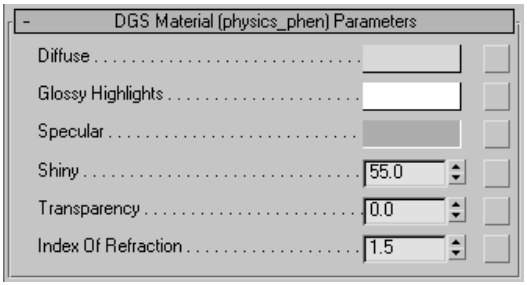
Note: The DGS material appears in the Browser only if the mental ray renderer is the currently active renderer.

DGS stands for Diffuse, Glossy, Specular. This material is a mental ray phenomenon (a scripted shader tree) that provides a physically accurate simulation of a surface. Using DGS materials with the mental ray renderer is comparable to using the *Architectural material (page 2–1535)* with the default scanline renderer.

The DGS material's basic components appear on the Parameters rollout. You can assign shaders to these components, creating a material of greater complexity. Buttons for assigning shaders are on both the Parameters rollout and the Shaders rollout.

Interface

Parameters rollout



Note: The button to the right of each control is a shortcut shader button. Clicking one of these buttons displays the *Material/Map Browser (page 2–1412)* so you can assign a shader to this component. When a map or a shader has been assigned to a component, this button displays the letter "M," and the comparable button on the Shaders rollout displays the map or shader name.

Diffuse—Click the color swatch to display a *Color Selector (page 1–161)* and change the material's diffuse color.

Glossy Highlights—Click the color swatch to display a Color Selector and change the color of glossy highlights.

Specular—Click the color swatch to display a Color Selector and change the color of mirror reflections. When the specular color is white, the material is 100 percent reflective, like a mirror. When the specular color is black, the material does not reflect any of its surroundings.

Shiny—Sets the width of glossy highlights. The larger this value, the smaller the highlights. Default=30.0.

Transparency—Specifies the transparency. The effective range of Transparency is from 0.0 to 1.0. At 0.0 the material is fully opaque. At 1.0 it is fully transparent. Default=0.0.

Warning: You can set the value of Transparency to be greater than 1.0, but this has no effect. An anomaly of the user interface for shaders in the mental ray and lume libraries is that spinner values are not "clamped" to lie within their effective ranges, as they are for controls in 3ds Max.

The value of Transparency also indirectly specifies the reflectivity of the material, which is calculated as 1.0 minus the Transparency value.

Index Of Refraction—Specifies the IOR. In the physical world, the IOR results from the relative speeds of light through the transparent material and the medium the eye or the camera is in. Typically this is related to the object's density. The higher the IOR, the denser the object. Default=1.5.

See *Extended Parameters Rollout (Standard Material) (page 2–1471)* for a list of IOR values for commonly encountered materials.

Note: When the IOR equals 1.0, there is no refraction, and calculating the transparency can take less time than when the material is refractive.

Shaders rollout

The controls on this rollout let you assign a map or shader to one of the basic parameters of the DGS material. This is comparable to mapping a component of a standard material; by adding shaders, you can create a shader tree that generates complex effects.



Click the button for a component to display the *Material/Map Browser (page 2–1412)* and assign the map or shader. Use the toggle at the left to turn the effect of the map off or on.


The button to the right of each main shader button is for shaders that can return multiple parameters. If a shader that returns multiple parameters is assigned to the component, the button's tooltip shows the parameter name. Clicking the button displays a *Connect Parameter To Shader dialog (page 2–1713)*, which lets you change which parameter is being used.

Important: *UV Coordinates (page 2–1728)* and *XYZ Coordinates (page 2–1730)* are the only shaders with multiple return values provided with 3ds Max. You might encounter multiple return values in shaders provided with other shader libraries or custom shader code.

For all the DGS material components, the available mental ray shaders are the same:

Shader	Library
Bump (page 2–1716)	3ds Max
DGS Material (page 2–1717)	3ds Max
Dielectric	base
Dielectric Material (page 2–1719)	3ds Max
Edge	lume
Facade	lume
Glass	lume
Glow	lume
Landscape	lume
Material to Shader (page 2–1723)	3ds Max
Metal	lume
Ocean	lume
Opacity	base
Reflect	base
Refract	base
Shader List (page 2–1723)	3ds Max
Stain	lume
Translucency	lume

Shader	Library
Transmat	physics
Transparency	base
Two Sided	base
UV Generator (page 2–1724)	3ds Max
Water Surface	lume
Wet-Dry Mixer	lume
XYZ Generator (page 2–1729)	3ds Max

Glass Material (mental ray)

Material Editor > Type button > Material/Map Browser > Glass (physics_phen)

Note: The Glass material appears in the Browser only if the mental ray renderer is the currently active renderer.

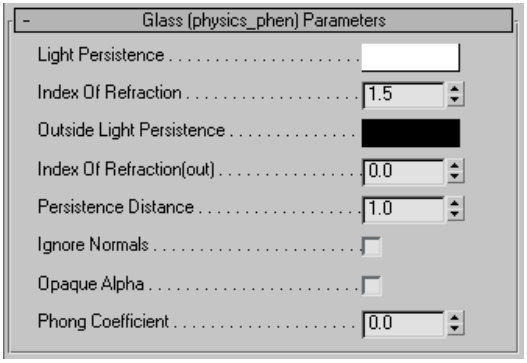
The Glass material simulates both the surface properties and the light-transmitting (photon) properties of glass. It is a mental ray phenomenon (a scripted shader tree) that is equivalent to a *mental ray material (page 2–1544)* with a *Dielectric Material shader (page 2–1719)* assigned to both its Surface and Photon components, with the parameter settings identical for both. A dielectric material, such as glass, is a material whose surface transmits most light that strikes it at angles close to perpendicular (90 degrees), but reflects most light that strikes at glancing angles (close to zero degrees).

Note: This material does not use a shadow shader, so shadows will always be opaque unless you generate caustics.

Adjacent Refractive Materials

Two controls, Outside Light Persistence and Index Of Refraction (Out), are for situations where you are modeling two adjacent refractive materials. Consider a drink in a martini glass. The glass has an index of refraction (IOR) of 1.5, while the alcohol in the glass has an IOR of about 1.3. To create a physically accurate model of this situation, use *three* glass materials: one for the glass itself, one for the alcohol, and a third material for the surfaces where they touch each other. For this third material, set the "inside" IOR to 1.3, and the outside IOR to 1.5.

Interface



Light Persistence—In conjunction with the Persistence Distance, controls the percentage of light that the volume transmits. For example, if the color is set to R=G=B=0.5 and the Persistence Distance is set to 2.0, then objects with a thickness of 2.0 units will appear 50 per cent transparent. Default=white (R=G=B=1.0).

Because transparency depends on the thickness of the object, objects with varying thickness show different transparency depending on the angle from which they are viewed.

Index Of Refraction—Specifies the Index Of Refraction (IOR). In the physical world, the IOR results from the relative speeds of light through the transparent material and the medium the eye or the camera is in. Typically this is related to the object's density. The higher the IOR, the denser the object. Default=1.5.

See *Extended Parameters Rollout (Standard Material) (page 2–1471)* for a list of IOR values for commonly encountered materials.

Outside Light Persistence—In conjunction with the Persistence Distance, controls the percentage of light transmitted on the other side of a surface. When set to the default of black, this control has no effect. See the section "Adjacent Refractive Materials," above. Default=black (R=G=B=0.0).

Index Of Refraction (out)—Sets the IOR on the other side of a surface. When set to the default of zero, this control has no effect. See the section "Adjacent Refractive Materials," above. Default=0.0.

Persistence Distance—In conjunction with the Light Persistence color, controls the percentage of light that the volume transmits. It is the distance at which light transmission is reduced to the percentage specified by the Light Persistence RGB values. Default=1.0.

If you specify an Outside Light Persistence color, that setting also uses the Persistence Distance.

Ignore Normals—When on, the renderer does not use normals to decide whether a light ray is entering or leaving the object. Normally, the material uses normals to decide whether a ray is entering or leaving an object. (It is entering if the normal points toward the ray, leaving if the normal points away from the ray.) This can present a problem for rendering objects whose normals are not unified. hen Ignore Normals is on, the material decides whether a ray is entering or leaving the object by counting the number of times the ray has intersected the object. Default=off.

Opaque Alpha—When on, refracted rays that touch the environment don't generate a transparent alpha value. (This is how 3ds Max usually treats the environment.) When off, refracted rays that touch the environment render a transparent alpha value, which can help if you plan to use the rendering as part of a composite. Default=off.

Phong Coefficient—When greater than zero, generates Phong highlights on the material. The highlights appear in the sample slot. In general this value must be greater than 10 for highlights to be apparent. Default=0.0.

Subsurface Scattering (SSS) Materials

Material Editor > Type button > Material/Map Browser > Choose SSS Fast Material (mi), SSS Fast Skin Material (mi), SSS Fast Skin Material+Displace (mi), or SSS Physical Material (mi), and then click OK.

Note: The SSS materials appear in the Browser only if the mental ray renderer is the currently active renderer.

The subsurface scattering (SSS) materials are provided especially to model skin and other organic materials whose appearance depends on more than one layer of light scattering. 3ds Max provides four of these materials. Each material is a top-level wrapper (a "phenomenon") for shaders whose controls are documented in the *Standard mental ray Shader Libraries* document. Click a link to see the mental images documentation for the shader.

Tip: When you follow a link to the documentation for mental images library shaders, scroll up a bit in your browser. The links tend to go past the title of the section, and there might be introductory content above the link location. If the link goes to the beginning of a section, scroll down instead.

Material Name	mi Library Shader Name
SSS Fast Material (mi)	misss_fast_simple_phen
SSS Fast Skin Material (mi)	misss_fast_skin_phen
SSS Fast Skin Material+Displace (mi)	misss_fast_skin_phen_d
SSS Physical Material (mi)	misss_physical

Also see Subsurface Scattering Shaders and the topics that follow it for more general information. See Physically Correct Subsurface Scattering for background information and tips about the Physical material.

Note: The SSS Physical Material can also be used as a shader for the Surface and Photon components of a *mental ray material (page 2–1544)*.

For a downloadable tutorial offering a practical demonstration of using the SSS Fast Skin material, see this Web page .

Light Controls for the SSS Physical Material

The SSS Physical material includes light controls that correspond to the *lights* array in the parameters for the *misss_physical* shader.

Lights—When on, the material is illuminated only by those lights specified in the list. When Lights is turned off, all lights in the scene affect the material. Default=off.

The remaining light controls are available only when Lights is on.

- **List of lights—**Displays the lights you have chosen to illuminate this material.
- **Add—**Adds a light to the list. Click Add to turn it on, then click the light object in a viewport.
- **Replace—**Replaces a light in the list. Highlight a light's name in the list, click Replace to turn it on, then click the replacement light object in a viewport.
- **Delete—**Deletes a light from the list. Highlight a light's name in the list, then click Delete.

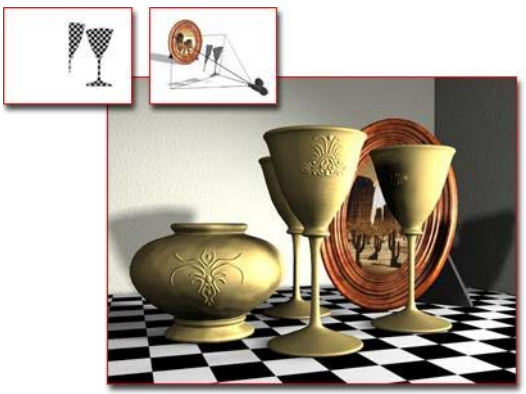
Matte/Shadow Material

Material Editor > Type button > Material/Map Browser > Matte/Shadow

The Matte/Shadow material allows you to make whole objects (or any subsets of faces) into *matte objects (page 3–971)* that reveal the current *environment map (page 3–934)*.

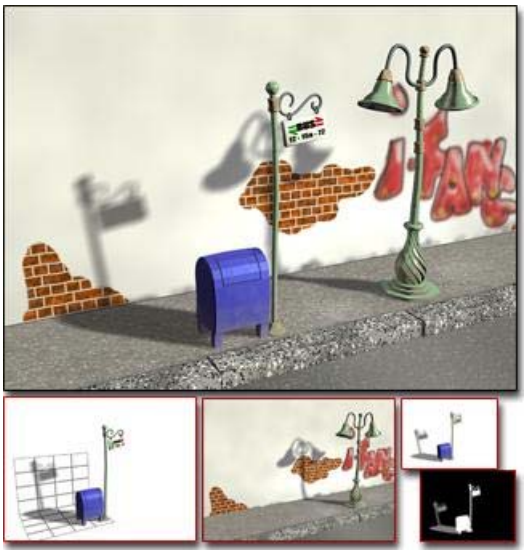


Simply rendering the framed photo against a background shows the photo in front of the background.



A matte object hides parts of the photo, revealing the background to make it appear the photo is behind the goblets.

It can also receive shadows cast on it from non-matte objects in the scene. Using this technique, you can cast shadows on backgrounds by building matte proxy objects and placing them in front of similarly shaped objects in the background.



Creating a matte object for casting shadows against a background image

Matte/Shadow materials can also have reflections.

Note: The Matte/Shadow effect is visible only when you render the scene. It isn't visible in viewports.

Procedures

To render objects seamlessly against a background environment:

There are three ways you can render objects to blend seamlessly into a background environment:

1. By assigning a Matte/Shadow Material
2. By assigning a 100% self-illuminated diffuse texture to an object using Camera Mapping
3. By assigning a 100% self-illuminated diffuse texture using Environment/Screen projection

Use the Plate Match/MAX R2.5 antialiasing whenever you are trying to match foreground objects with an unfiltered background, or trying to match the antialiasing qualities of the 3ds Max 2.5 renderer. See the description of Antialiasing Filter in the description of the *Default Scanline Renderer (page 3–38)*.

Example: To see the effect of Affect Alpha and Matte Reflection:

1. Create a scene with one or more objects on a box platform, and one or more shadow-casting spotlights.

   Tip: If you use only one spotlight, increase its Multiplier value.

2. Assign a Matte/Shadow material to the box, and render the scene with default Matte/Shadow parameters (Opaque Alpha is on, and Affect Alpha is off).
3. In the *rendered frame window (page 3–5)* that is displayed, click Display Alpha Channel.

   The objects all appear as white silhouettes, including the platform.

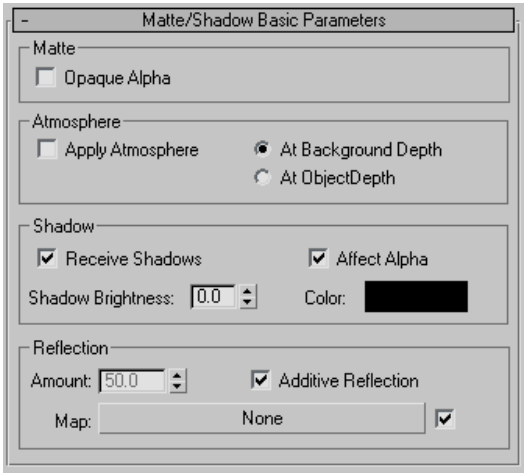
4. On the Matte/Shadow Basic Parameters rollout, turn off Opaque Alpha, and turn on Receive Shadows and Affect Alpha.
5. Press F9 (*Render Last (page 3–25)*).

   The silhouette of the box no longer appears, but the other objects and their shadows appear.

6. On the Matte/Shadow Basic Parameters rollout, click the Map button in the Reflection group, and use the *Material/Map Browser (page 2–1412)* to assign a *Flat Mirror (page 2–1695)* map to the box's material.
7. Press F9 (*Render Last (page 3–25)*).
8. In the rendered frame window, turn off Display Alpha Channel.

   The reflections of the objects appear on the box, even though the box itself is invisible.

Interface



Matte group

Opaque Alpha—Determines whether or not the matte material appears in the alpha channel. If you turn off Opaque Alpha, the matte object will not make an alpha channel, and the image can be used for compositing, just as if there are no matte objects in the scene. Default=off.

Atmosphere group

These options determine whether fog effects are applied to the matte surfaces, and how they are applied.

Apply Atmosphere—Turns the fogging of matte objects on and off.

When applying fog, you can choose between two different methods. You can either apply fog as if the matte surface is at an infinite distance from the camera or you can apply it as if the matte surface is actually at that point on the object being shaded. In other words, you can apply the fog to the matte surface in either 2D or 3D. The following controls determine how this is applied:

At Background Depth—This is the 2D method. The *scanline renderer (page 3–1006)* fogs the scene, and then renders its shadows. In this case, the shadows won't be lightened by the fog. If you want to lighten the shadows, you need to turn up the shadow brightness.

At Object Depth—This is the 3D method. The renderer first renders the shadows, and then fogs the scene. Since this varies the amount of fog over the 3D matte surface, the generated matte/alpha channels don't blend perfectly into the background. Use At Object Depth when the matte object is meant to be a 3D object in the scene that the 2D background represents.

Shadow group

This group determines whether the matte surfaces receive shadows that are cast upon them, and how they receive them.

Receive Shadows—Renders shadows on the matte surfaces. Default=on.

Affect Alpha—When on, shadows cast on a matte material are applied to the alpha channel. This lets you render bitmaps with alpha channels that you can composite later. Default=on.

Affect Alpha is available only when Opaque Alpha (in the Matte group box) is turned off.

When Affect Alpha is on, the higher the Shadow Brightness value, the more transparent the shadow, allowing the background to show through more, and making the shadow appear brighter.

Shadow Brightness—Sets shadow brightness. At 0.5, the shadows will not be attenuated on the matte surface; at 1.0, the shadows are brightened to the color of the matte surface; and at 0.0 they are darkened to completely obliterate the matte surface.

Color—Displays a *Color Selector (page 1–161)* to let you choose the color of the shadow. Default=black.

Setting shadow color is especially useful when you're using a Matte/Shadow material to composite

your shadows against a background image, such as video. It lets you tint your shadows to match pre-existing shadows in the image.

Reflection group

Controls in this group determine whether the matte surfaces can have reflections. You create matte reflections using a shadow map.

Tip: Matte reflections don't successfully create an alpha channel unless you render them against a black background.

Amount—Controls the amount of reflection to use. This is a percentage that can range from 0 to 100. This control is unavailable unless you have assigned a map. Default=50.

You can animate this parameter.

Map—Displays the *Material/Map Browser (page 2–1412)* so you can assign a map to use for reflections. The reflection is independent of the environment unless you choose a *Reflect/Refract (page 2–1695)* or *Flat Mirror map (page 2–1695)*.

Compound Materials

Compound materials combine two or more sub-materials. Compound materials are similar to *compositor maps (page 2–1687)*, but they exist at the material level. Applying a compound material to an object creates a compound effect that often uses mapping. You load or create compound materials using the Material/Map Browser.

Using a filter control, you can choose whether the Browser lists maps or materials or both.

Different types of materials create different effects, behave in particular ways, or are provided as ways to combine multiple materials.

Note: The sub-material buttons and sub-map buttons for most materials and maps have check boxes beside each button. These let you turn that branch of the material or map on or off. For example, in the Top/Bottom material, the Top Material and Bottom Material buttons each have check boxes. Similarly, the Checker map has two map buttons, one for each color. Each button has a check box beside it that lets you disable that color's map.

These are the different types of compound materials:

Blend (page 2–1588): Combines two materials by mixing their pixel colors, as Mix maps do.

Composite (page 2–1589): Mixes up to 10 materials, using additive colors, subtractive colors, or opacity mixing.

Double-Sided (page 2–1591): Stores two materials. One material is rendered on the object's outer faces (the usual side for one-sided materials, as determined by face normals), and the other is rendered on the object's inner faces.

Morpher Material (page 2–1592): Morpher materials use the *Morpher modifier (page 1–729)* to manage multiple materials over time.

Multi/Sub-Object (page 2–1594): Lets you assign more than one material to the same object. Stores two or more sub-materials, which you assign at the sub-object level by using the *Mesh Select (page 1–719)* modifier. You can also assign the sub-materials to whole objects by using the *Material modifier (page 1–714)*.

Shellac (page 2–1597): Superimposes one material on another.

Top/Bottom (page 2–1599): Stores two materials. One material is rendered on an object's top faces and the other is rendered on the object's bottom faces, depending on whether a face normal points up or down.

Procedure

To have the Browser list only materials:

- At the top of the Show group, leave Materials on and turn off Maps.

  Note: This option is available only at the material level in the Material Editor.

Blend Material

Material Editor > Type button > Material/Map Browser > Blend



Blend material combines bricks and stucco.

The Blend material lets you mix two materials on a single side of the surface. Blend has an animatable Mix Amount parameter that lets you draw material *morphing (page 3–974)* function curves to control the way that the two materials are blended over time.

Note: If even one sub-material has its shading set to Wire (see *Shader Basic Parameters Rollout (page 2–1466)*), the *entire material* displays and renders as a wire material.

Procedures

To create a Blend material:

1. Activate a sample slot in the Material Editor.
2. Click the Type button.
3. In the *Material/Map Browser (page 2–1412)*, choose Blend and then click OK.

   A *Replace Map (page 2–1456)* dialog is displayed. This dialog asks whether you want to discard the original material in the slot, or retain it as a sub-material.

   Blend materials have similar controls to Mix maps.

To specify a component material:

- In the Blend Basic Parameters rollout, click one of the two material buttons.

  The parameters for the sub-material are displayed. By default, a sub-material is a Standard material with Blinn shading.

To control the mix amount:

- In the Basic Parameters rollout, adjust the Mix Amount value.

  You can also control the mix amount by using a map.



Map used to reveal brick beneath stucco

To control the mix amount using a map:

- In the Basic Parameters rollout, click the map button next to Mask.

  The Browser is displayed so you can select a map type.

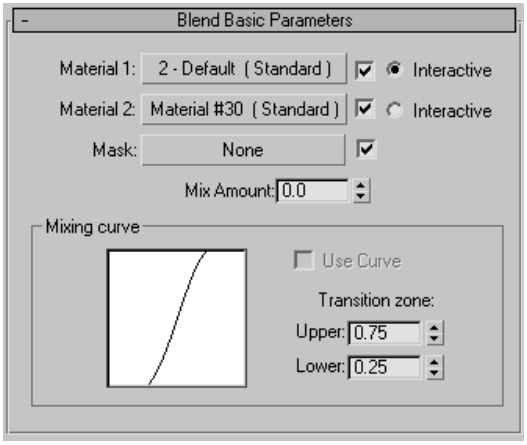
  The intensity of pixels in this mixing map controls the mix. When the intensity is close to zero, one of the component colors or maps

is visible; when it is close to full intensity, the other component is visible.

Tip: Using a Noise map for the mixing map can give good effects that have a natural appearance.

Mix Amount is unavailable while a map is assigned to this parameter. If Use Curve is turned off, the mixing map is used as is. If Use Curve is on, you can shift the effect of the mixing map's gradient ramp to reveal more of one material and less of the other.

Interface



Material 1/Material 2—Set the two materials to be blended. Use the check boxes to turn the materials on and off.

Interactive—Chooses which of the two materials is displayed on object surfaces in viewports by the interactive renderer.

If one material has *Show Map in Viewport (page 2–1445)* on, this takes precedence over the Interactive setting. Only one map at a time can be displayed in viewports.

Mask—Sets a map to use as a mask. The degree of blending between the two materials depends on the intensity of the mask map. Lighter (whiter) areas of the mask show more of Material 1, while darker (blacker) areas of the mask show more of Material 2. Use the check box to turn the mask map on or off.

Mix Amount—Determines the proportion of the blend (percentage). 0 means only Material 1 is visible on the surface; 100 means only Material 2 is visible. Unavailable if you have assigned a mask map and the mask's check box is on.

You can animate this parameter. *Create Material Preview (page 2–1434)* is useful for testing the effect.

Mixing Curve group

The mixing curve affects how gradual or how sharp the transition between the two colors being blended will be. It affects the blend only when a mask map is assigned.

Tip: For mottled effects, blend two standard materials using a noise map as a mask.

Use Curve—Determines whether the Mixing Curve affects the mix. This control is available only when a mask is assigned and active.

Transition Zone—These values adjust the level of the Upper and Lower limits. If the two values are the same, the two materials meet at a definite edge. Wider ranges give more gradual blending from one sub-material to the other. The mixing curve displays the effect of changing these values.

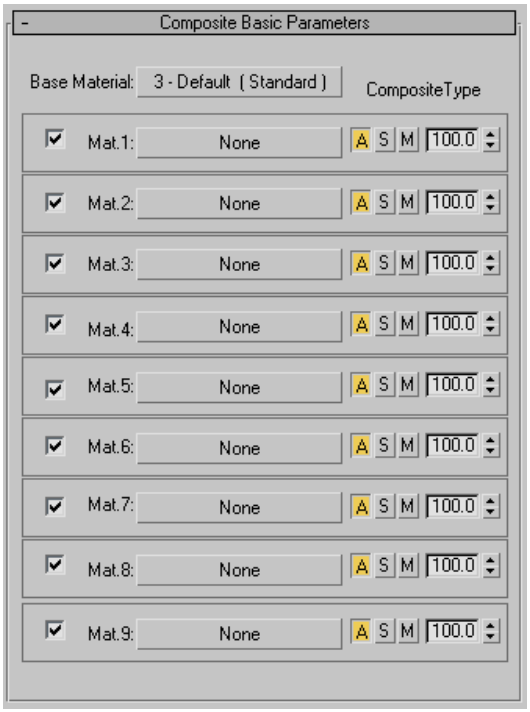
Composite Material

Material Editor > Type button > Material/Map Browser > Composite

Composite material *composites (page 3–922)* up to 10 materials. The materials are superimposed from top to bottom, as listed in the rollout. Materials are combined using *additive opacity (page 3–906)*, *subtractive opacity (page 3–1017)*, or mixed using an Amount value.

Note: If even one sub-material has its shading set to Wire (see *Shader Basic Parameters Rollout (page 2–1466)*), the *entire material* displays and renders as a wire material.

Interface



Base Material—Displays the *Material/Map Browser (page 2–1412)*, where you assign the base material. By default, the base material is a Standard material. The other materials are composited by superimposing them on top of this material, in order from top to bottom.

Mat 1 through Mat 9—Each of these nine groups contains controls for a material to composite. By default, no materials are assigned.

Check box—When on, uses the material in the composite. When off, doesn't use it. Default=on.

Button—Displays the *Material/Map Browser (page 2–1412)*, where you assign a material to composite.

ASM buttons—These buttons control how the material is composited. Default=A.

A—When on, this material uses *additive opacity (page 3–906)*. Colors in the material are summed based on their opacity.

S—When on, this material uses *subtractive opacity (page 3–1017)*. Colors in the material are subtracted based on their opacity.

M—When on, this material mixes materials based on the Amount. Both color and opacity are blended as they are when you use a *Blend material (page 2–1588)* with no mask.

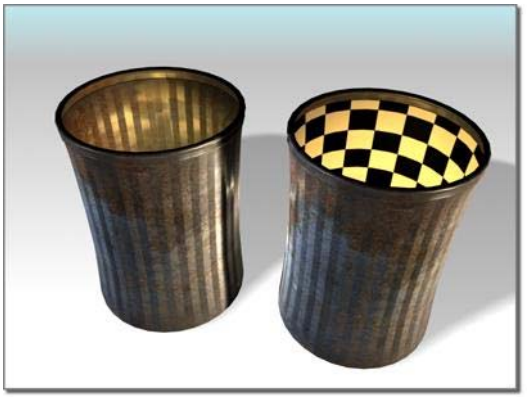
Amount spinner—Controls the amount of mixing. Default=100.0.

For additive (A) and subtractive (S) compositing, the Amount can range from 0 to 200. When the Amount is 0, no compositing happens, and the material below is not visible. When the Amount is 100, the composite is complete. When the amount is greater than 100, compositing is "overloaded": transparent portions of the material become more opaque, until the material below is no longer visible.

For mix (M) compositing, the Amount can range from 0 to 100. When the Amount is 0, no compositing happens, and the material below is not visible. When the amount is 100, compositing is complete, and only the material below is visible.

Double-Sided Material

Material Editor > Type button > Material/Map Browser > Double-Sided



On the right, a double-sided material creates a pattern for the inside of the trash can.

The Double-Sided material lets you assign two different materials to the front and back faces of an object.

Note: If even one sub-material has its shading set to Wire (see *Shader Basic Parameters Rollout (page 2–1466)*), the *entire material* displays and renders as a wire material.

Procedures

To create a double-sided material:

1. Activate a sample slot in the Material Editor.
2. Click the Type button.
3. In the Material/Map Browser, choose Double-Sided and then click OK.

   A *Replace Map dialog (page 2–1456)* is displayed. This dialog asks whether you want to discard the original material in the slot, or retain it as a sub-material.

   The Double-Sided material controls let you choose the two materials, and the translucency of the material overall.

To choose the outer material:

- Click the button labeled Facing Material.

  The parameters for the sub-material are displayed. By default, a sub-material is a Standard material with Blinn shading.

To choose the inner material:

1. Go back to the parent material (parameters for the Double-Sided material).
2. On the Double-Sided Basic Parameters rollout, click the button labeled Back Material.

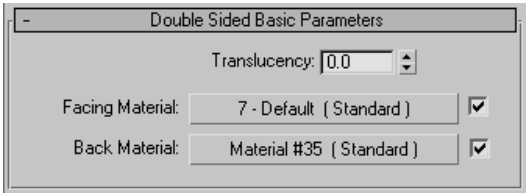
   The parameters for the sub-material are displayed. By default, a sub-material is a Standard material with Blinn shading.

To make the material translucent:

- Set Translucency to a value greater than 0.

  The Translucency control affects the blending of the two materials. When Translucency is 0, there is no blend. When Translucency is 100.0 percent, the outer material is visible on inner faces and the inner material is visible on outer faces. At intermediate values, the specified percentage of the inner material "bleeds through" and is visible on outer faces.

Interface



Translucency—Sets the amount that one material shows through the other. This is a percentage that can range from 0.0 to 100.0. At 100 percent, the outer material is visible on inner faces and the inner material is visible on outer faces. At intermediate values, the specified percentage of

the inner material "bleeds through" and is visible on outer faces. Default=0.0.

You can animate this parameter.

Facing Material and Back Material—Click to display the *Material/Map Browser (page 2–1412)* and choose a material for one side or the other.

Use the check boxes to turn the materials on or off.

Morpher Material

Material Editor > Type button (labeled Standard by default) > Material/Map Browser > Morpher material

Procedures (page 2–1592) Interface (page 2–1593)

The Morpher material works hand-in-hand with the Morpher modifier. You can use it to make the cheeks of a character blush, or to wrinkle a character's forehead when the eyebrows are raised. With the Morpher modifier's channel spinners, you can blend materials the same way you morph the geometry.

The Morpher material has 100 material channels that map directly to the 100 channels in the Morpher modifier. After you apply the Morpher material to an object and bind it to the Morpher modifier, you use the channel spinners in the Morpher modifier to morph materials and geometry. Empty channels in the Morpher modifier, with no geometry morph data, can be used to morph materials only.

Note: The *mental ray renderer (page 3–78)* does not support the Morpher material.

Note: If even one sub-material has its shading set to Wire (see *Shader Basic Parameters Rollout (page 2–1466)*), the *entire material* displays and renders as a wire material.

See also

Morpher Modifier (page 1–729)

Applying the Morpher Material

An object must have at least one Morpher modifier in its modifier stack. You can assign the material to an object and bind it to the object's Morpher modifier in either of two ways.

- After the Morpher modifier is applied to an object, use the Assign New Material command in the Global Parameter rollout of the Morpher modifier. This is the simplest way, and applies the Morpher Material to the object and binds the material to the Morpher modifier at the same time.
- Open the Material Editor, select the Morpher material, and click Choose Morph Object in the Parameters rollout, then click the object in the viewports. After clicking the object, a dialog displays in the viewports, select the Morpher modifier from the dialog (an object may have multiple Morpher modifiers). This binds the Morpher material to the Morpher modifier.

Note: You can bind a Morpher material to only one Morpher modifier.

Procedures

Example: To apply and use the Morpher material:

1. Create a sphere in the Perspective viewport.
2.  On the Modify panel, right-click the sphere's entry in the modifier stack display, and choose Convert To: Editable Mesh.
3. From the Modifier List, choose Morpher.

   This applies the Morpher modifier to the sphere.
4. On the Morpher modifier's Global Parameters rollout, click Assign New Material.

   The Morpher material is now applied to the object and bound to the Morpher modifier.

5.  Open the Material Editor, and click Pick Material from Object (the eyedropper), then click the sphere in the viewports.

   The Material Editor displays the Morpher material parameters.

6. On the Morpher Material Parameters rollout, click the Mat 1 slot.
7. On the Material/Map Browser choose Standard.
8. On the Basic Parameters rollout, click the Diffuse color swatch.
9. On the Color Selector, choose a bright yellow, and close the color selector. Leave the Material Editor open.
10. 

Turn on the Auto Key button, then move the time slider to frame **50**.
11.  Select the sphere, and then open the Modify panel.
12. On the Morpher modifier's Channel List rollout, set the channel 1 spinner to **100**.

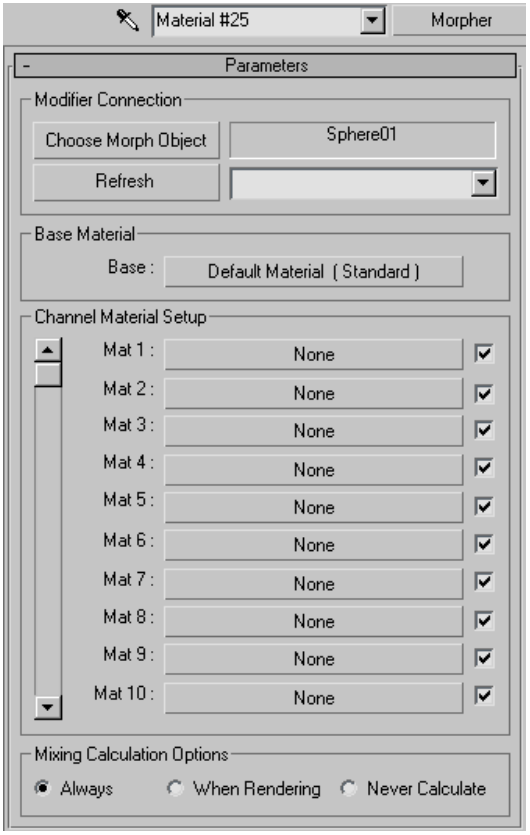
    In the Material Editor, the color of the sample sphere changes to yellow.

13. On the main toolbar, click Render.

    The sphere is yellow. If you render an animation the sphere changes from a grey color to yellow.

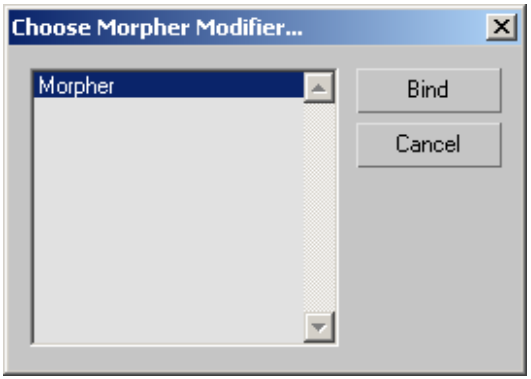
Interface

The Morpher material interface is on a Parameters rollout in the Material Editor.



Modifier Connection group

Choose Morph Object—Click this option, then select an object in the viewports that has a Morpher modifier applied to it. Clicking an object in the viewports displays the Choose Morpher modifier dialog. Choose a Morpher modifier, and click Bind.



Name Field—Displays the name of object to which the Morpher material is applied. If no object has been specified, the field displays "No Target".

Refresh—Updates the channel data.

Marker List—This list is identical to the marker list in the Morpher modifier. Markers you save in the Morpher modifier appear here.

Base Material group

Base material button—Click to apply a base material to the object. The base material represents what the model looks like before any channel blending takes place.

Channel Material Setup group

Map #—100 material channels are available. The scroll bar allows you to scroll through all the channels. Double-click a channel to jump to the material parameters for that channel.

There is a one-to-one correspondence between the channels in the Morpher material and the Morpher modifier. A material in channel 1 of the Morpher material is controlled by the channel 1 spinner in the Morpher modifier.

Material on/off toggle—Turns a channel on and off. Channels that are off do not affect the morph result.

Mixing Calculation Options group

The system can slow down if there are many active materials being blended. Options in this group allow you to control when the morph result will be computed.

Constantly—Choose to compute the material morph result all the time.

When Rendering—Choose to compute the material morph result at render time.

Never Calculate—Choose to bypass material blending.

Multi/Sub-Object Material

Material Editor > Type button > Material/Map Browser > Multi/Sub-Object

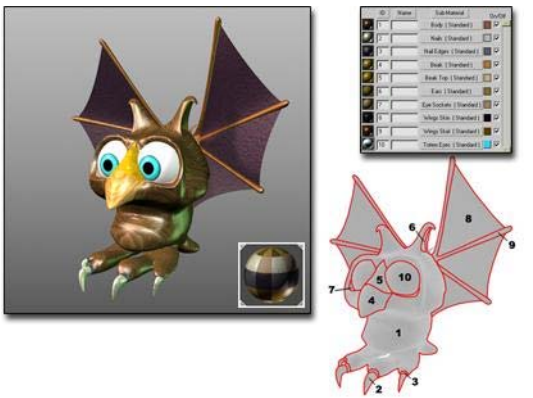


Figure mapped using a Multi/Sub-Object material

The Multi/Sub-Object material lets you assign different materials at the sub-object level of your geometry. You create a multi-material, assign it to an object, and then use the *Mesh Select modifier (page 1–719)* to select faces and choose which of the sub-materials in the multi-material are assigned to the selected faces.

If the object is an *editable mesh (page 1–996)*, you can drag and drop materials to different selections of faces, building a Multi/Sub-Object material on

the fly. See *Drag and Drop Sub-Object Material Assignment (page 2–1424)*.

You can also create a new Multi/Sub-Object material by dragging to faces selected with the *Edit Mesh modifier (page 1–634)*.

Sub-material IDs do not depend on the order of the list, and you can enter new ID values.

The *Make Unique button (page 2–1442)* in the Material Editor lets you make an instanced sub-material into a unique copy.

At the Multi/Sub-Object material level, the sample slot's sample object shows a patchwork of the sub-materials. When you edit a sub-material, the sample slot display depends on the setting of the Simple Multi Display Below Top Level toggle in the *Material Editor Options dialog (page 2–1436)*.

Using Multi/Sub-Object Materials

Here are some usage tips with regards to mesh editing and managing sub-materials.

- When working at sub-object levels of Editable Meshes, Polys, Patches and Splines, or with objects that have Edit Mesh, Spline or Patch modifiers applied to them, you can browse by sub-material names if the object has a multi-sub-object material applied to it.
- Sub-materials that are not assigned to an object, or surface of an object, can be 'cleaned' from the Multi-Sub-Object material by using the *Clean MultiMaterial utility (page 2–1742)*.
- Duplicate maps, assigned to materials, can be changed to instances by using the *Instance Duplicate Maps utility (page 2–1744)*.

Procedures

To create a Multi/Sub-Object material:

1. Activate a sample slot in the Material Editor.
2. Click the Type button.
3. In the *Material/Map Browser (page 2–1412)*, choose Multi/Sub-Object and then click OK.

   A *Replace Map dialog (page 2–1456)* is displayed. This dialog asks whether you want to discard the original material in the slot, or retain it as a sub-material.

   The controls for a Multi/Sub-Object material are essentially a list of the sub-materials it contains.

To assign a sub-material:

- On the Multi/Sub-Object Basic Parameters rollout, click a sub-material button.

  The parameters for the sub-material appear. By default, a sub-material is a Standard material with Blinn shading.

To make one of the sub-materials a solid color:

- On the Multi/Sub-Object Basic Parameters rollout, click the color swatch next to the sub-material button.

  In the *Color Selector (page 1–161)*, choose a color.

  The color swatches for sub-materials are shortcuts. They assign the color you choose to the sub-material's Diffuse component.

To assign one of the sub-materials to a sub-object selection:

1. Select the object, and assign a Multi/Sub-Object material to it.
2.  On the *Modify panel (page 3–758)*, apply *Mesh Select (page 1–719)* to the object.
3. Click Sub-Object and choose Face as the sub-object category.
4. Select the faces to which you will assign a sub-material.

5. Apply a *Material modifier (page 1–714)*, and set the material ID value to the number of the sub-material you want to assign.

   The viewport updates to show the sub-material assigned to the selected faces.

   The material ID values in the Multi/Sub-Object material and the material ID numbers in the Select Face rollout correspond. If you set the ID to a number that doesn't correspond to a material contained in the Multi/Sub-Object material, the faces render as black.

Warning: Some geometric primitives do not use 1 as the default material ID, and some, such as hedra or box, have multiple material IDs by default.

Tip: You can also use the *Edit Mesh modifier (page 1–634)* to assign a contained material to selected faces. Apply Edit Mesh to the object, go to the Face sub-object level, and select the faces to assign. Then on the Edit Surface rollout, set the material ID value to the ID of the sub-material. (You can *drag and drop (page 2–1424)* a Multi/Sub-Object material to an Edit Mesh modifier as you can to an editable mesh object.)

To add a new sub-material:

- Click Add.

  A new sub-material is added to the end of the list. By default, the new sub-material's ID number is one greater than the highest material ID already in use.

To remove a sub-material:

1. Select the sub-material by clicking its small sample sphere in the Multi/Sub-Object Basic Parameters rollout.

   The small sample sphere is surrounded by a black and white border to show the sub-material is selected.

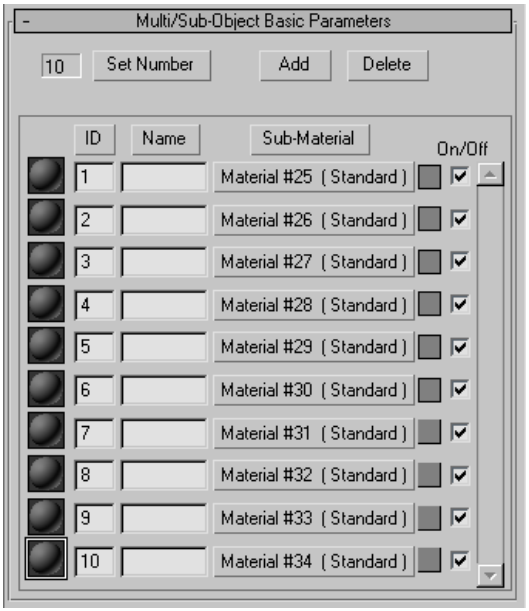
   If the list of sub-materials is longer than the rollout will hold, you can use the scroll bar at the right to display other parts of the list.

2. Click Delete.

   The sub-material is removed.

   Deleting a sub-material is an undoable operation.

Interface



Number—This field displays the number of sub-materials contained in the Multi/Sub-Object material.

Set Number—Sets the number of sub-materials make up the material. At the Multi/Sub-Object material level, the sample slot's sample object shows a patchwork of the sub-materials. (When you edit a sub-material, the sample slot display depends on the setting of the Simple Multi Display Below Top Level toggle on the *Material Editor Options dialog (page 2–1436)*.)

Reducing the number of sub-materials removes sub-materials from the end of the list. You can

undo Set Number when you have used it to delete materials.

Add—Click to add a new sub-material to the list. By default, the new sub-material's ID number is one greater than the highest material ID already in use.

Delete—Click to delete the currently chosen sub-material from the list. You can undo deleting a sub-material.

Sort list controls

These buttons appear above three of the columns in the sub-materials list.

ID—Click to sort the list so it begins with the sub-material that has the lowest material ID, and ends with the sub-material that has the highest material ID.

Name—Click to sort the list by the names you have entered in the Name column.

Sub-Material—Click to sort the list by the sub-material names that appear on the Sub-Material buttons.

List of sub-materials

Each sub-material has a single entry in this list. The rollout displays up to 10 sub-materials at a time. If the Multi/Sub-Object material contains more than 10 sub-materials, you can scroll the list using the scrollbar at the right.

Each sub-material in the list has the following controls:

Small sample sphere—The small sample sphere is a "mini-preview" of the sub-material. Click it to select this sub-material. You must select a sub-material before you delete it.

ID—Shows the ID number assigned to this sub-material. You can edit this field to change the ID number. If you assign two sub-materials the same ID, a warning message appears at the top of the rollout.

When the Multi/Sub-Object material is applied to an object, faces in the object assigned the same material ID number render with this sub-material.

You can click Sort by ID to sort the sub-material list by this value, from lowest to highest.

Note: Sometimes the Sub-Material button shows a material number. This is *not* the sub-material ID.

Name—Lets you enter a custom name for the material. A sub-material name appears in the *Name (page 2–1448)* field when you're at the level of the sub-material. It also appears in the Browser and the Navigator.

Sub-Material button—Click the sub-material button to create or edit one of the sub-materials. Each of the sub-materials is a complete material in its own right, with as many *maps (page 3–968)* and levels as you want.

By default, each sub-material is a *Standard material (page 2–1465)* with *Blinn shading (page 2–1480)*.

Color swatch—Click the color swatch to the right of the Sub-Material button to display the *Color Selector (page 1–161)* and choose a diffuse color for the sub-material.

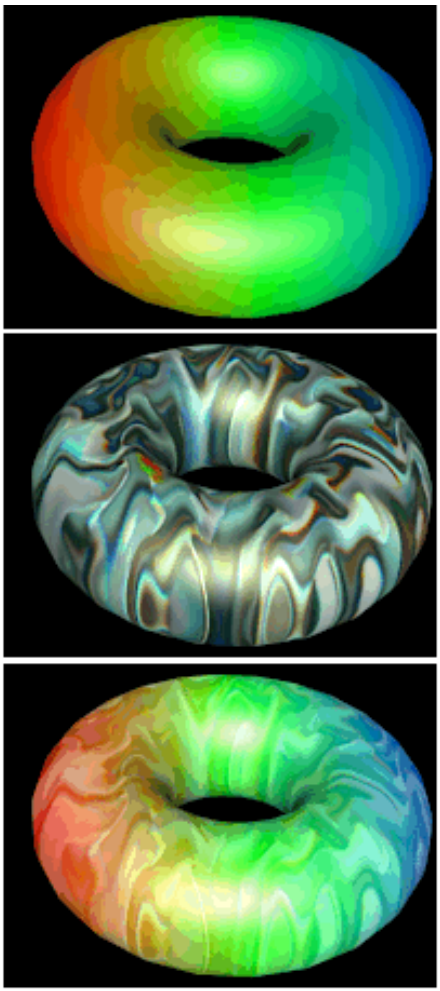
On/Off toggle—Turns the sub-material on or off. When a sub-material is off, it appears black in the sample slot and on objects in the scene. Default=on.

Shellac Material

Material Editor > Type button > Material/Map Browser > Shellac

Shellac material mixes two materials by superimposing one over the other. Colors in the superimposed material, called the "shellac" material, are added to colors in the base material.

A Shellac Color Blend parameter controls the amount of color mixing.



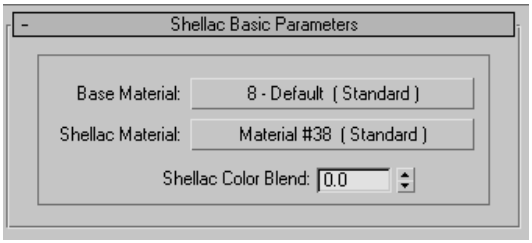
Top: Base material

Middle: Shellac material

Bottom: Materials combined with a shellac color blend value of 50%

Note: If even one sub-material has its shading set to Wire (see *Shader Basic Parameters Rollout (page 2–1466)*), the *entire material* displays and renders as a wire material.

Interface



Base Material—Goes to the level of the base sub-material. By default, the base material is a Standard material with Blinn shading.

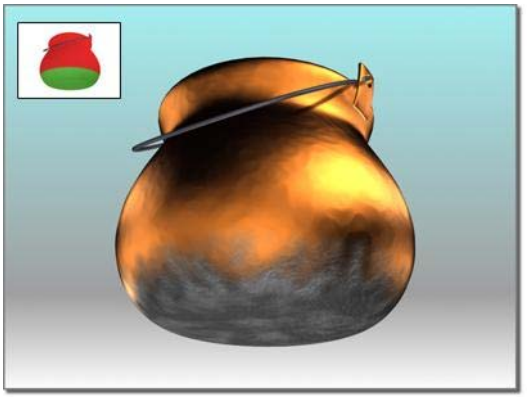
Shellac Material—Goes to the level of the shellac material. By default, the shellac material is a Standard material with Blinn shading.

Shellac Color Blend—Controls the amount of color mixing. At 0.0, the shellac material has no effect. Increasing the Shellac Color Blend value increases the amount of shellac material color blended into the base material color. There is no upper limit on this parameter. Large values "overload" the shellac material colors. Default=0.0.

You can animate this parameter.

Top/Bottom Material

Material Editor > Type button > Material/Map Browser > Top/Bottom



Top/bottom material gives the pot a charred bottom.

The Top/Bottom material lets you assign two different materials to the top and bottom portions of an object. You can blend the materials into one another.

The object's top faces are those whose normals point up. The bottom faces have normals that point down. You can choose whether "up" and "down" refer to the scene's world coordinates or to the object's local coordinates.

Note: If even one sub-material has its shading set to Wire (see *Shader Basic Parameters Rollout (page 2–1466)*), the *entire material* displays and renders as a wire material.

Procedures

To create a top/bottom material:

1. Activate a sample slot in the Material Editor.
2. Click the Type button.
3. In the *Material/Map Browser (page 2–1412)*, choose Top/Bottom and then click OK.

   A *Replace Map (page 2–1456)* dialog is displayed. This dialog asks whether you want to discard the original material in the slot, or retain it as a sub-material.

   The Top/Bottom material controls let you choose the two materials, and also the transition between them.

To choose the top or bottom material:

- On the Top/Bottom Basic Parameters rollout, click the Top Material button or the Bottom Material button.

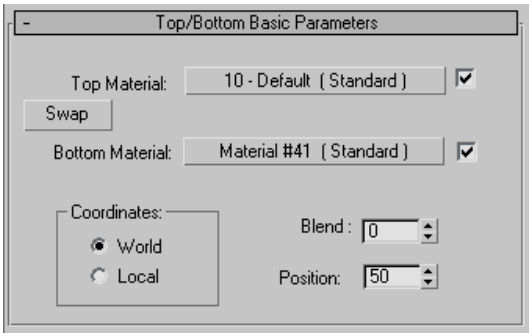
  The parameters for the sub-material appear. By default, a sub-material is a Standard material with Blinn shading.

To swap the two component materials:

- In the Basic Parameters rollout, click Swap.

  The remaining controls, described in the "Interface" section, affect the transition between top and bottom.

Interface



Top Material and Bottom Material—Click to display the parameters for the top or bottom sub-material. By default, a sub-material is a Standard material with Blinn shading.

The check box to the right of each button lets you turn off that material, making it invisible in the scene and in the sample slot.

Swap—Swaps the position of the top and bottom materials

Coordinates group

Controls in this group let you choose how the software determines the boundary between top and bottom.

World—Faces point up or down according to the scene's world coordinates. When you rotate the object, the boundary between top and bottom faces remains in place.

Local—Faces point up or down according to the object's local coordinates. When you rotate the object, the material rotates with it.

Blend—Blends the edge between the top and bottom sub-materials. This is a percentage that can range from 0 to 100. At 0, there is a sharp line between the top and bottom sub-materials. At 100, the top and bottom sub-materials tint each other. Default=0.

You can animate this parameter.

Position—Determines where the division between the two materials lies on an object. This is a percentage that can range from 0 to 100. 0 is at the bottom of the object, and displays only the top material. 100 is at the top of the object, and displays only the bottom material. Default=50.

You can animate this parameter.

Shell Material

Render to a texture. > Material Editor > Pick Material from Object > Click object with "baked" material.

The Shell material is for use with *texture baking (page 3–144)*. When you use Render To Texture to bake a texture, it creates a Shell material that contains two materials: the original material used in the rendering, and the baked material. The baked material is a bitmap that is saved to disk by Render To Texture. It is "baked," or attached to an object in the scene.

The Shell material is a container for other materials, like Multi/Sub-Object. It also lets you control which material is used in which renderings.

Note: The Material/Map Browser lists the Shell material when you assign a new material. You can apply two materials to a single object this way, but changing a material's type to Shell does not generate a baked texture that is saved to disk.

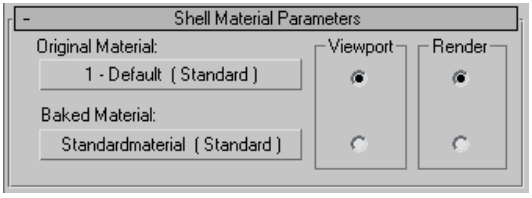
Procedures

To load a shell material into a sample slot:

1. Click an unused sample slot.
2.  Click Pick Material From Object.
3. In a viewport, click an object that has a baked material.

   The sample slot now contains the baked material, and the Shell Material Parameters rollout is displayed.

Interface



Original Material—Displays the name of the original material. Click the button to view that material and adjust its settings.

Baked Material—Displays the name of the baked material. Click the button to view that material and adjust its settings.

In addition to the color and mapping of the original material, the baked material can include shadows from lighting, and other information. Also, a baked material has a fixed resolution.

Viewport—Use these buttons to choose which material appears in shaded viewports: the original

material (upper button) or the baked material (lower button).

Render—Use these buttons to choose which material appears in renderings: the original material (upper button) or the baked material (lower button).

Advanced Lighting Override Material

Material Editor > Type button > Material/Map Browser > Advanced Lighting Override

This material lets you directly control the radiosity properties of a material. Advanced Lighting Override is always a supplement to a *base material,* which can be any renderable material. The Advanced Lighting Override material has no effect on ordinary renderings. It affects the *radiosity solution (page 3–51)* or *light tracing (page 3–44)*. Advanced Lighting Override has two main uses:

- Adjusting the material properties used in a radiosity solution or light tracing
- Creating special effects such as having self-illuminating objects contribute energy to the radiosity solution

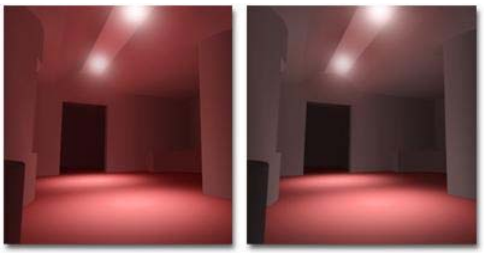
As the rollout for the Advanced Lighting Override material states, you don't have to apply this material to obtain a radiosity solution, and most models will never require it.

Important: The *mental ray renderer (page 3–78)* does not support the Advanced Lighting Override material.

Obtaining a Better Image

Materials that use default settings can be highly reflective. This can lead to overexposed or washed-out radiosity solutions. In general, the best way to adjust this is to reduce the HSV Value (V) of a material color; or, for a bitmapped material, reduce the RGB Level. In some situations, Radiosity Override can improve the appearance of the radiosity solution. Examples of situations where Radiosity Override can help include color bleeding and large dark areas:

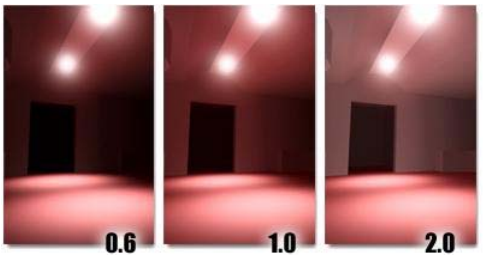
- You might want to reduce Reflectance Scale or Color Bleed when a large area of color (for example, a red carpet in a room with white walls) creates excessive color bleeding. This might be physically accurate, but the eye adjusts for such effects, and the radiosity result might look better with less reflectance or less color bleeding.



Left: Excessive bleeding of the floor color onto the walls and ceiling.

Right: Radiosity Override material reduces the floor's reflectance, causing less bleeding.

- You might want to increase Reflectance Scale when the scene includes a large dark area (for example, a black floor). This can lead to a very dark radiosity result. You can maintain the floor's color but increase reflectance, giving the solution the colors you want while increasing its brightness.

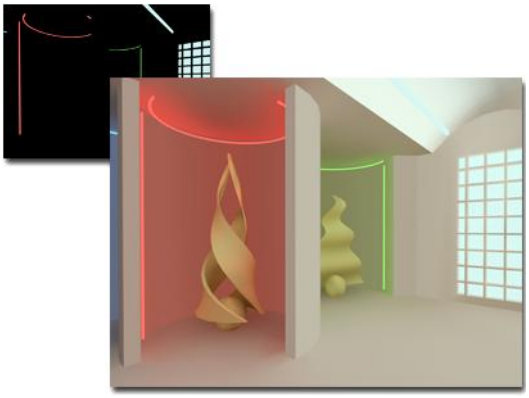


The room is lit only by spotlights pointed at the floor. Increasing reflectance of the floor brightens the entire room.

Tip: Check the *reflectance and transmittance display (page 2–1430)* to get an idea of how the current material will affect the radiosity solution or light-traced rendering.

Creating Special Effects

Self-illumination (page 2–1487) makes an object appear to glow in ordinary renderings, but does not contribute energy to the radiosity solution. To have radiosity processing take a self-illuminating material into account, make this material the base material of Advanced Lighting Override, then increase the value of Luminance Scale.



Upper left: By default, self-illuminated neon lights do not influence the scene light.

Right: Advanced Lighting Override material scales the neon lights' Luminance so the radiosity solution can take it into account.

Luminance scale takes *self-illumination mapping (page 2–1502)* into account. You can use this to model effects such as a computer monitor in a darkened room.

The Special Effects group of the Advanced Lighting Override material also has a control for adjusting the quality of *bump mapping (page 2–1506)* in areas of indirect lighting.

Procedures

To adjust a material's reflectance and transmittance:

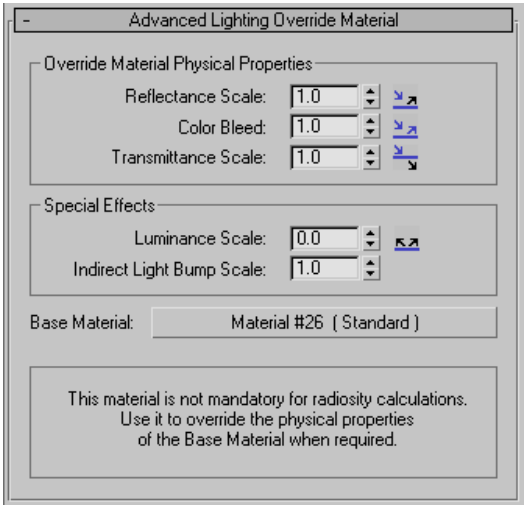
1. Create a material for your scene.
2. Click the Type button and choose Advanced Lighting Override.
3. In the *Replace Material dialog (page 2–1456)*, choose Keep Old Material As Sub-Material, and click OK.
4. Adjust the Reflectance Scale and Transmittance Scale parameters. As you do, watch the Reflectance and Transmittance display, and make sure the values are good for a radiosity solution. For example, 85 percent reflectance is about the highest that will work with radiosity. Most real-world materials have much lower reflectance.

   See *Reflectance and Transmittance Display (page 2–1430)* for some reflectance properties of real-world materials.

To make a self-illuminating material emit radiosity energy:

1. Create a material that is self-illuminating.
2. Click the Type button and choose Advanced Lighting Override.
3. In the *Replace Material dialog (page 2–1456)*, choose Keep Old Material As Sub-Material, and click OK.
4. Increase the value of Luminance Scale to have the material emit energy for radiosity processing.

Interface

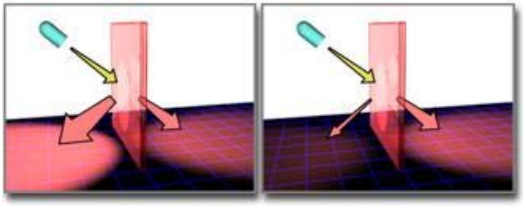


Override Material Physical Properties group

These parameters directly control the base material's advanced lighting properties.

Warning: There is no problem with reducing the default scale, but increasing it for any of these parameters might cause colors to "burn out": if the scale is too great, they render as pure white, appearing overexposed.

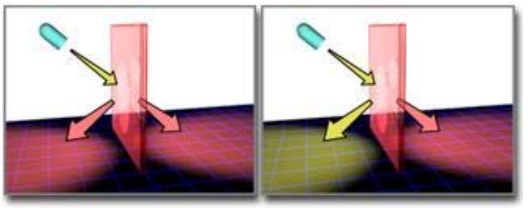
Reflectance Scale—Increases or decreases the amount of energy the material reflects. Range=0.1 to 5.0. Default=1.0.



Reflectance Scale increases or decreases the energy of reflected rays.

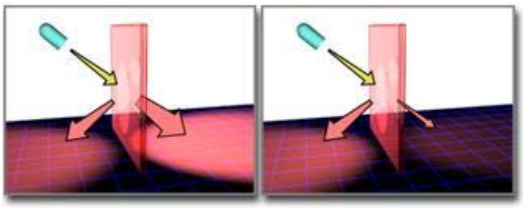
Tip: Don't use this control to increase self-illumination. Use the Luminance Scale instead. Luminance Scale is in the Special Effects group (below).

Color Bleed—Increases or decreases the saturation of reflected color. Range=0.0 to 1.0. Default=1.0.



Color Bleed increases or decreases the saturation of reflected color.

Transmittance Scale—Increases or decreases the amount of energy the material transmits. Range=0.1 to 5.0. Default=1.0.



Transmittance Scale increases or decreases the energy of transmitted rays.

Note: This parameter affects only radiosity. It has no effect on light tracing.

Special Effects group

These parameters relate to specific components in the base material.

Luminance Scale (cd/m^2)—When greater than zero, scales the *self-illumination component (page 2–1487)* of the base material. Use this parameter to have self-illuminating objects contribute energy to the radiosity or light-traced solution. Cannot be less than zero. Default=0.0.

Typically, a value of 500 or more will give good results.

Indirect Light Bump Scale—Scales the effect of the base material's *bump mapping (page 2–1506)* in areas lit by indirect light. When this value is

zero, no bump mapping is done for indirect light. Increasing Indirect Light Bump Scale increases the bump effect under indirect lighting. This value does not affect the Bump amount in areas where the base material is lit directly. Cannot be less than zero. Default=1.0.

Tip: This parameter is useful because indirect bump mapping is simulated and not always accurate. Indirect Light Bump Scale lets you adjust the effect by hand.

Base Material—Click to go to the base material and adjust its components. You can also replace the base material with a different material type.


To return from the base material to the Advanced Lighting Override level, click Go To Parent.

Lightscape Material

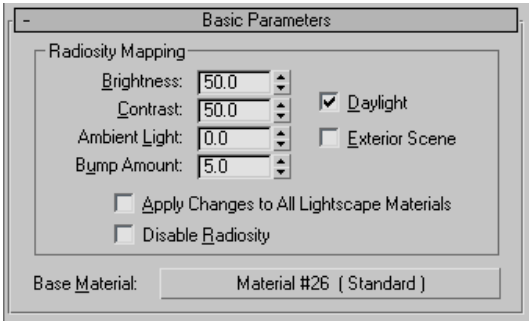
Material Editor > Type button > Material/Map Browser > Lightscape Mtl

The Lightscape material lets you set radiosity behavior for 3ds Max materials you want to use in existing Lightscape radiosity meshes.

Note: This material is for use with Lightscape. It is *not* meant for use with the advanced lighting solution in 3ds Max. For adjusting a material's radiosity properties in 3ds Max, use the *Advanced Lighting Override material (page 2–1601)*.

Important: The *mental ray renderer (page 3–78)* does not support the Lightscape material.

Interface



Radiosity Mapping group

Brightness—Controls the brightness of the displayed image on your monitor. The setting of this control does not affect the actual lighting levels in the model. Note: The default value is exactly the same as the values set in the Import Lightscape Solution dialog.

Contrast—Controls the contrast between light and dark regions in the model. Note: The default value is exactly the same as the values set in the Import Lightscape Solution dialog.

Ambient Light—Controls the amount of 3ds Max ambient light that will be mixed in with the radiosity calculations. If the value is 0, none of the 3ds Max ambient light is used. If the value is 1, the 3ds Max ambient light value is added into the radiosity calculations. Default=0.0.

Bump Amount—This value controls the strength of the bump map that is applied to the lighting from Lightscape, as opposed to lighting done by 3ds Max. It can be separately controlled from the bump amount in the 3ds Max material, so you can adjust the bumps in Lightscape lighting to match the bumps in 3ds Max lighting. Default=5.0.

Daylight—Determines whether you want natural daylight to be used in the calculation. Note: The default value is exactly the same as the values set in the Import Lightscape Solution dialog.

Exterior Scene—Used for exterior daylight simulations. Note: The default value is exactly the same as the values set in the Import Lightscape Solution dialog.

Apply Changes to All Lightscape Materials—Determines whether or not the changes you make in this dialog are applied to every Lightscape Radiosity material. Default=off.

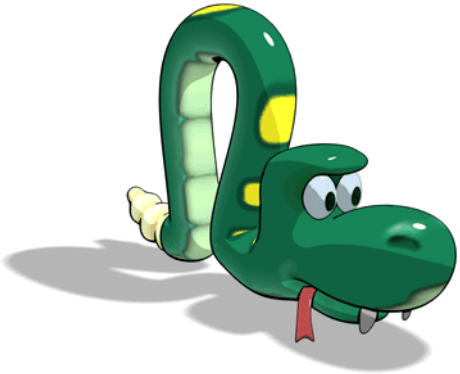
Disable Radiosity—Determines whether or not the radiosity is calculated. This option provides a quick way to see the effect of the Lightscape-specific illumination. Default-off.

Base Material—Displays the base 3ds Max material to which the radiosity illumination is applied. (The base material displays on the object while the Lightscape Radiosity material adds radiosity effects.) Use this button to change the base material on the object.

Ink 'n Paint Material

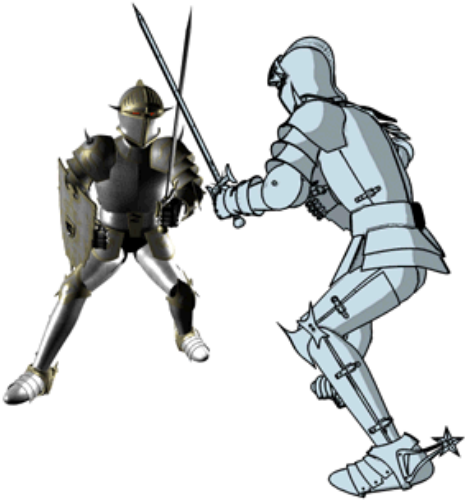
Material Editor > Type button > Material/Map Browser > Ink 'n Paint

The Ink 'n Paint material creates cartoon effects. Rather than the three-dimensional, realistic effect most other materials provide, Ink 'n Paint provides flat shading with "inked" borders.



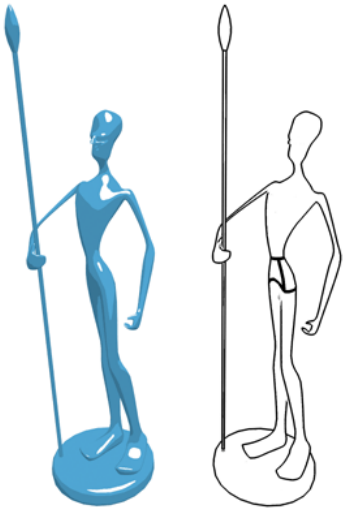
Snake rendered with ink 'n paint

Because Ink 'n Paint is a material, you can create a scene that combines 3D-shaded objects with flat-shaded cartoon objects.



Rendering that combines realistic shading with cartoon shading

In the Ink 'n Paint material, ink and paint are two separate components, with customizable settings.



Left: The paint component only
Right: The ink component only

Tip: Ink 'n Paint uses the *Raytracer Settings (page 2–1528)*, so adjusting raytrace acceleration can have an effect on the speed of Ink 'n Paint. Also, while you work with Ink 'n Paint, disabling antialiasing can speed up the material, until you're ready to create final renderings. (Turning off Ink *really* speeds it up.)

Note: Motion blur does not work with Ink 'n Paint. (Typically, hand-drawn cartoons are not motion blurred.)

Note: Shadows don't appear on objects shaded with Ink 'n Paint unless the value of Paint Levels is 4 or greater.

Warning: Ink 'n paint will only give correct results when rendered from a camera or perspective view. It does not work in orthographic views.

Using Ink 'n Paint

You can use Ink 'n Paint on multiple objects, but in general, it tends to work best if you do the following:

1. Collect the objects for cartoon rendering into a single surface model such as an Editable Mesh.
2. Assign different *material ID values (page 3–969)* to portions of the model you want to color differently.

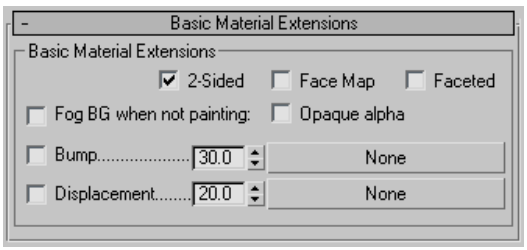
   Typically, you would do this at the Element sub-object level, although you can certainly apply different material IDs to faces and polygons as well.
3. Create a *Multi/Sub-Object material (page 2–1594)*. In it, create a sub-material for each of the colors in the model. Make each sub-material an Ink 'n Paint material, then assign colors and maps using each sub-material's Paint controls.

   If necessary, adjust the Ink controls as well.

Tip: *ActiveShade (page 3–17)* works with the Ink 'n Paint material, and can be a good way to preview the material's effect.

Interface

Basic Material Extensions rollout



2-Sided—Makes the material *2-sided (page 3–901)*. Applies the material to both sides of selected faces.

Face Map—Applies the material to the faces of the geometry. If the material is a mapped material, it requires no *mapping coordinates (page 3–967)*. The map is automatically applied to each facet of the object.

Faceted—Renders each face of a surface as if it were flat.

Fog BG when not painting—When paint is turned off, the painted areas of the material color are the same as the background. This toggle, when on, lets the background in paint areas be affected by fog between the camera and the object. Default=off.

Opaque alpha—When on, the alpha channel is opaque even if ink or paint is turned off. Default=off.

Bump—Adds bump mapping to the material.

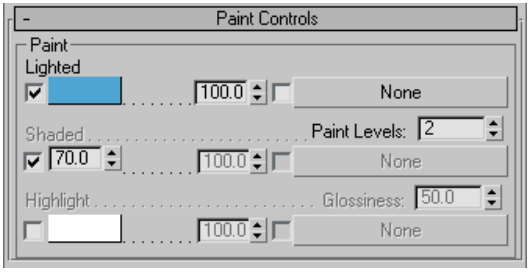
- **Toggle**—When on, enables the bump map.
- **Spinner**—Controls the bump map amount.
- **Map button**—Click to assign a map to use for bump mapping.

Displacement—Adds displacement mapping to the material.

- **Toggle**—When on, enables the displacement map.
- **Spinner**—Controls the displacement map amount.
- **Map button**—Click to assign a map to use for displacement mapping.

Paint Controls rollout

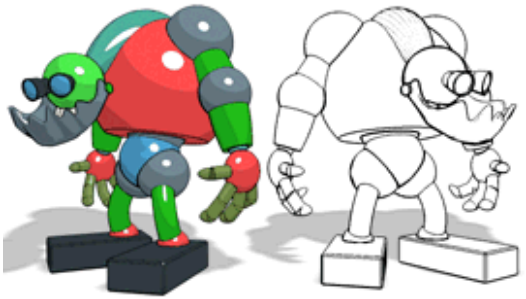
Paint is the main color of the material.



There are three main components of the "paint" of Ink 'n Paint. Each has several associated controls, most of which are documented toward the end of this section.

Lighted—The fill color for the lighted side of objects. Default=light blue.

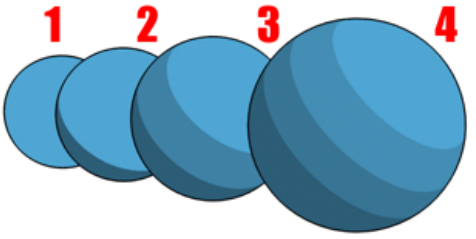
Turning off this component makes the object invisible, except for the ink. Default=on.



Left: A lighted character

Right: Lighted and Highlight both turned off to render only the ink

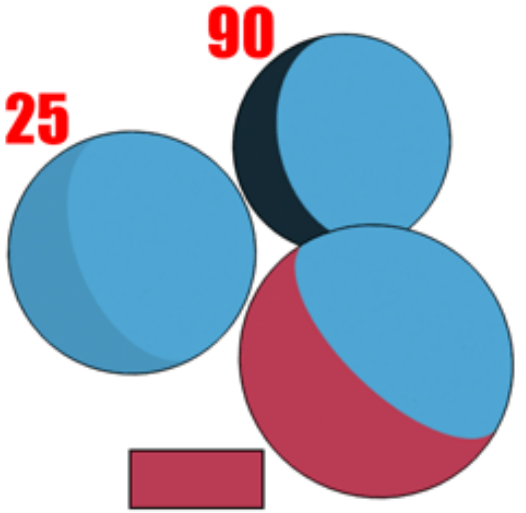
- **Paint Levels**—The number of shades of color that are rendered, from light to dark. Lower values make objects look flatter. Range=1 to 255. Default=2.



Increasing the value of Levels increases the number of shades of the basic color seen in the lighted area.

Shaded—The value in the spinner at the left is the percent of the Lighted color that appears on the unlighted side of objects. Default=70.0.

Turning off this component displays a color swatch, which you can use to assign a distinct color to shaded areas. Default=on.



Increasing the value of Shaded increases the saturation of the shaded area. You can also use Shaded to assign a distinct color for shading.

Highlight—The color of the specular highlight. Default=white.

When this component is off, there is no specular highlight. Default=off.

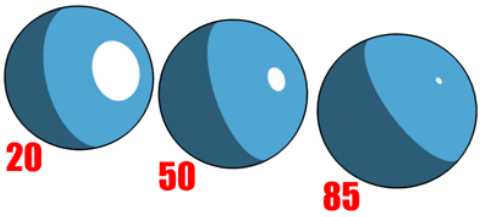
Tip: A specular highlight can destroy the illusion of 2D. Use this component sparingly.



Left: No highlights

Right: Highlight on

- **Glossiness—**The size of the specular highlight. The greater the Glossiness, the smaller the highlight. Default=50.0.



Increasing glossiness decreases the size of the highlight.

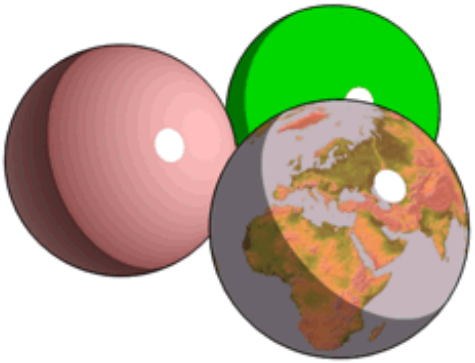
Color component controls:

These are the controls that are duplicated for each of the paint components. Each has an on/off toggle, a main control, and then on the right, a set of map controls.

- **Check box—**The check box at the left of the rollout enables or disables that particular component. In the case of Shaded, it toggles between a percentage value (of the Lighted color) or a distinct Shaded color.
- **Color swatch or spinner—**The main control for each component. Click a color swatch to display a *Color Selector (page 1–161)* and set the color of the component. In the case of Shaded, this control can also be a percentage spinner.



- **Map spinner—**The spinner to the right of the main control is the percentage of the map to use. Default=100.0.



Mapping the Lighted component

Right rear: The original, unmapped material

Left: Lighted component with a falloff map applied

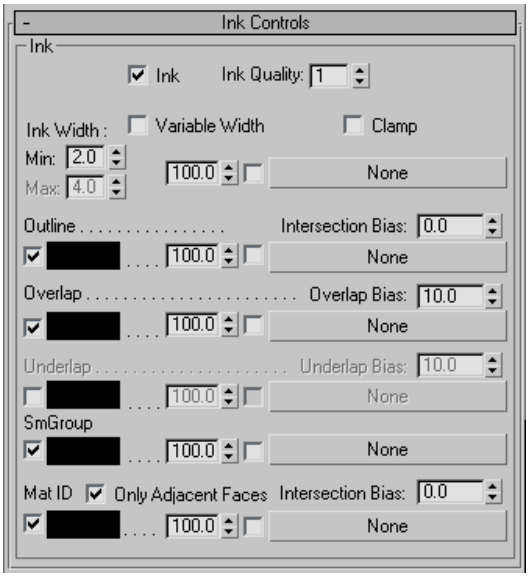
Right front: Lighted component with a bitmap applied

- **Map check box—**The check box between the spinner and the button enables or disables the map. Default=off until a map is assigned, then on.
- **Map button—**Click the button to assign a map to this component.

  While a map is assigned and enabled, at 100 percent it completely overrides the main color component. At lower percentages, the map is blended with the color.

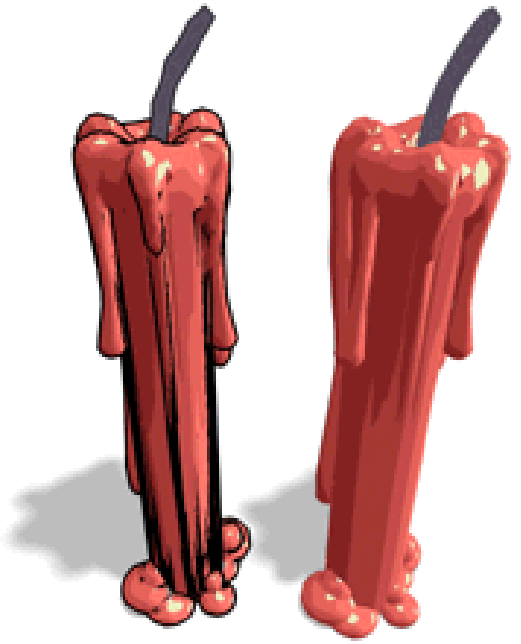
Ink Controls rollout

Ink is the linework, the outlines, in the material.



Except for Ink Width, each of the ink components has an on/off toggle and a color swatch. Click the color swatch to display a *Color Selector (page 1–161)* and change the ink component's color. Each ink component, Ink Width included, also has a set of map controls.

Ink—When on, the rendering is "inked." When off, no ink lines appear. Default=on.



Left: Rendering with ink
Right: Ink turned off

Ink Quality—Affects the shape of the brush and the number of samples it uses. When Quality equals 1, the brush is a "+" shape, and samples are taken over an area of 5 pixels. When Quality equals 2, the brush is octagonal and the samples are taken over an area of 9 to 15 pixels. When Quality equals 3, the brush is nearly circular, and samples are taken over an area of 30 pixels. Range=1 to 3. Default=1.

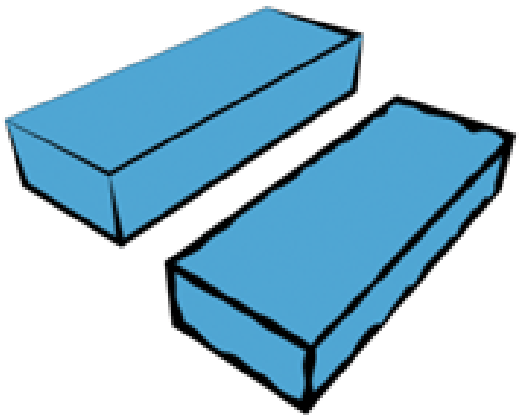
Tip: For most models, increasing the Quality value introduces only a very subtle change, and can take considerably longer to render. Do so only when a sub-object's ink shows too many artifacts in the finished rendering, using the default Ink Quality. (Don't rely on the ActiveShade preview, which will tend to be aliased.)

Ink Width—The width of the ink, in pixels. This is specified by the spinner labeled Min (minimum) unless Variable Width is on. When Variable Width is on, the Max (maximum) spinner is also enabled, and the ink width can vary between the minimum and maximum values. Default: Min=2.0, Max=4.0.



Left: One-pixel ink width
Middle: Five-pixel ink width
Right: Ink width varies from one to five pixels.

Variable Width—When on, the ink's width can vary between the minimum and maximum Ink Width values. Ink with Variable Width looks a bit more streamlined than ink with a constant width. Default=off.

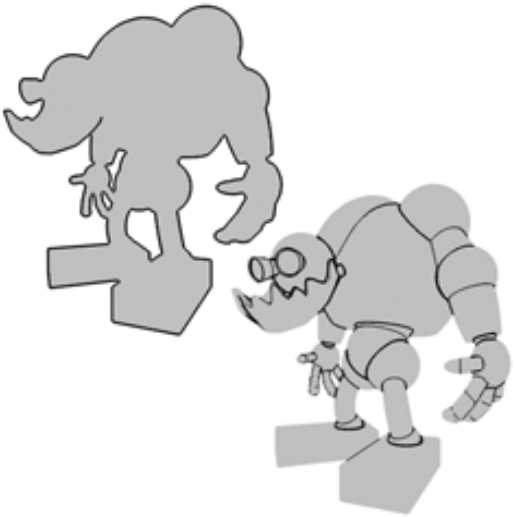


The thickness of ink can be mapped.
Left: Thickness mapped with a gradient map
Right: Thickness mapped with a noise map

Clamp—When Variable Width is on, sometimes the scene lighting causes some ink lines to become so thin they nearly disappear. If this happens, turn on Clamp, which forces the ink width to

always remain between the Min and Max values, regardless of the lighting. Default=off.

Outline—The ink where the outer edges of the object appear against the background or in front of a different object. Default=on.



Left: Rendering the outline only

Right: Rendering only the overlap and underlap

- **Intersection Bias—**Use this to adjust artifacts that might appear when two objects intersect each other. In effect, this moves the inked object closer to the rendering point of view, or farther away, so Ink 'n Paint can decide which object is in front. Positive values push the object away from the point of view, negative values pull it closer. Default=0.0.

Overlap—The ink used when a portion of an object overlaps itself. Default=on.

- **Overlap Bias—**Use this to adjust artifacts that might appear in ink that traces the overlap. It says how far the overlap has to be in front of the rear surface for Overlap ink to turn on. Positive values push the object away from the point of view, negative values pull it closer. Default=10.0.

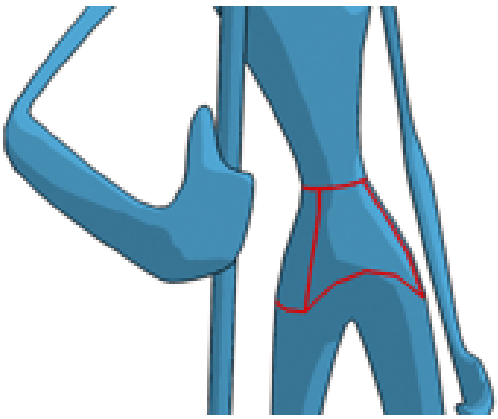
Underlap—Similar to Overlap, but applies ink to the farther surface rather than the nearer one. Default=off.

- **Underlap Bias—**Use this to adjust artifacts that might appear in ink that traces the underlap. It says how far the underlap has to be behind the front surface for Underlap ink to turn on. Positive values push the object away from the point of view, negative values pull it closer. Default=0.0.

SmGroup—The ink drawn between the boundaries of *smoothing groups (page 3–1013)*. In other words, it inks the edges of the object that have not been smoothed. Default=on.

Mat ID—The ink drawn between different *material ID values (page 3–969)*. Default=on.

Tip: If two Ink 'n Paint materials overlap in the viewport, and both have Mat ID on, you will often get a doubly thick ink line where they overlap. To correct this, turn off the Mat ID component for one of these materials.



Inking the edges between sub-materials

- **Only Adjacent Faces—**When on, inks the material ID edge between adjacent faces, but not between one object and another. When off, inks the material ID edge between two objects or other non-adjacent faces. Default=on.

- **Intersection Bias—**When Only Adjacent Faces is turned off, use this to adjust any artifacts that appear at the boundary between two objects with different material IDs. Default=0.0.



Map controls—There are map controls for each of the ink components: Width, Outline, Overlap, Underlap, SmGroup, and Mat ID. These work the same as they do for the material's paint components, as described above.



Mapping the outline and overlap components to simulate the look of drawing on paper

Troubleshooting

Here are some commonly encountered problems, and potential solutions:

- Internal ink lines are missing.

  The Overlap bias is probably too high. Decrease it. If Underlap is turned on, this might also have too high a bias.

  Another possible reason is that you have a self-intersecting object, or an object built by attaching smaller objects, thus creating intersecting faces. In this case, set up the objects to use the Mat ID or SmGroup ink components. If elements already have differing material IDs, try turning off Only Adjacent Faces.

- Ink looks sloppy on sloping parts of the object.

  The Overlap or Underlap bias might be too low. Try increasing it.

- Ink looks sloppy between interpenetrating objects.

  Find out which ink component is the sloppy one. Then adjust its bias control.

- Ink lines disappear or are too narrow when Variable Width is on.

  Turn on Clamp. You can also try to see if reducing the lighting level helps. Or, you can try turning off Variable Width, then assigning a *Falloff map (page 2–1670)* to the Ink Width component.

Tip: To isolate which ink component is causing a problem, you can try assigning each component a different, distinctive (and easy to read) color, then rendering the image.

Warning: Ink 'n paint will only give correct results when rendered from a camera or perspective view. It does not work in orthographic views.

DirectX 9 Shader Material

Material Editor > Type button > Material/Map Browser > DirectX 9 Shader

Note: Typically, this material is visible in the Browser only if DirectX 9 is available on your system, and you are using the Direct3D display driver with DirectX 9.0 chosen as the Direct3D version. If this material is not visible, you can see it (in gray) by turning on Incompatible in the Show group.

The DirectX 9 Shader material enables you to shade objects in viewports using DirectX 9 (DX9) shaders. With DirectX shading, materials in a viewport more accurately represent how the material will appear in another application, or on other hardware such as a game engine. You can use this material only when you are using the *Direct3D Display driver (page 3–844)* and DirectX 9.0 is chosen as the Direct3D version.

DX9 shaders are FX files. Several sample FX files are provided in the *\3ds Max 9\maps\fx* directory (the name of the 3ds Max root directory might differ on your system).

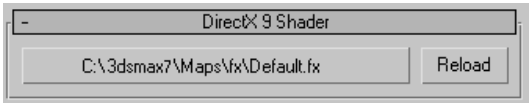
See also

DirectX Shader group (page 2–1438)

FX File (page 3–946)

Interface

DirectX 9 Shader rollout



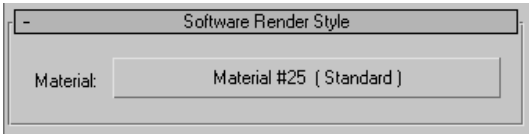
Shader button—Click to display a file dialog that lets you select a DX9 FX effect (FX) file. By default, the *default.fx* file is chosen.

Reload—Click to reload the active FX file. To update an FX file, you can edit it and then click Reload. You don't have to restart 3ds Max to see the effect of the changes to the shader.

Shader-specific rollouts

The rollouts that appear below the DirectX 9 Shader rollout and above the Software Rendering rollout are the interface to the shader you chose. These rollouts are specific to each shader.

Software Rendering rollout



Specifies a material that controls software shading and rendering of objects to which the DirectX 9 Shader material is applied. Viewports use DX9 shading unless the Software or OpenGL driver is active. Renderings *always* use software shading.

Usually you will want to choose a material that clearly identifies which objects in your scene have the DirectX 9 Shader material applied.

Note: The DirectX 9 Shader material has no specific settings for software shading. Any type of 3ds Max material will do. Scenes from previous versions that used DX9-specific settings are assigned a Standard material with equivalent rendering properties.

Note: If DirectX 9 is not available on your system, but you assign the DirectX 9 Shader material anyway (by using the Material/Map Browser's Incompatible option), this is the only rollout that appears in the Material Editor.

DirectX Viewport Shaders

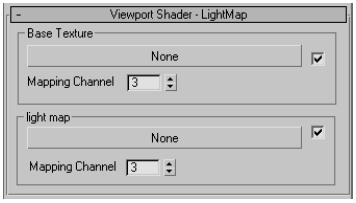
LightMap Shader Rollout

Material Editor > Viewport Manager rollout > Choose LightMap from the drop-down list. > LightMap Shader rollout appears.

When you have chosen LightMap as the *DirectX viewport shader (page 2–1464)*, this rollout appears. The LightMap shader can display both a base texture and a lighting map. Typically both these maps come from *rendering to textures (texture baking) (page 3–144)* The base texture typically would be a completed map, a blend map, or a diffuse map. You can choose these map types, as well as lighting map, to render as *elements of a baked texture (page 3–146)*.

Note: In order to use the LightMap shader, you must have 3ds Max configured to use the Direct3D graphics driver. To change the graphics driver configuration, refer to the *Viewport Preferences (page 3–821)* topic.

Interface



Base Texture group

Button—Shows the name of the base texture. Click the button to display that material's parameters, and adjust them if necessary.

Toggle—When on, shaded viewports display the base texture. When off, it is not displayed.

If both the Base Texture and Light Map toggles are off, the material appears black in viewports.

Mapping Channel—Shows the *map channel (page 3–966)* this texture uses.

Light Map group

Button—Shows the name of the lighting map.

Toggle—When on, shaded viewports display the lighting map. When off, it is not displayed.

If both the Base Texture and Light Map toggles are off, the material appears black in viewports.

Mapping Channel—Shows the map channel this texture uses.

Metal Bump Shader Rollout

Material Editor > Viewport Manager rollout > Choose LightMap from the drop-down list. > Metal Bump Shader rollout appears.

When you have chosen Metal Bump as the *DirectX viewport shader (page 2–1464)*, this rollout appears. The Metal Bump shader can display a variety of *texture-baked maps (page 3–144)*, including normal maps for an embossed effect. It is good for displaying shiny surfaces.

Note: In order to use the Metal Bump shader, you must have 3ds Max configured to use the Direct3D graphics driver. To change the graphics driver configuration, refer to the *Viewport Preferences (page 3–821)* topic.

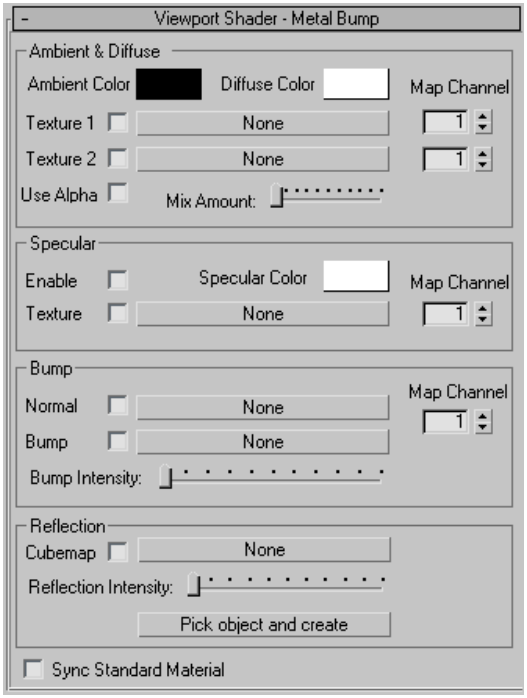
The Metal Bump shader's results are alway visible in viewports, regardless of the object type.

Warning: The Metal Bump shader lets you adjust settings to get various effects in shaded viewports. These settings will not necessarily apply when you display the texture-baked object on other Direct3D devices.

See also

Baked Texture Elements (page 3–146)

Interface



Ambient & Diffuse group

Ambient Color—When not black, tints the object's ambient color. Click the color swatch to display a *Color Selector (page 1–161)* and choose the ambient color. Default=black.

Diffuse Color—When not white, tints the diffuse color. Click the color swatch to display a Color Selector and choose the diffuse color. Default=white.

Texture 1—Displays a texture map for the diffuse color. Typically this would be a texture-baked diffuse map, completed map, or blend map.

See the section "Map Controls," below, for a description of the individual controls.

Texture 2—Displays a second texture map for the diffuse color. Typically this would be a texture-baked lighting map or shadows map.

See the section "Map Controls," below, for a description of the individual controls.

Use Alpha—When on, displays the alpha channel. When off, does not. Default=off.

Mix Amount—Adjusts the mixing of the two texture maps in shaded viewports.

Specular group

Enable—When on, enables specular highlights for the object. Default=off.

Specular Color—Specifies a specular color for the object. Click the color swatch to display a Color Selector and choose a color. Default=white.

Texture—Displays a specular map for the object. Typically this would be a texture-baked specular map.

See the section "Map Controls," below, for a description of the individual controls.

Bump group

Normal—Displays a normal map for the object. Typically this would be a texture-baked normals map.

See the section "Map Controls," below, for a description of the individual controls.

Bump—Displays a bump map for the object. Typically this would be a bump map used for the original material.

See the section "Map Controls," below, for a description of the individual controls.

Bump Intensity—Adjusts the intensity of the bumps in shaded viewports.

Reflection group

Cubemap—Displays a reflection map projected cubically (around the scene). Typically this would be an environment map.

Reflection Intensity—Adjusts the intensity of reflections in shaded viewports.

Pick object and create—Click to choose an object and have the program generate the reflections used in the viewport.

Sync Standard Material

When on, adjustments you make to the Metal Bump shader update settings in the active standard material, letting you save the changes you made. When off, the standard material is unchanged. Default=off.

Map Controls

In this rollout, all texture maps have the same general controls. The rollout appears only when you are using the DirectX viewport shader.

Toggle—When on, the map is used in viewports. When off, it is not used. Default=on if a map is assigned, off otherwise.

Map button—Click to choose the texture map to use for this component of the object.

Map Channel—Specifies the *map channel (page 3–966)* used by this map.

This control is not present for the cubic reflection map.

XRef Material

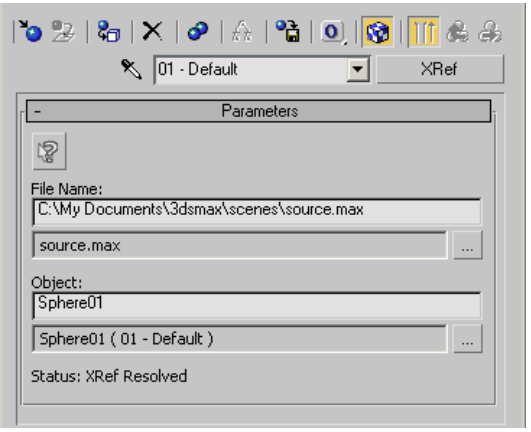
Material Editor > Type button > Material/Map Browser > XRef Material

The XRef material lets you externally reference a material applied to an object in another scene file. As with *XRef objects (page 3–394)*, the material resides in a separate source file. You can set the material properties only in the source file. When you change them in the source file and then save it, the material's appearance can change in the master file that contains the XRef.

Note: If an XRef object has a material applied to it in the original source file (and Merge Materials is turned off when you reference the source file), then that material is automatically externally referenced in the scene, and can be loaded in the Material Editor if you browse from the scene. When you explicitly use the Material Editor to create an XRef material, you don't have to have any XRef objects from that particular source file. However, the record's source file and material *do* appear in the *XRef Objects dialog (page 3–397)*.

The Show Map In Viewport button works for an XRef material only if the same button is turned on in the source file. Otherwise, it is disabled.

Interface



Highlight Corresponding XRef Record in the XRef Objects Dialog—Click to open the *XRef Objects dialog (page 3–397)* that highlights the source file's current record with its object displayed in the XRef Entities list. If no file and object have yet been selected for the material, the XRef Objects dialog is displayed, and lets you browse for the file and material to use.

File name field—Displays the path and file name of the scene file containing the source of the XRef material. You can edit this to point to a different path and file.

File name display—Displays the file name only, without the path.

 **Path button—**Click to display the Open File dialog from which you can specify a different path and file name for the source file. After you choose the file, 3ds Max displays the *XRef Merge dialog (page 3–406)* that lets you choose the object whose material you want to reference.

Object name field—Displays the name of the source object pointed to in the source file.

Object name and material—Displays the name of the source object followed by the material name in parentheses. For example, "Shaker (Chrome)."

 **Path button—**Click to display the *XRef Merge dialog (page 3–406)* pointing to the scene in the XRef File Name field. Here, you can specifiy a different object whose material you want to reference.

Status line—Displays the status of the material. For example, if the file and object are both found, this field says "Status: XRef Resolved."

Maps

The most common use for *maps (page 3–968)* is to improve the appearance and realism of *Materials (page 3–971)*. You can also use maps to create *environments (page 3–272)* or projections from lights (see *Advanced Effects Rollout (page 2–1341)*).

Maps can simulate textures, applied designs, reflections, refractions, and other effects. Used with materials, maps add details without adding complexity to the geometry of an object. (*Displacement mapping (page 2–1511)* can add complexity.)

Accessing Map Types

You use the *Material/Map Browser (page 2–1412)* to load a map or create a map of a particular type. The Browser groups maps into categories according to their type. You can choose whether the Browser lists maps, materials, or both; you can choose which map types.

To list map types by category:

1.  Open the Material Editor.
2.  Click Get Material to display the Material/Map Browser.
3. On the left side of the Material/Map Browser, in the Show group, turn off Materials. Leave Maps turned on.
4. In the lower set of buttons, choose the category you want listed, or choose All to show all map types.

   The list is displayed in the right panel of the Browser.

To open the Browser from the Material Editor:

As you work with materials, you can open the Material/Map Browser from the Material Editor.

1.  Open the Material Editor.
2.  On the Material Editor toolbar, click Get Material to display the Material/Map Browser.

Understanding Map Types

Different types of maps create different effects and behave in particular ways.

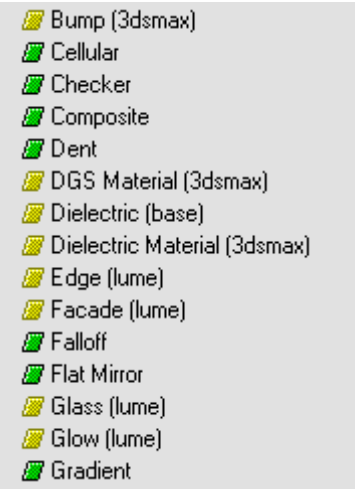
- *2D maps (page 2–1624)* are two-dimensional images that are typically mapped onto the surface of geometric objects, or used as environment maps to create a background for the scene. The simplest 2D maps are

bitmaps; other kinds of 2D maps are generated procedurally.

- *3D maps (page 2–1662)* are patterns generated procedurally in three dimensions. For example, Marble has a grain that goes through the assigned geometry. If you cut away part of an object with Marble assigned as its texture, the interior matches the exterior.
- *Compositors (page 2–1687)* are for compositing other colors or maps.
- *Color Modifiers (page 2–1692)* alter the color of pixels in a material.
- The "Other" category includes maps that create *reflections and refractions (page 2–1695)*. It also includes *mental ray shaders (page 2–1710)* , *Normal Bump maps, (page 2–1731)* and *Camera Per-Pixel Projection maps (page 2–1732)*.

Maps and mental ray Shaders

When the *mental ray renderer (page 3–78)* is active, the Material/Map Browser also lists mental ray shaders. Shaders are similar to maps, but use yellow icons. You assign them the way you do maps. The mental ray shaders don't fit into the map categories described in the previous section, and aren't described in this topic. See *mental ray Shaders (page 2–1710)* for links to shader descriptions.



mental ray maps in the browser's list are shown with yellow icons.

Tip: When the default scanline renderer is active, you can view mental ray shaders in the Browser list, and assign them, by turning on Incompatible in the Show group. Incompatible shaders in the list are displayed in gray.

Maps and Mapping Coordinates

Maps have a spatial orientation. When you apply a material with maps in it to an object, the object must have mapping coordinates. These are specified in terms of UVW axes local to the object.

Most objects have a Generate Mapping Coordinates toggle. You can turn this on to provide default mapping coordinates. If the object has this toggle, it is also turned on automatically when you render the scene, or use *Show Map In Viewport (page 2–1445)*.

Some objects, such as editable meshes, don't have automatic mapping coordinates. For these types of objects, you can assign coordinates by applying a *UVW Map modifier (page 1–922)*. If you assign a map that uses a mapping channel, but don't apply a UVW Map modifier to the object, the renderer displays a warning that lists objects that require

mapping coordinates. You can also use UVW Map to change an object's default mapping.

See *Mapping Coordinates (page 2–1405)*.

UVW Mapping Coordinate Channels

Each object can have from 1 to 99 UVWmapping coordinate channels. The default mapping (from the Generate Mapping Coordinates toggle) is always UVW 1. The UVW Map modifier can send coordinates to any of these channels.

Each map in a material can use any UVW channel (if present), or other type of mapping that depends on whether the map is 2D or 3D.

You can set the mapping channel used by *NURBS surface (page 1–1101)* sub-objects in their creation or modification parameters.

Mapping for 2D Maps

You can position a 2D map on the surface of an object by using a map channel, any assigned vertex color, or the local or world coordinate systems. You can also choose different environment mappings. See *Coordinates Rollout (2D) (page 2–1625)*.

Mapping for 3D Maps

You can position a 3D map within the volume of an object by using a map channel, any assigned vertex color, or the local or world coordinate systems. See *Coordinates Rollout (3D) (page 2–1663)*.

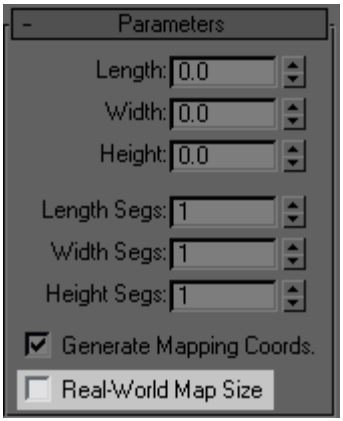
Noise for Maps

Random noise values increase the complexity of maps and can give them a more natural look. For 2D maps, see *Noise Rollout (2D) (page 2–1630)*. For 3D maps, you can assign a separate Noise map. See *Noise Map (page 2–1674)*.

Real-World Mapping

Real-world mapping is an alternative mapping paradigm in 3ds Max that is off by default. The idea behind real-world mapping is to simplify the correct scaling of texture-mapped materials applied to geometry in the scene. This feature lets you create a material and specify the actual width and height of a 2D texture map in the Material Editor. When you assign that material to an object in the scene, the texture map appears in the scene with correct scaling.

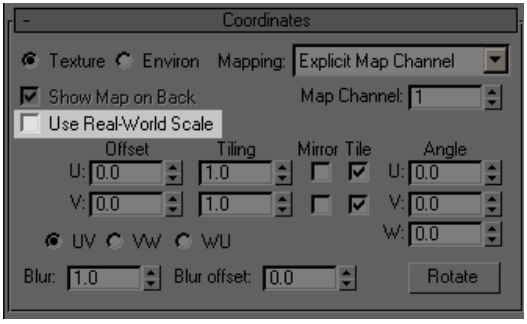
For real-world mapping to work, two requirements must be met. First, the correct style of UV texture coordinates must be assigned to the geometry. Essentially, the size of the UV space needs to correspond to the size of the geometry. To this end, a new option called Real-World Map Size has been added to many of the dialogs and rollouts that let you generate texture coordinates (see list at the end of this topic). Any dialog or rollout that lets you turn on Generate Mapping Coords also lets you enable Real-World Map Size. Also, you can toggle this option globally on the Preferences dialog > *General panel (page 3–815)*.



Most object parameters settings now include a Real-World Map Size toggle.

The second requirement is available in the Material Editor. All 2D texture maps, such as

Bitmap, provide a Use Real-World Scale check box on the *Coordinates rollout (page 2–1625)*. Like Real-World Map Size, this check box is off by default, but when on, the U/V parameter names change to Width/Height and the Tiling label changes to Size. You can then specify the horizontal/vertical offsets and size of the texture map in current *system units (page 3–850)*.



2D map coordinates settings now include a Use Real-World Scale toggle.

Following is a list of affected features (note links to topics with more extensive information):

- Bevel Modifier
- Bevel Profile Modifier
- Box Primitive
- Capsule Extended Primitive
- C-Ext Extended Primitive
- ChamferBox Extended Primitive Create panel > Geometry button
- ChamferCyl Extended Primitive
- Cone Primitive
- CV Curve
- Cylinder Primitive Create panel
- Doors
- Editable Spline
- Extrude Modifier
- Gengon Extended Primitive
- GeoSphere Primitive
- Importing AutoCAD Drawing
- Importing DXF Files
- Lathe Modifier
- L-Ext Extended Primitive
- L-Type Stair
- Loft > Surface Parameters Rollout
- Material Editor Coordinates Rollout
- *Material Editor Options (page 2–1436)*
- OilTank Extended Primitive
- Plane Primitive
- Point Curve
- Pyramid Primitive
- Railing
- Renderable Spline Modifier
- Sphere Primitive
- Spindle Extended Primitive
- Spline Rendering File Link Settings
- Splines and Extended Splines
- Stairs
- Sweep Modifier
- Teapot Primitive
- Torus Primitive
- Tube Primitive
- *UVW Map Modifier (page 1–922)*
- Wall
- Windows

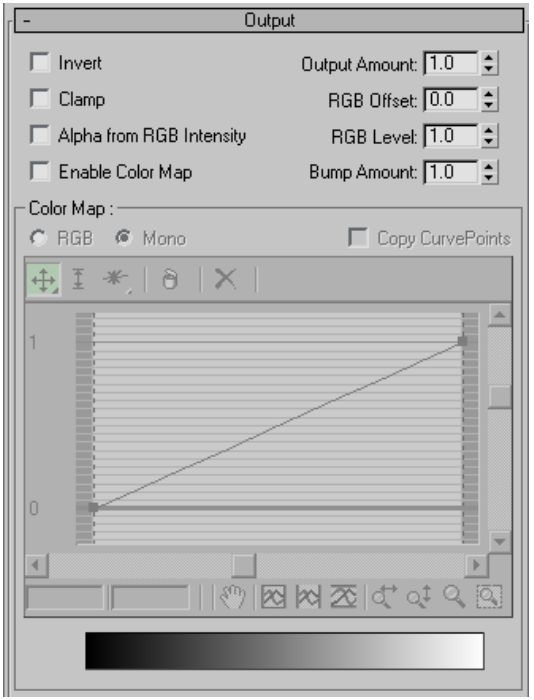
Output Rollout

Material Editor > Select sample slot. > Get Material > Material/Map Browser > Turn on 2D Maps and 3D Maps. > Choose a map type (Bitmap, Cellular, Falloff, Gradient, Gradient Ramp, Mix, Noise, or Output). > Double-click a map type to apply it to sample slot. > Output rollout (displayed near the bottom of the Material Editor)

After applying a map and setting its internal parameters, you can adjust its output parameters to determine the final appearance of the map.

Note: Most controls on the Output rollout are for color output, and don't affect *bump mapping (page 2–1506)*. Only the Invert toggle is considered; it reverses the direction of the bumps.

Interface



These controls appear on the Output rollout for a number of 2D and 3D maps:

Invert—Reverses the hues of the map, like a negative color photo. Default=off.

Clamp—When on, this parameter limits the values of the colors to no greater than 1.0. Turn this on when you're increasing the RGB Level, but don't want the map to appear self-illuminated. Default=off.

Note: If you set the RGB Offset to a value greater than 1.0 while Clamp is on, all colors become white.

Alpha from RGB Intensity—When on, an alpha channel is generated based on the intensity of the RGB channels in the map. Black becomes transparent and white becomes opaque. Intermediate values are translucent according to their intensity. Default=off.

Enable Color Map—Turn on to use the Color Map. See "Color Map group." Default=off.

Output Amount—Controls the amount of the map being mixed into a composite material. Affects the saturation and alpha value of the map. Default=1.0.

RGB Offset—Adds to the RGB values of the map colors by the amount set by the spinner, which affects the tonal value of the colors. Eventually the map becomes white and self-illuminated. Lowering the value decreases the tonal value toward black. Default=0.0.

RGB Level—Multiplies the RGB values of the map colors by the amount set by the spinner, which affects the saturation of the color. Eventually the map becomes fully saturated and *self-illuminated (page 3–1009)*. Lowering the value decreases the saturation and makes the map colors grayer. Default=1.0.

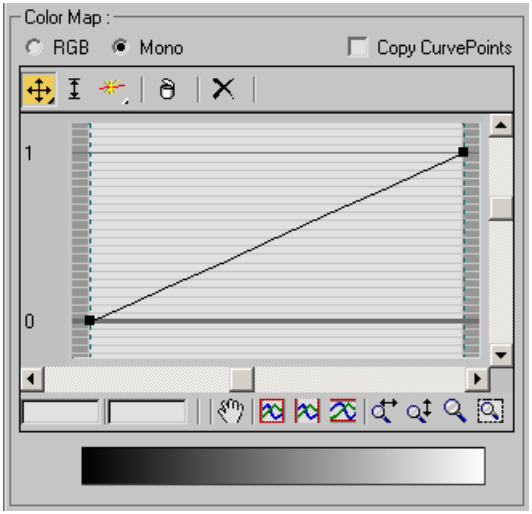
Bump Amount—Adjusts the amount of bumpiness. This value has an effect only when the map is used as a bump map. Default=1.0.

For example, suppose you have a map instanced for both the Diffuse and the Bump components. If you want to adjust the amount of bumpiness

without affecting the Diffuse colors, adjust this value, which changes the amount of bumpiness without affecting the map's use in other material components.

Color Map group

At the bottom of the Output rollout, the Color Map group is inactive until you turn on Enable Color Map.



The Color Map graph lets you adjust the tonal range of an image. The point at 1,1 controls highlights, the point at 0.5,0.5 controls midtones, and the point at 0,0 controls shadows.

You adjust the shape of the graph by adding points to the line and moving or scaling them. You can add Corner, Bezier-Smooth, or Bezier-Corner points. When a move or scale option is active, points can be selected much like objects in a viewport, by clicking a point, dragging a region around one or more points, and holding down Ctrl to add or subtract from the selection.

When you select an individual point, its exact coordinates are displayed in the two fields below the graph at the lower left. You can enter values directly in these fields, but the values are automatically constrained as they are when you manually move or scale a point.

You can zoom into the graph to make detailed adjustments. As you zoom in, the graph updates to show decimal measurements along the left vertical axis. You can pan anywhere on the graph with the horizontal and vertical scroll bars, use a button option, or the middle mouse button. Points can be deleted, and you can reset the graph to its default at any time.

RGB/Mono—Specifies a map curve to either filter the RGB channels separately (RGB) or in combination (Mono).

Copy CurvePoints—When turned on, points added to a Mono graph are copied when you switch to an RGB graph. If you start with an RGB graph, the points are copied to a Mono graph. You can animate the control points but not the Bezier handles.

Important: When Copy CurvePoints is on, animation created in Mono mode is carried over to RGB mode and you can switch channels. The reverse doesn't work.

The following controls affect the points on the graph:

Move flyout

 Moves a selected point in any direction, limited by the unselected points on either side.

 Constrains movement to the horizontal.

 Constrains movement to the vertical.

On a Bezier smooth point, you can move the point or either handle.

 **Scale Point—**Changes the output amount of control points while maintaining their relative position. On a Bezier corner point, this control is effectively the same as a vertical move. On a

Bezier smooth point, you can scale the point itself or either handle. As with the move controls, scale is limited by the unselected points on either side.

Add Point flyout

 Adds a Bezier corner point anywhere on the graph line. The point makes a sharp angle when moved.

 Adds a Bezier smooth point anywhere on the graph line. Handles attached to the point create smooth curves when moved.

When either Add Point button is active, you can use Ctrl+click to create the other type of point. This eliminates the need to switch between buttons.

 **Delete Point**—Removes selected points.

 **Reset Curves**—Returns graph to its default, a straight line between 0,0 and 1,1.

The following controls affect the view of the graph. The change in view does not affect the graph's results.

 **Pan**—Drags the graph in any direction within the viewing window.

 **Zoom Extents**—Shows the entire graph.

 **Zoom Horizontal Extents**—Shows the entire horizontal range of the graph. The scale of the curve will be distorted.

 **Zoom Vertical Extents**—Shows the entire vertical range of the graph. The scale of the curve will be distorted.

 **Zoom Horizontally**—Compresses or expands the graph in a horizontal direction.

 **Zoom Vertically**—Compresses or expands the view of the graph in a vertical direction.

 **Zoom**—Zooms in or out around the cursor.

 **Zoom Region**—Draws a rectangular region around any area of the graph, then zooms to that view.

Missing Map Coordinates Dialog

Material Editor > Assign a mapped material to an object that has no UVW Map modifier. > Map level > Coordinates rollout > Change Map Channel to a value other than 1. > Render

Menu bar > File menu > Open > One or more maps in MAX file can't be found.

The Missing Map Coordinates dialog is displayed when you attempt to render one or more objects with materials whose maps use a *map channel (page 3–966)* other than channel 1. (Channel 1 is an exception because channel 1 mapping coordinates are automatically turned on when you assign a mapped material to an object.) Only maps with Show Map In Viewport toggled on are listed as missing.

To use other map channels, you must assign a *UVW Map modifier (page 1–922)* to the object. In UVW Map, set Map Channel to match the value used in the material.

If the material has multiple maps that use multiple channels, you must assign a separate UVW Map modifier for each channel besides channel 1.

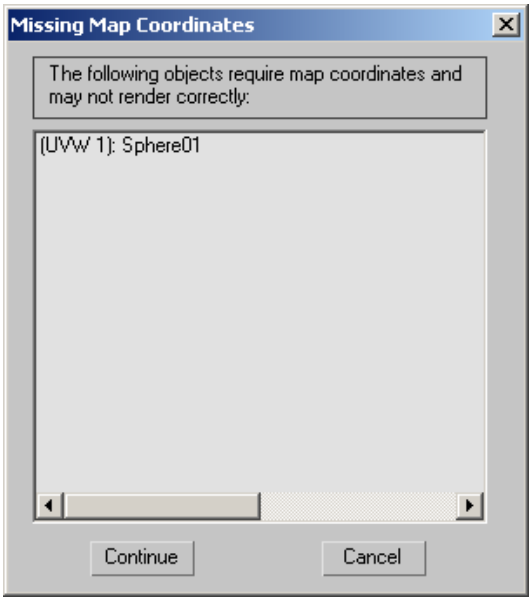
If the object is a *NURBS surface sub-object (page 1–1177)*, you don't need to use UVW Map. Instead, set the map channel on the surface sub-object's *Material Properties rollout (page 1–1149)*.

The Missing Map Coordinates dialog is also displayed when you open a MAX file that

references bitmaps that can't be found in their original location, or are at a location not specified via the *Configure User Paths dialog (page 3–808)*. To open the MAX file, click the Browse button on the Missing Map Coordinates dialog. This opens the Configure User Paths dialog so you can assign one or more paths for the MAX file to access. These settings are then stored with the MAX file.

Note: When you open a MAX file that references bitmaps that can't be found, you might also see a *Missing Map/Photometric Files dialog (page 3–503)*, which lets you browse for the missing files.

Interface



The list shows the map channel followed by the name of the object.

Continue—Proceeds with the rendering. The objects listed will not show maps in the material assigned to them, and might not be visible at all.

Cancel—Cancels rendering.

2D

2D Maps are two-dimensional images that are typically mapped onto the surface of geometric objects, or used as environment maps to create a background for the scene. The simplest 2D maps are bitmaps; other kinds of 2D maps are generated procedurally.

Bitmap (page 2–1631): An image saved as an array of pixels in one of a number of still-image file formats, such as *.tga*, *.bmp*, and so on, or an animation file such as *.avi*, *.mov*, or *.ifl*. (Animations are essentially sequences of still images.) Any of the bitmap (or animation) file types that 3ds Max supports can be used as a bitmap in a material.

Checker (page 2–1638): Combines two colors in a checker pattern. You can replace either color with a map.

Combustion (page 2–1639): Works in conjunction with the Autodesk Combustion product. You can paint directly on a bitmap or object and have the result update in the Material Editor and viewports. The map can include other Combustion effects. Painting and other effects can be animated.

Gradient (page 2–1650): Creates a linear or radial ramp of three colors.

Gradient Ramp (page 2–1652): Creates a great variety of ramps, using as many colors, maps, and blends as you choose.

Swirl (page 2–1656): Creates swirled (spiraling) patterns of two colors or maps.

Tiles (page 2–1658): Creates bricks or other tiled materials with colors or material mappings. Includes commonly defined architectural brick patterns, but you can also customize patterns.

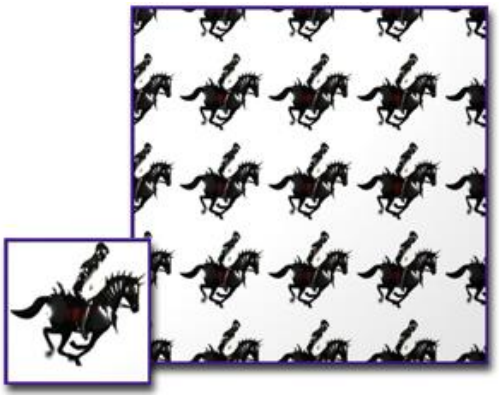
Coordinates Rollout (2D)

Material Editor > Select sample slot. > Get Material > Material/Map Browser > Turn on 2D Maps. > Double-click a map type to apply it to sample slot. > Coordinates rollout is displayed on the Material Editor.

In the Coordinates rollout, by adjusting coordinate parameters, you can move a map relative to the surface of the object to which it is applied and achieve other effects.

Tiling

Often when you apply a bitmap, especially as a texture pattern, you want the pattern to repeat. This effect is known as *tiling*, as in a tiled floor or fountain. You control tiling directly from the Coordinates rollout for any 2D map.



Tiling a map

In default mapping, tiling is active, but because the map is scaled to fit the object, you don't see the effect of tiling unless you offset the UV coordinates or rotate the map. In this case, the portions of the surface from which the bitmap has moved away are filled by other portions of the map. Tiling wraps the object with the map image.

Mirroring

Mirroring a map is an effect related to tiling. It repeats the map and flips the repeated copy.

As with tiling, you can mirror in the U dimension, the V dimension, or both. The Tiling parameter for each dimension specifies how many copies of the map are shown. Each copy is flipped relative to its neighbors.

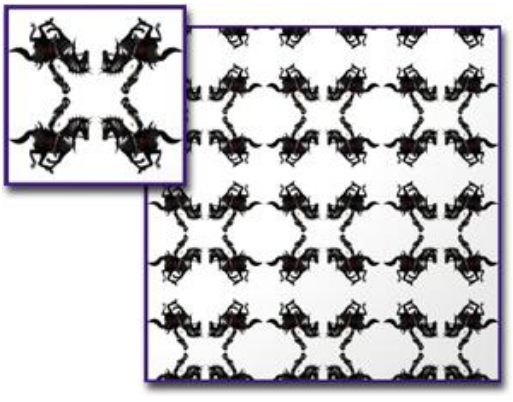


Mirroring a map

Tiling and Mirroring Combined

Because mirroring defaults to two reflected images of the map, the meaning of the Tiling value differs when Mirror is set.

In a single mapping dimension (U or V), a value of 1.0, the default, shows two copies of the bitmap; a value of 2.0 shows four copies; a value of 1.5 shows three copies; and so on. Mirroring in both dimensions multiplies the effect.



Mirroring and tiling a map

Decals

Decals are useful for mapping single designs, small elements such as stickers, or light switches.

A 2D map used as a decal appears only once and is not repeated as with tiling. Wherever the decal doesn't appear on the surface, the surface is rendered as a basic material, with the component colors specified at the material level. With map trees, a decal might appear on top of a different bitmap or other map type.



Decal mapping

Procedures

To set tiling:

1. In the Coordinates rollout, make sure Tile is turned on for the U or V coordinate, or for both.
2. Set the Tiling value for the corresponding coordinates.

   In the Material Editor, the sample slot changes to show the tiling value you chose.

   - The Tiling value is the number of times the bitmap repeats along the specified dimension. A value of 1.0, the default, maps the bitmap exactly once; a value of 2.0 maps the bitmap twice, and so on. Fractional values map a fractional portion of the bitmap in addition to copies of the whole map. For example, a value of 2.5 maps the bitmap two and a half times.
   - Tiling values less than one increase the size of the map relative to the object. For example, a value of 0.5 maps half of the bitmap.
   - Tiling is uniform if both the U and V dimensions are tiled by the same amount.

To preview the effect of tiling:

-  In the Material Editor, use the Tiling flyout to choose a 1x, 2x, 3x, or 4x tiling preview.

  The button you choose changes tiling in the active sample slot to 1.0, 2.0, 3.0, or 4.0 in both U and V.

  The flyout setting has no effect on the material or its mapping. It does not change the Tiling value or the check box setting. It only helps you preview the effect of changing these settings.

To set mirroring:

1. On the Coordinates rollout, make sure Mirror is turned on for the U or V coordinate, or both.

   Mirror and Tile are exclusive settings: if one is set when you choose the other, the original setting turns off.
2. Set the Mirror value for the corresponding coordinate or coordinates.

   The material preview changes to show the Mirror value you chose.

To create a decal:

1. In the Material Editor, choose a bitmap as a diffuse map.
2. On the map's Coordinates rollout, turn off both Mirror and Tile for both the U and V coordinates. Change the mapping to Planar from Object XYZ.
3. Adjust the U and V Tiling parameters to scale the decal.

Changing the Tiling value changes the size of the mapped bitmap as it does when you tile the map, but with Mirror and Tile turned off, the bitmap appears only once.

4. Adjust the U and V Offset parameters to position the decal.

   Note: You can also use the Parameters rollout's Cropping and Placement controls to achieve a decal effect.

To offset the position of a map:

- On the Coordinates rollout, set the U and V values to be between 0.0 and 100, or between -100 and 0.0.

  On the surface of the object, the map shifts in the directions you chose.

  UV offsets are especially useful when you have turned tiling off and want the bitmap to appear in a single location.

To rotate the map:

- On the Coordinates rollout, set the Angle U, V, and W spinners.

  Positive angles rotate the map in the clockwise direction; negative angles rotate it counterclockwise.

  The angle can be up to 360 degrees, which rotates the map completely and has no visible effect unless you are animating the map's rotation.

  You can also click Rotate to use the Rotate Mapping Coordinates dialog, which lets you change the rotation by dragging the mouse.

To increase or decrease antialiasing:

- On the map's Coordinates rollout, increase or decrease the Blur value.

  For diffuse maps and other maps besides bump maps, the Blur value is most effective in the 0.5-20 range. Lower values decrease antialiasing; higher values increase it.

  The Blur Offset parameter adjusts the image before antialiasing Blur is applied. If all you need is antialiasing, leave Blur Offset at its default of 0.0.

To make a map image fuzzier:

- On the map's Coordinate's rollout, increase the Blur Offset value.

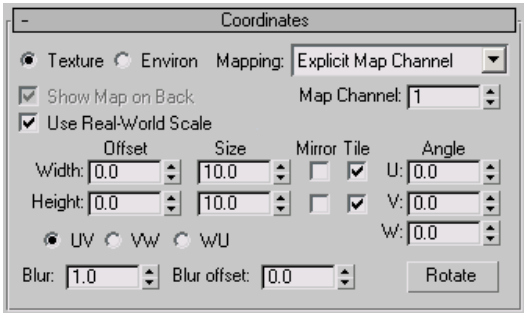
  Blur Offset is a very strong parameter. The Blur Offset spinner has increments of 0.001 Values greater than 0.1 are likely to be too high.

To make a map image sharper:

- On the map's Coordinates rollout, decrease the Blur Offset value to a value below 0.0.

  The negative Blur Offset value sharpens the image.

Interface



These controls appear on the Coordinates rollout for many 2D maps:

Texture—Applies the map as a texture map to a surface. Select the type of coordinates from the Mapping list:

Environ—Uses the map as an *environment map (page 3–934)*. Select the type of coordinates from the Mapping list:

Mapping list—Options vary depending on choice of Texture or Environ mapping:

- **Explicit Map Channel—**Uses any map channel. When selected, the Map Channel field becomes active, and you can choose any channel from 1 to 99.
- **Vertex Color Channel—**Uses assigned vertex colors as a channel. See *Editable Mesh (page 1–996)* for details on assigning vertex color.

  See also *Vertex Color Map (page 2–1693)* and *Assign Vertex Colors Utility (page 2–1734)*.
- **Planar from Object XYZ—**Uses planar mapping based on the object's local coordinates (disregarding the pivot point location). For rendering purposes, planar mapping doesn't project through to the back of the object unless you turn on Show Map On Back.
- **Planar from World XYZ—**Uses planar mapping based on the scene's world coordinates (disregarding the object's bounding box). For rendering purposes, planar mapping doesn't project through to the back of the object unless you turn on Show Map On Back.
- **Spherical Environment, Cylindrical Environment, Shrink-wrap Environment—**Project the map into the scene as though it were mapped to an invisible object in the background. See *Environment Map (page 3–934)*.
- **Screen—**Projects as a flat backdrop in the scene.

Show Map on Back—When on, planar mapping (Planar from Object XYZ, or with the UVW Map modifier) projects through to render on the back of the object. When off, planar mapping doesn't render on the object's back. Default=on.

This toggle is available only when Tiling is off in both dimensions. Its effect is visible only when you render the scene.

Note: In viewports, planar mapping always projects to the back of the object, whether Show Map On Back is turned on or not. To override this, turn off Tiling.

 **Use Real-World Scale—**When turned on, applies the map to objects using the real-world Width and Height values instead of UV values. Default=on.

When Real-World Scale is on, the texture placement is relative to the corner of the texture map so alignment with architectural objects likes walls is more efficient. When off, the texture placement is relative to the center of the texture map.

Offset—Changes the position of the map in *UV coordinates (page 3–1028)*. The map moves in relation to its size. For example, if you want to shift the map its full width to the left, and half its width downward from its original position, you enter **-1** in the U Offset field and **0.5** in the V offset field.

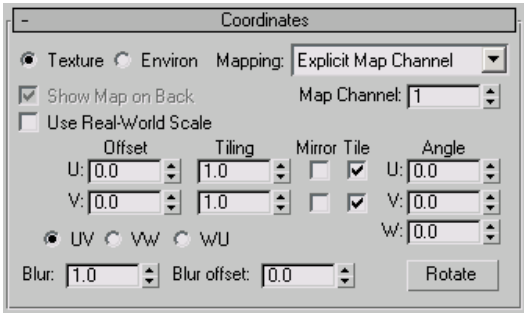
UV/VW/WU—Changes the *mapping coordinate (page 3–967)* system used for the map. The default UV coordinates project the map onto the surface like a slide projector. The VW and WU coordinates rotate the map so that it is perpendicular to the surface.

Tiling—Determines the number of times the map is *tiled (page 3–1022)* (repeated) along each axis.

Mirror—*Mirrors (page 3–1022)* the map left-to-right (U axis) and/or top-to-bottom (V axis).

Tile—Turns tiling on or off in the U or V axis.

When Use Real-World Scale is off



Offset (UV)—Changes the position of the map in *UV coordinates (page 3–1028)*. The map moves in relation to its size. For example, if you want to shift the map its full width to the left, and half its width downward from its original position, you enter **-1** in the U Offset field and **0.5** in the V offset field.

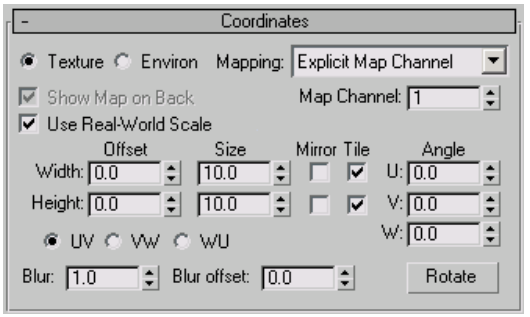
UV/VW/WU—Changes the *mapping coordinate (page 3–967)* system used for the map. The default UV coordinates project the map onto the surface like a slide projector. The VW and WU coordinates rotate the map so that it is perpendicular to the surface.

Tiling—Determines the number of times the map is *tiled (page 3–1022)* (repeated) along each axis.

Mirror—*Mirrors (page 3–1022)* the map left-to-right (U axis) and/or top-to-bottom (V axis).

Tile—Turns tiling on or off in the U or V axis.

When Use Real-World Scale is on



Offset (Width/Height)—Move the map horizontally or vertically along the width or height of the object to which the material is applied. The offset distance is relative to the lower-left corner of the map.

UV/VW/WU—Changes the *mapping coordinate (page 3–967)* system used for the map. The default UV coordinates project the map onto the surface like a slide projector. The VW and WU coordinates rotate the map so that it is perpendicular to the surface.

Size—Determines the real world width and height of the map.

For example, if you scan a piece of marble that is 12” x 8” and then assign this image as the Diffuse Map, you can type 12” (or 1’) and 8” as the Width and Height. This ensures that the scale of the marble is correct in the rendered scene.

Note: The default setting for the texture size can be set using the Default Texture Size option in the *Material Editor Options Dialog (page 2–1436)*.

Mirror—*Mirrors (page 3–1022)* the map horizontally and/or vertically.

Tile—Turns horizontal and/or vertical tiling on or off.

Note: If the Use Real-World Size switch is turned off in the Material Editor, the Real-World Map Size settings in modifiers like UVW Map or for primitives like Box will not work. Likewise, moving vertices at a sub-object level or scaling an object, in general, will not honor the Use Real-World Scale settings.

Angle U/V/W—Rotates the map about the U, V, or W axis (in degrees).

Rotate—Displays a schematic Rotate Mapping Coordinates dialog that lets you rotate the map by dragging on an arcball diagram (similar to

the arcball used to rotate viewports, although dragging inside the circle rotates along all three axes, and dragging outside it rotates about the W axis only). The Angle UVW values change as you drag in the dialog.

Blur—Affects the sharpness or blurriness of the map based on its distance from the view. The farther away the map is, the greater the blurring. The Blur value blurs maps in world space. Blur is primarily used to avoid *aliasing (page 3–907)*.

Blur Offset—Affects the sharpness or blurriness of the map without regard to its distance from the view. Blur Offset blurs the image itself in object space. Use this option when you want to soften or defocus the details in a map to achieve the effect of a blurred image.

See *Blur/Blur Offset (page 3–918)*.

Noise Rollout (2D)

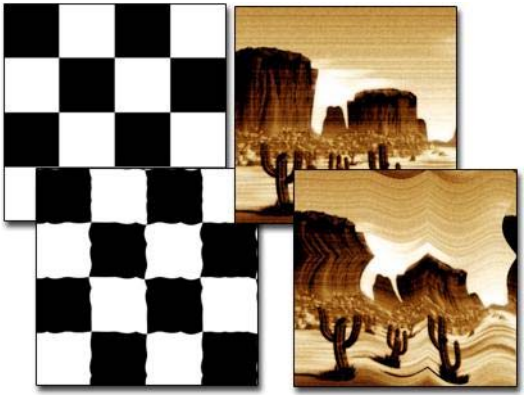
Material Editor > Select sample slot. > Get Material > Material/Map Browser > Turn on 2D Maps. > Select a map type. > Drag map from Browser thumbnail to sample slot. > Noise rollout is displayed in the Material Editor.

You can add a random noise to the appearance of your material. Noise perturbs the UV mapping of pixels by applying a fractal noise function.

Noise patterns can be very complex and are a versatile way to create apparently random patterns. They are also good for simulating surfaces found in nature, as is characteristic of fractal images.

Noise parameters interact closely with each other. Slight variations in each can create noticeably different effects.

Note: Noise settings aren't displayed in viewports.



Above: A checker map and a bitmap

Below: The same maps with noise applied

Procedures

To add noise to a material:

1. In the Noise rollout, select On.
2. Adjust the three noise parameters to get an effect you like.

To remove noise from a material:

- In the Noise rollout, turn off On.

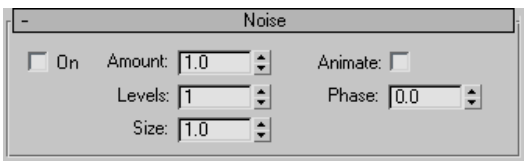
  Noise is no longer applied to the map.

To animate the noise effect:

1. Turn on the Auto Key button.
2. Move to a non-zero frame.
3. In the Noise rollout, turn on Animate.

   By default, animation keys are set at either end of the active frame range.
4. Change the Phase value at different keyframes.

Interface



These controls appear on the Noise rollout for many 2D maps:

On—Determines whether the Noise parameters affect the map.

Amount—Sets the strength of the fractal function, expressed as a percentage. If the amount is 0 there is no noise. If the amount is 100 the map becomes pure noise. Default=1.0.

Levels—Or iterations: the number of times the function is applied. The effect of the level is dependent on the Amount value. The stronger the amount, the greater the effect of increasing the Levels value. Range=1 to 10; Default=1.

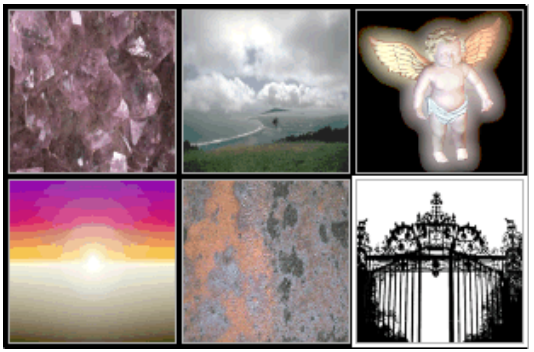
Size—Sets the scale of the noise function relative to geometry. At very small values, the noise effect becomes white noise. At large values, the scale can exceed the scale of the geometry, in which case it has little or no effect. Range=0.001 to 100; Default=1.0.

Animate—Determines whether animation is on the noise effect. This parameter must be turned on if you intend to animate the noise.

Phase—Controls the speed of the animation of the noise function.

Bitmap 2D Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Bitmap



Bitmaps shown in Material Editor sample slots

A *bitmap* is an image produced by a fixed matrix of colored pixels, like a mosaic. Bitmaps are useful for creating many kinds of materials, from wood grains and wall surfaces to skin and feathers. You can also use an animation or video file instead of a bitmap to create an animated material.

When you assign the Bitmap map, the *Select Bitmap Image File dialog (page 2–1635)* opens automatically. Use this dialog to specify a file or sequence as the bitmap image.

The Bitmap map can synchronize the frames of a bitmap sequence to the age of particles to which the map is applied. With this effect, each particle displays the sequence from the start when it is born, rather than being assigned whichever frame is current. This is achieved by turning on the *Sync Frames To Particle Age check box*. Also, when using Particle Flow, assign the material containing the Bitmap map to a Material Dynamic operator. For more details and a procedure, see *Material Dynamic Operator (page 2–191)*.

Note: To save loading time, if a map with the same name is in two different locations (in two different paths), it is loaded only once. This poses a problem only if your scene includes two maps

that have different content but the same name. In this case, only the first map encountered will appear in the scene.

Supported File Types

The Material Editor supports the following file formats:

AVI files (page 3–609)

BMP files (page 3–610)

CIN files (Kodak Cineon) (page 3–610)

DDS files (page 3–611)

GIF files (page 3–613)

IFL files (page 3–616)

JPEG files (page 3–620)

MOV files (QuickTime Movies) (page 3–621)

MPEG files (page 3–621)

PNG files (page 3–628)

PSD files (page 3–629)

RGB files (SGI format) (page 3–633)

RLA files (page 3–630)

RPF files (page 3–631)

TGA files (Targa) (page 3–633)

TIFF files (page 3–634)

YUV files (page 3–635)

Note: If your scene includes animated bitmaps with materials, projector lights, or environments, the animation file is reloaded once per frame. If your scene uses multiple animations, or if the animations are large files, rendering will be slower.

See also

For functionality shared with other 2D maps, see the following topics:

Coordinates Rollout (2D) (page 2–1625)

Noise Rollout (2D) (page 2–1630)

Output Rollout (page 2–1621)

Procedures

To crop an image:

1. On the Bitmap Parameters rollout, click the Bitmap button and assign a bitmap.
2. In the Cropping/Placement group, turn on Apply to see the results of cropping in the sample slot (and in shaded viewports, if Show Map In Viewport is active).
3. Turn on Crop.
4. Click View Image to display the bitmap.

   A frame window appears, displaying the image surrounded by a region outline (a dashed line at the outer edges of the image, with handles on the sides and corners).
5. Specify a cropping region by adjusting the spinners at the top of the window, or by dragging the region outline.

To place an image:

1. On the Bitmap Parameters rollout, click the Bitmap button and assign a bitmap.
2. In the Cropping/Placement group, turn on Apply to see the results of cropping in the sample slot (and in shaded viewports if Show Map In Viewport is active).
3. Turn on Place.
4. Click View Image.

   A frame window appears, displaying the image surrounded by a region outline (a dashed line at the outer edges of the image, with handles on the sides and corners).
5. Move the image by adjusting the spinners at the top of the window, or by dragging the region outline.

The reduced image "decals" on the sample sphere. The diffuse color is visible around the image.

To use the alpha channel that is part of the bitmap:

1. On the Maps rollout, assign the map to the Opacity component.

   (You can assign a copy or instance of this map to other components, such as Diffuse, as well.)

2. Click the map button for the Opacity component.

   This lets you adjust the settings for the Opacity map.

3. In the Bitmap Parameters rollout > Alpha Source group, choose Image Alpha.

   This option is not available if the bitmap does not have an alpha channel.

4. In the Bitmap Parameters rollout > Mono Channel Output group, choose Alpha.

   This option is not available if the bitmap does not have an alpha channel.

   The bitmapped material will now have the transparency specified by the alpha channel. This will appear in renderings. Transparency does not appear in viewports or ActiveShade renderings.

To create an alpha channel based on intensity:

- In Bitmap Parameters rollout > Alpha Source group, turn on RGB Intensity.

  The software creates an alpha channel. Full-intensity areas of the image are opaque, zero-intensity areas are transparent, and intermediate colors become partially transparent.

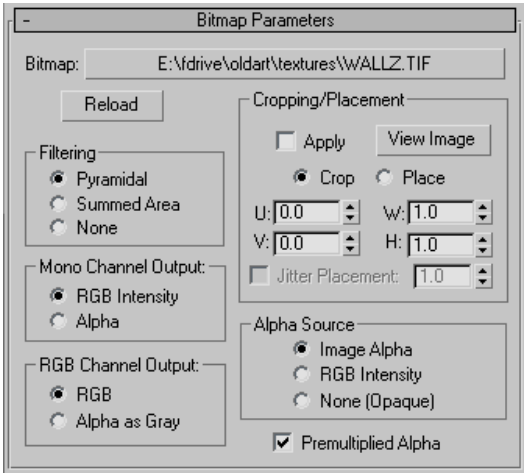
To use a completely opaque bitmap:

- In Bitmap Parameters rollout > Alpha Source group, turn on None (opaque).

  The software ignores the bitmap's alpha channel, if present, and does not create a new one.

Interface

Bitmap Parameters rollout



Bitmap—Selects the bitmap using the standard file browser. After selection, the full path name appears on this button.

Reload—Reloads the bitmap file using the same name and path. You don't need to use the file browser to reload the bitmap after you've updated it in your paint program.

Clicking reload for any instance of the map updates the map in all sample slots and in the scene.

Filtering group

Filtering (page 3–939) options let you select the method of pixel averaging used in *antialiasing (page 3–907)* the bitmap.

Pyramidal—Requires less memory and is adequate for most purposes.

Summed Area—Requires much more memory, but yields generally superior results.

None—Turns off filtering.

Mono Channel Output group

Some parameters, such as opacity or specular level are a single value as opposed to a material's three-value color components. Controls in this group determine the source of the Output mono channel in terms of the input bitmap.

RGB Intensity—Uses the intensity of the red, green, and blue channels for mapping. The color of the pixels is ignored and only the value or luminance of the pixels is used. The colors are computed as gray values in the range between 0 (black) and 255 (white).

Alpha—Uses the intensity of the *alpha channel (page 3–907)* for mapping.

RGB Channel Output group

The RGB Channel Output determines where the output RGB part comes from. The controls in this group affect only maps for material components that display color: Ambient, Diffuse, Specular, Filter Color, Reflection, and Refraction.

RGB—Displays the full color values of the pixels. (Default)

Alpha as Gray—Displays tones of gray based on the levels of the alpha channel.

Cropping/Placement group

The controls in this group let you crop the bitmap or reduce its size for custom placement. Cropping a bitmap means to reduce it to a smaller rectangular area than it originally had. Cropping doesn't change the scale of the bitmap.

Placing a bitmap lets you scale the map and place it anywhere within its tile. Placing can change the bitmap's scale, but shows the entire bitmap. The four values that specify the placement and size of the cropping or placement region are all animatable.

Cropping and placement settings affect the bitmap only as it's used for this map and any instances of the map. They have no effect on the bitmap file itself.

Apply—Turn on to use the cropping or placements settings.

View Image—Displays a *rendered frame window (page 3–5)* that shows the bitmap surrounded by a region outline. The region outline has handles at its sides and corners. When cropping is on, dragging the handles changes the size of the crop area. You can also drag within the region area to move it.

The frame window has U/V and W/H (width/height) spinners on its toolbar. Use these to adjust the location and size the image or crop area.

When Place is turned on, dragging the region area handles changes the scale of the bitmap (hold down Ctrl to preserve the bitmap's aspect ratio), and dragging the image changes its location within the tile area.

When Crop is turned on, the UV or XY button at the right of the rendered frame window toolbar lets you switch between using UV or XY coordinates in the toolbar spinners. Also, you can zoom out by pressing Shift+Z and zoom in by pressing Z.

Crop—Makes cropping active.

Place—Makes placement active.

U/V—Adjusts the bitmap location.

W/H—Adjusts the width and height of the bitmap or crop area.

Jitter Placement—Specifies the amount of random offset. At 0, there is no random offset. Range = 0.0 to 1.0

When Place is turned on, the size and position specified by the spinners or editing window are ignored. The software then chooses a random size and tile position for the image.

Alpha Source group

Controls in this group determine the source of the Output alpha channel in terms of the input bitmap.

Image Alpha—Uses the image's alpha channel (disabled if the image has no alpha channel).

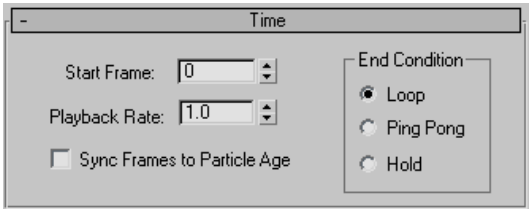
RGB Intensity—Converts the colors in the bitmap to grayscale tonal values and uses them for transparency. Black is transparent and white is opaque.

None (Opaque)—Does not use transparency.

Premultiplied Alpha—Determines how alpha is treated in the bitmap. When turned on, the default, *premultiplied alpha (page 3–997)* is expected in the file. When turned off, the alpha is treated as non-premultiplied, and any RGB values are ignored.

Tip: If you apply an alpha image as a Diffuse map, for example, and it doesn't decal correctly, the bitmap file probably contains non-premultiplied alpha; the RGB values are maintained separately from the alpha values. To correct this, turn off Premultiplied Alpha.

Time rollout



These controls let you change the start time and speed of animation (*AVI (page 3–609)* or *MOV (page 3–621)*) files used as animated texture maps. They make it easier to use sequences of images as maps in scenes, because you can control the timing very precisely

Start Frame—Specifies the frame where the playback of the animated map will begin.

Playback Rate—Lets you speed up and slow down the rate that the animation is applied to the map (for example, 1.0 is normal speed, 2.0 is twice as fast, .333 is 1/3 as fast).

Sync Frames to Particle Age—When on, the software synchronizes the frames of a bitmap sequence to the age of particles to which the map is applied. With this effect, each particle displays the sequence from the start when it is born, rather than being assigned whichever frame is current. Default=off.

When using Particle Flow, assign the material containing the Bitmap map to a Material Dynamic operator. For more details and a procedure, see *Material Dynamic Operator (page 2–191)*.

Note: This functionality is not supported by the mental ray renderer.

End Condition group

Determines what happens after the last frame of the bitmap animation if the animation is shorter than the scene.

Loop—Causes the animation to repeat over and over again from the beginning.

Ping-Pong—Causes the animation to be played forward and then backward repeatedly, making every animated sequence "loop smoothly."

Hold—Freezes on the last frame of the bitmap animation.

Select Bitmap Image File Dialog

Material Editor > Maps Rollout > Click any map selector button. > Material/Map Browser > Double-click Bitmap. > Select Bitmap Image.

Material Editor > Bitmap map > Bitmap Parameters rollout > Bitmap button

The Select Bitmap Image dialog allows you to choose a file or sequence of files for a map. If

a sequence of files is selected by turning on Sequence, the *Image File List Control dialog (page 3–618)* is opened when you click Setup or Open.

Procedures

To select a bitmap image for a map:

1. In the Material Editor, open the Maps rollout.
2. Click any button in the Map column.

   This adds a map into the channel you've selected. For example, clicking in the Map column of the Diffuse channel creates a diffuse or texture map.

   The Material Map Browser is displayed.
3. In the Material/Map Browser, double-click Bitmap.

   The Select Bitmap Image dialog is displayed.
4. In the Select Bitmap Image dialog, navigate the Look In field to select the appropriate directory.

   Note: The Select Bitmap Image File dialog uses the last location where a bitmap was chosen, rather than the default bitmap path defined in Customize > Configure User Paths.
5. Highlight the file name in the file list window.
6. Click Open to close the dialog.

To select a set of still images as a bitmap sequence:

1. In the Material Editor, open the Maps rollout.
2. Click any button in the Map column.

   This adds a map into the channel you've selected. For example, clicking in the Map column of the Diffuse channel creates a diffuse or texture map.

   The Material Map Browser is displayed.
3. In the Material/Map Browser, double-click Bitmap.

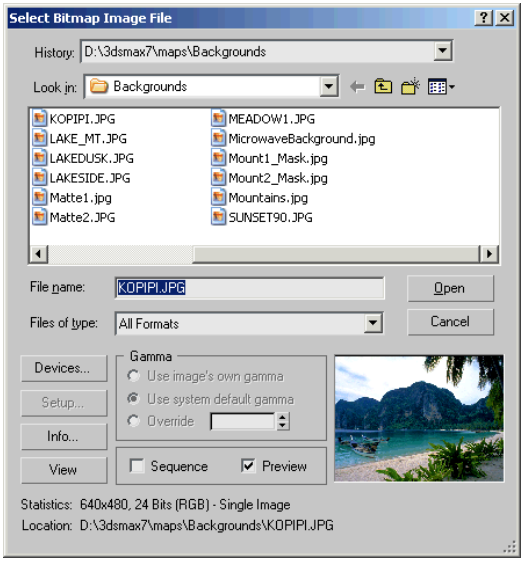
   The Select Bitmap Image dialog is displayed.
4. In the Select Bitmap Image dialog, navigate the Look in field to select the directory containing the sequence of files.
5. If necessary, change file type to match the file extension of the sequence, or choose All Formats.
6. Turn on Sequence, and choose the name of the first sequential file.
7. Click the Setup button.

   The Image File List Control Dialog opens.
8. Click the Browse button and set the Target path to a writable directory on your hard disk. Do not set the path to a CD-ROM drive.
9. Choose the options you want and click OK.

   The IFL file is written to the target directory.

Interface

Tip: You can resize the dialog by dragging an edge or a corner.



History—Displays a list of the most recent directories searched. Whenever an image is selected, the path used is added to the top of the history list as the most recently used path. The

history information is saved in the *3dsmax.ini (page 1–18)* file.

Look In—Opens a navigation window to browse other directories or drives.


Up One Level—Moves you up a level in the directory structure.


Create New Folder—Lets you create a new folder while in this dialog.


List—Displays the contents of a directory without the details.


Details—Displays the contents of a directory with all the details.

List Window—When details are turned on, the contents of the directory are displayed with Name, Size, Type, Date Modified, and Attributes. You can sort based on each of these columns by clicking the column label.

File Name—Displays the file name of the file selected in the list.

Files of Type—Displays all the file types that can be displayed. This serves as a filter for the list.

Open—Selects the highlighted file and closes the dialog.

Cancel—Cancels the selection and closes the dialog.

Devices—Lets you choose the hardware output device, for example, a digital video recorder. The device, its driver, and its 3ds Max plug-in must all be installed on your system to use the device.

Setup—When Sequence is turned on, and there are sequential files in the directory displayed, this Setup displays an *Image File List Control dialog (page 3–618)* to create an IFL file.

Info—Displays expanded information about the file, such as frame rate, compression quality, file size, and resolution. The information here depends on the type of information that is saved with the file type.

View—Displays the file at its actual resolution. If the file is a movie, the Media Player is opened so the file can be played.

Gamma Group

Gamma—Selects the type of gamma to be used for the selected file. This is unavailable unless Enable Gamma Correction is turned on in the *Gamma panel (page 3–824)* of the Preferences dialog.

Use Image's Own Gamma—Uses the gamma of the incoming bitmap.

Use System Default Gamma—Ignores the image's own gamma and uses the system default gamma instead, as set in the Gamma panel of the Preferences dialog.

Override—Defines a new gamma for the bitmap that is neither the image's own, nor the system default.

Sequence—Creates an "Image File List " using the given information. Note that each time an image is selected, an evaluation is done to see if an IFL sequence can be created. If the selected image does not yield a list, this option box is unavailable. In the past, it was necessary to enter a wild card in order to create a list. Now it is possible to use wild card to filter files in the file selector.

Preview—Displays the image as a thumbnail in the image window.

Image Window—Displays a thumbnail of the selected file.

Statistics—Displays the resolution, color depth, file type, and number of frames of the selected file.

Location—Displays the full path for the file. With this information at the bottom of the dialog, you always know exactly where you are.

Checker Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Checker



Checker map used for the tablecloth and (in a composite) for the floor of the ice-cream shop

The Checker map applies a two-color checkerboard pattern to the material. The default checker map is a pattern of black and white squares. Checker maps are 2D procedural maps. The component checkers can be either colors or maps.

Tip: Turning on Noise for a Checker map can be an effective way to create irregular patterns with a natural appearance.

See also

For functionality shared with other 2D maps, see:

Coordinates Rollout (2D) (page 2–1625)

Noise Rollout (2D) (page 2–1630)

Procedures

To create a Checker map:

1. Click a map button to assign a map.
2. Choose Checker in the Material/Map browser, and then click OK.

To change the color of one set of squares:

1. In the Checker Parameters rollout, click a color swatch to display the *Color Selector (page 1–161)*.
2. Adjust the color.
3. Click Close.

To use a map for one set of squares:

- In the Checker Parameters rollout, click a map button to assign a map to a color.

To soften the edge between the two checker colors:

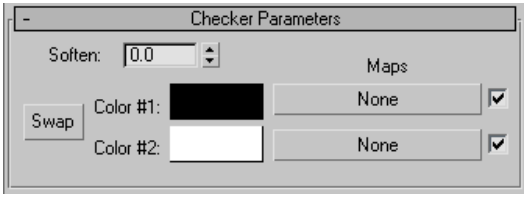
- In the Checker Parameters rollout, increase the value of Soften.

  When Soften equals 0.0, there is a hard edge between the checker colors. Low positive values soften or blur the checker boundary. Larger Soften values can blur the entire material.

To swap the two checker components:

- In the Checker Parameters rollout, click Swap.

Interface



Soften—Blurs the edges between the checkers. A little blurs a lot.

Swap—Switches the position of the two checkers.

Color #1—Sets the color of one of the checkers. Click to display the *Color Selector (page 1–161)*.

Color #2—Sets the color of one of the checkers. Click to display the *Color Selector (page 1–161)*.

Maps—Selects a map to use within the area of the checker color. For example, you could put an additional checkerboard within one of the checker colors.

The check boxes enable or disable their associated map.

Combustion Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Combustion

With the Combustion map, you can create maps interactively using the Autodesk Combustion software and 3ds Max at the same time. You use Combustion to paint on a bitmap, and the material updates automatically in the Material Editor and in shaded viewports.

Important: The Combustion map works only if Autodesk Combustion is installed on your system.

Only Combustion 2.1 and later formats are supported. Maps in the Combustion 1 format are not supported in 3ds Max.

Important: The *mental ray renderer (page 3–78)* does not support the Combustion map.

See also

CWS (Combustion Workspace) Files (page 3–611)

Noise Rollout (2D) (page 2–1630)

About the 3ds Max and Combustion Integration

You can use Combustion as a material map in 3ds Max. With a Combustion map, you can create a material from a Paint or composite operator, and in turn apply that material to objects in a 3ds Max scene. The Combustion map can include Combustion effects, and it can be animated.

In addition, with Combustion you can import 3ds Max scenes that have been rendered to a *rich pixel file (page 3–631)* (*RPF (page 3–631)* or *RLA (page 3–630)* file). The imported rich pixel rendering becomes an element of your composite. You can adjust its 3D position relative to video elements of the composite, and you can apply Combustion 3D Post effects to objects within it. See the *Combustion User's Guide* for more information.

Note: Because 3ds Max runs only on Windows, you cannot use Combustion to create material maps on a Macintosh.

Note: The environmental atmospheric effect known as "Combustion" in versions of 3ds Max prior to v4 is now known as the *Fire effect (page 3–276)*.

3ds Max Materials and the Combustion Map

In 3ds Max, a *material (page 3–971)* is data that you assign to the surface or faces of an object so that it appears a certain way when rendered. Materials affect the color of objects, their shininess, their opacity, and so on.

The *Material Editor (page 2–1409)* is the portion of 3ds Max that creates and manages materials. In the Material Editor, you can assign maps to a material's color components and to its numeric components such as opacity. Maps add images, patterns, color adjustments, and other effects to the visual properties of the material.

In the 3ds Max Material Editor, you assign a map by clicking the map button for a component color or other component. This displays the Material/Map Browser, which lets you choose the map type.

3ds Max provides several types of *maps (page 3–968)*. The most basic is a 2D map, a two-dimensional image that is typically mapped onto the surface of geometric objects.

Other uses of 2D maps are as environments to create a background for the scene, as projections from lights, and as displacements to "emboss" geometry.

A Combustion map is a *2D map (page 2–1624)*. It is a Combustion project used by the 3ds Max Material Editor, so like any Combustion project, it is vector-based, animatable, and fully editable. From within the Material Editor, you can have Combustion create a new project from scratch, or use an existing composite or Paint branch. You can synchronize the Combustion Timeline with the 3ds Max time slider so animated materials synchronize with your 3D scene.

With a Combustion map, you can paint in either program: that is, you can paint either in the Combustion viewport or on 3ds Max objects. Both programs update the paint display. You also have the option of using Combustion to paint on an "unwrapped" projection of 3ds Max object geometry.

In addition, with Combustion effects that require you to pick a point, such as Lens Flare or Ripple, you can use either program, Combustion or 3ds Max, to pick the point.

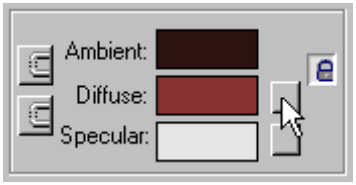
Tips for Working with a Combustion Map in 3ds Max

- If you have a dual-screen configuration, you can set it up so you can see both the 3ds Max and the Combustion windows at the same time. Otherwise, you need to use Alt+Tab to switch between the two windows.
- To work with Combustion, the 3ds Max object must have *mapping coordinates (page 3–967)*. Primitive objects have a Generate Mapping Coordinates toggle, which is automatically enabled when you assign a mapped material to the object. Some objects, such as editable meshes, do not have a Generate Mapping Coordinates toggle. For these kinds of objects, go to the Modify panel and apply a *UVW Map modifier (page 1–922)*.
- Sometimes it can be hard to see how the Combustion operator is oriented to the 3ds Max object's mapping coordinates. It can help to paint some temporary strokes in Combustionto see how they are aligned in 3ds Max viewports. Displaying the mapping coordinates in Combustion can help. See the procedure, "To display an unwrapped mesh." It can also help to paint directly on the object in a 3ds Max viewport. See the procedure, "To paint directly on the 3D object."

Procedures

To create a new Combustion map:

1. Open the Material Editor.
2. Drag an unused sample slot from the Material Editor to the object you want to paint.
3. In the Material Editor, click the map button for the Diffuse Color component. This button is on the material's Basic Parameters rollout.



All standard materials have a Basic Parameters rollout, whose controls vary depending on the chosen shader. The Strauss shader has only one color component, labeled Color.

The Material/Map Browser appears.

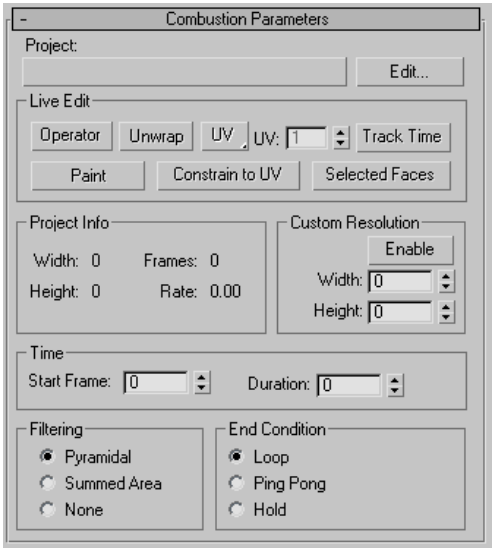
4. In the Material/Map Browser, choose Combustion, and click OK.

A Combustion map is assigned to the Diffuse Color, and a black material map appears in the active sample slot.

5.  Click to turn on Show Map In Viewport.

   In the scene, the object turns black in shaded viewports.

6. In the Parameters rollout, click Edit.

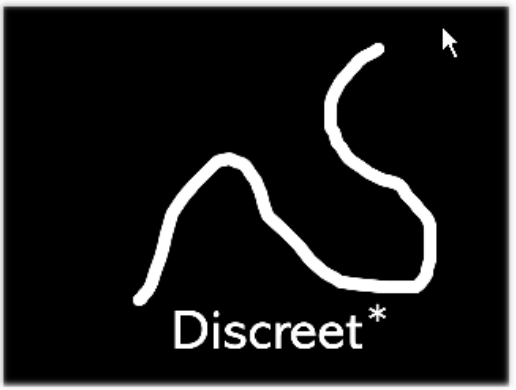


   This launches Combustion, which displays the New Workspace dialog.

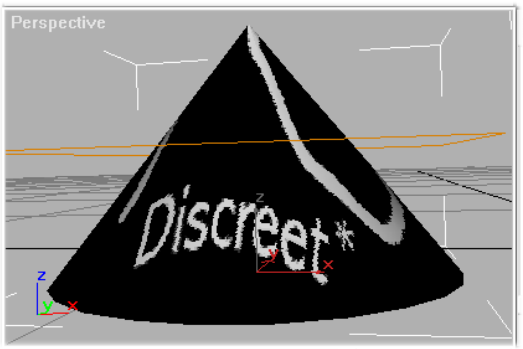
7. Set up the new project.

   The composite or Paint branch that you create in Combustion appears on the object in 3ds Max viewports, as well as in the sample slot for the material with the Combustion map. The workspace name and path are assigned to the material, and appear on the Project button in the material's Combustion Parameters rollout.

For example, you can use the Paint operator in Combustion. When you release the mouse, the stroke appears on the 3ds Max object.



Paint operator in Combustion



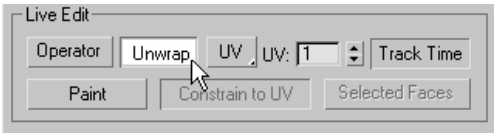
Painted object in 3ds Max

To display an unwrapped mesh:

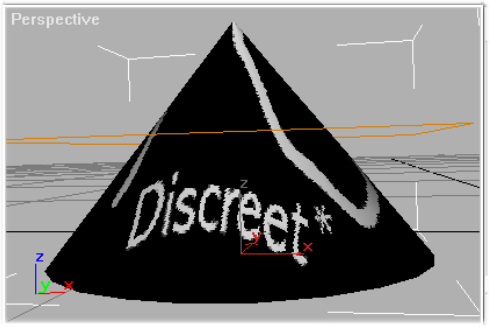
In the 3ds Max Material Editor, you can use the Unwrap Mesh feature to display your 3D object as a 2D mesh in Combustion. You can adjust the color and size of the mesh.

The mesh display is only an overlay to help you orient paint strokes and other Combustion effects. It is displayed in Combustion but is not a part of the composite or the map.

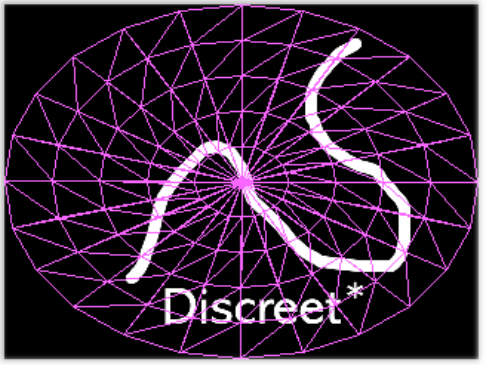
1. Create a Combustion map.
2. In the 3ds Max Material Editor, enable Unwrap Selected in the Live Edit group.



In Combustion, a mesh appears. This is an "unwrapped" projection of the 3D object.



3D object in 3ds Max



Corresponding mesh in Combustion

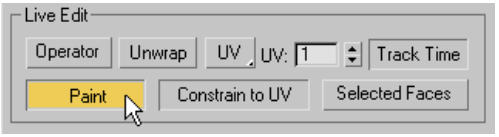
To set the mesh parameters:

- In Combustion, choose File > Preferences > Mesh.

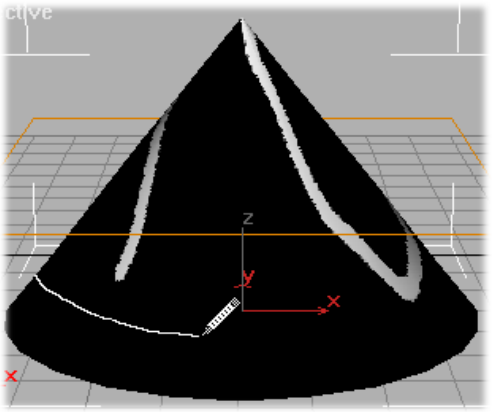
Use	To
Anti-Alias Mesh	Remove jaggies from the mesh.
Display During Playback	Display the mesh when you play back the animation.
Color	Click the color box to set the color of the mesh using a color picker.

To paint directly on the 3D object:

1. Create a Combustion map.
2. In Combustion, select one of the following drawing tools:
   -  Freehand
   -  Straight Line
   -  Rectangle
   -  Ellipse
3. In the 3ds Max Material Editor, enable Paint in the Live Edit group of the Combustion Parameters rollout.



In the 3ds Max viewport, a pen cursor appears. Drag the cursor over the object to paint on it.



When you release the cursor, the Paint object also appears in Combustion.

To animate Combustion Paint strokes:

1. Create a Combustion map.
2. In Combustion, set the time scale to start at frame number 0. Choose File > Preferences. In the Preferences dialog select General, set Display Time As to Frames (From 0), and then click OK.
3. In the 3ds Max Material Editor, enable Track Time in the Live Edit group of the Combustion Parameters rollout.

   Now the time slider in 3ds Max controls the Timeline indicator in Combustion.
4. In 3ds Max, move the time slider to a frame and create a Paint object.

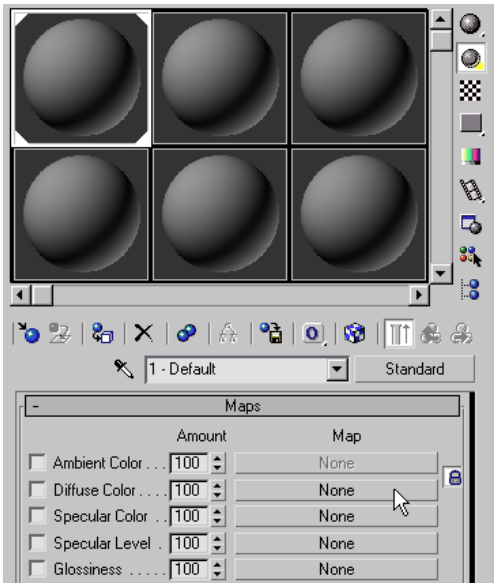
   The Paint object appears on that frame in both Combustion and 3ds Max.
5. Move to another frame and use Combustion to modify the Paint object. Combustion uses interpolation to determine the appearance of the Paint object between keyframes. If you add a new Paint object, that object simply appears, starting on the frame where you created it.
6. Continue advancing in the clip, adding and modifying Paint strokes and effects to create your animated material. For more information on animating objects in Combustion, refer to the *Combustion User's Guide*.

   You can add Paint strokes in either program, but to modify them you must use Combustion.

   Note: Remember, Combustion tracks the time slider in 3ds Max, but 3ds Max does not track the Timeline indicator in Combustion. If the 3ds Max viewport does not appear to be updating as you paint in Combustion, you might be painting on a different frame than the one displayed in 3ds Max. To find your Paint objects, move to the correct frame in 3ds Max.

To use an existing Combustion workspace as a material map in 3ds Max:

1. Open the Material Editor and select an unused sample slot.
2. In the Material Editor, expand the Maps rollout, and click the Map button next to the Diffuse Color component.



The Material/Map Browser appears.

3. In the Material/Map Browser, choose Combustion, and click OK.

   A Combustion map is assigned to the Diffuse Color, and a black material map appears in the active sample slot.

4. In the Combustion Parameters rollout, click the Project bar.



   The Open Project dialog appears.

5. Browse for the workspace file (*.cws)* that you want to use as a map, and click the Open button.

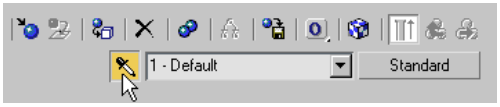
   The Combustion workspace name and path appear in the Project button.

   To apply the map to an object, drag the sample slot from the Material Editor to the object in a 3ds Max viewport.

   To edit the map, click the Edit button in the Parameters rollout. In Combustion, the workspace corresponding to the selected map opens, and you can edit the image.

To paint geometry with a bitmapped material already assigned to it:

1. In 3ds Max, select the object that you want to paint.
2. In the Material Editor, select an unused sample slot.
3. Click the Pick Material From Object button, then click the object in the viewport to put the object's material in the selected sample slot.



4. Open the Maps rollout and note the name of the bitmap file. Click Map to go to the bitmap level of the material.
5. Click the map's Type button.

   The Material/Map Browser appears.

6. In the Material/Map Browser, choose Combustion to change the type from Bitmap to a Combustion map.
7. On the Paint Parameters rollout, click the blank Project button, and then choose the same bitmap.



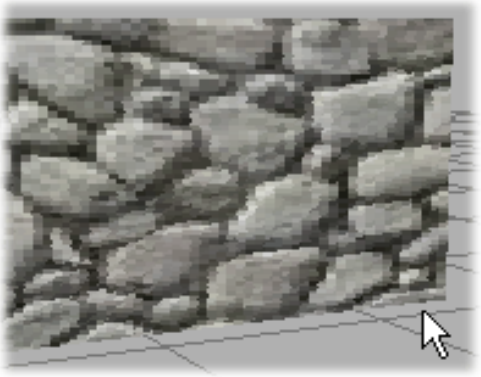
8. Click Edit.

   Combustion is launched and the Import Footage dialog appears. Import the same bitmap.

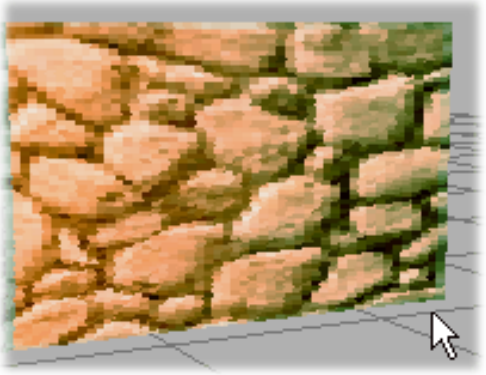
   To paint on the bitmap, select Paint. You can also key or color correct the bitmap, or use it to build a composite. For more information, see the *Combustion User's Guide*.

9.  In the 3ds Max Material Editor, click Show Map In Viewport.

   In the scene, the object is mapped in shaded viewports.



Object with original bitmap



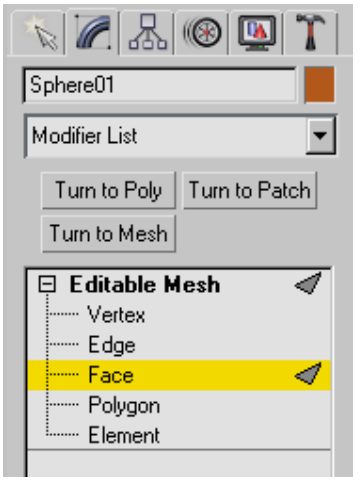
Object with painted bitmap

To paint selected faces:

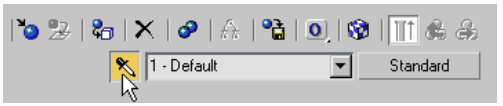
Use a Multi/Sub-Object material to control the location of your painting. Any sub-material can have a Combustion map, so you can use Combustion to affect only the selected faces.

1. In 3ds Max, select the object you want to paint.
2. In the Modify panel, apply an Edit Mesh modifier to the object. (Choose Edit Mesh from the Modifier drop-down list.)

   If you are working with an editable mesh object, or a patch or NURBS surface, skip step 2. For geometry primitives, an option is to convert the object to a mesh, patch, or NURBS surface before step 3. However, you then lose the ability to adjust object parameters (for example, the radius of a sphere, the height of a box).
3. Choose Face as the sub-object selection level. Select the faces on which you want to paint.

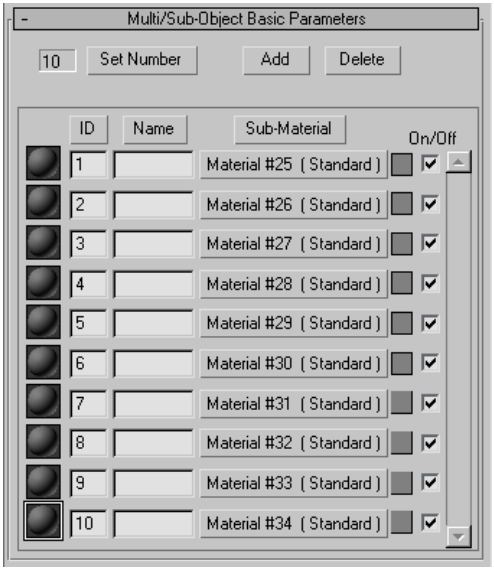


4. Drag a material from a sample slot in the Material Editor onto the selected faces.
5. In the modifier stack display, choose the object again, to disable sub-object selection.
6. In the Material Editor, use Pick Material From Object to grab the material from the geometry.



   You now have a new Multi/Sub-Object material. The original material appears as a sub-material applied to the selected faces.
7. In the Multi/Sub-Object material, go to the material assigned to the faces you want to paint.

   A Multi/Sub-Object material is simply a container for multiple sub-materials assigned to different faces of the same object. Click a Sub-Material button to go to a sub-material.



8. Assign a Combustion map to the Diffuse component of the sub-material applied to the selected faces.
9. Click Edit to launch Combustion.
10. Use the tools in Combustion to modify the material.

To modify a Combustion map:

1. In the Material Editor, select the material you want to modify.

   Material maps created in Combustion are vector-based and fully modifiable.
2. In the Combustion Parameters rollout, click the Edit button.

   The workspace corresponding to the Combustion map opens in Combustion. As you modify the workspace in Combustion the map is updated in 3ds Max.
3. In Combustion, save the workspace before you disable the Edit button in 3ds Max.

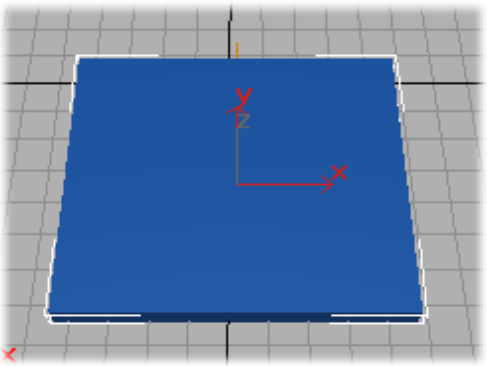
To create a displacement map:

In 3ds Max, the *Displace modifier (page 1–629)* acts as a force field to push and reshape an object's geometry. You can apply its variable force directly from the modifier gizmo, from a bitmap image, or from a Combustion workspace.

The grayscale component of the image is used to generate the displacement. Lighter colors in the image push outward more strongly than darker colors, resulting in a 3D displacement of the geometry.

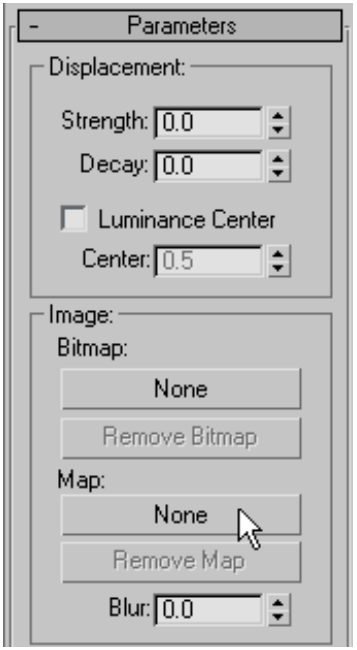
1. In 3ds Max, select the object to which you want to apply the displacement map.

   In this example, the displacement is applied to a box primitive.



   In the object's Parameters rollout, increase the number of Length and Width Segments. The closer the number of segments approaches the resolution of the displacement map, the more accurate is the result.

   In the example, 150 by 150 gives good results.
2. Apply a Displace modifier: in the Modify panel, choose Displace from the Modifier drop-down list.
3. In the Parameters rollout, Image group, click the Map button.

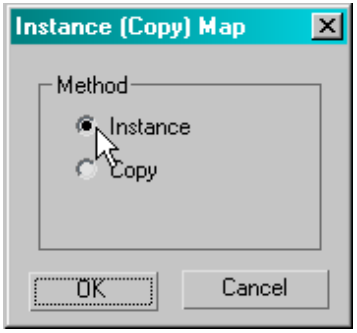


4. The Material/Map Browser appears. Select Combustion and click OK.

   The Map button now reads Map #1 (Combustion).

5. Open the Material Editor, and then click and drag the Map #1 (Combustion) button to an unused sample slot in the Material Editor.

   An Instance (Copy) Map dialog is displayed.



6. Select Instance and click OK.

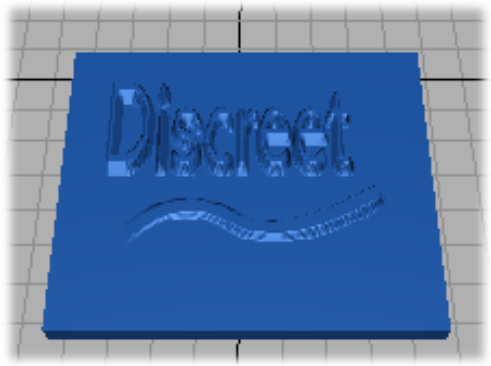
7. In the Material Editor, Combustion Parameters rollout, click Edit.

   This launches Combustion. In the New dialog, set the Type To Paint, and create a grayscale image to use as a displacement map. For more information, see the *Combustion User's Guide*.



8. In 3ds Max, increase the Displacement strength in the modifier Parameters rollout.

   As you increase the strength, you can see the result of the displacement map on the selected object.



9. In Combustion, save your project, then in 3ds Max, disable Edit in the Combustion Parameters rollout to exit Edit mode.

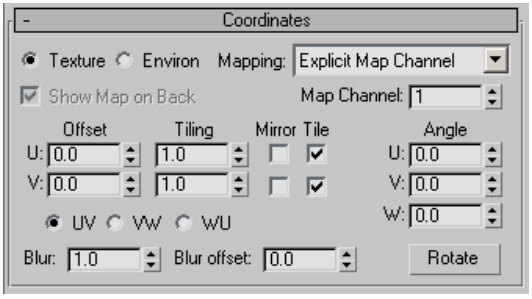
Interface

2D Mapping Coordinates

Like any 2D map in 3ds Max, mapping coordinates control how a Combustion map is positioned on objects.

For geometric primitives, mapping coordinates are usually provided automatically. For some kinds of geometry, such as *meshes (page 1–996)*, *patches (page 1–993)*, and *NURBS surfaces (page 1–1078)*, you must apply a *UVW Map modifier (page 1–922)* to provide mapping coordinates.

Controls in a *2D map's Coordinates rollout (page 2–1625)* affect how the map is positioned.

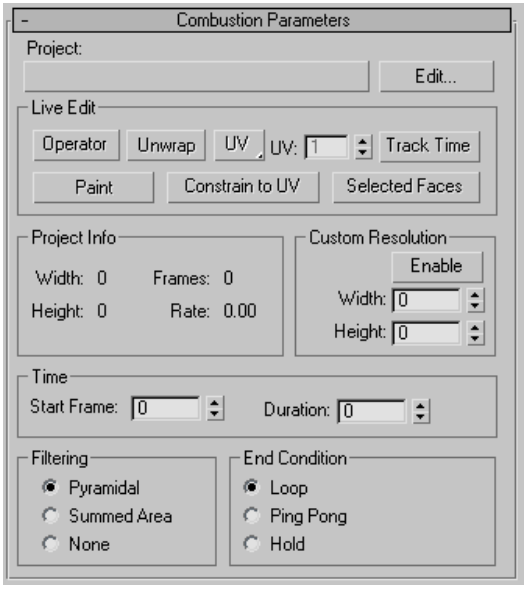


When you work with a Combustion map, these are the important points to remember:

- When you apply a Combustion map to an object, leave mapping set to the default values of Texture and Explicit Map Channel.
- When you use a Combustion map as an environment map, set mapping to Environ and then choose the mapping shape from the Mapping drop-down list.
- The offset, tiling, mirror, and angle controls are useful especially when the size of the projected Combustion map is smaller than the geometry.
- You can choose between UV, VW, and WU projections. (You can also do this from the Combustion Parameters rollout, as described below.) UV projects onto the surface of geometry like a slide projector. VW and WU project the map at right angles to the geometry. With a Combustion map, UV is almost always the most useful choice.

Combustion Parameters rollout

The Combustion Parameters rollout appears when you assign a Combustion map to a material.



Project—Loads the file to use in Combustion. You can load only file types supported by Autodesk Combustion, such as Combustion workspace files (*cws)*, or footage and image file formats supported by Combustion (see the *Combustion User's Guide* for information on supported footage formats).

Edit—Launches Combustion from the 3ds Max Material Editor. If a project is loaded, it is opened in Combustion. If no project is loaded, Combustion displays the New dialog. This dialog lets you specify a project type, name, video format, duration, and background color.

Live Edit group

These controls affect how you use Combustion with 3ds Max.

Operator—Switches control to Combustion, where you can select an operator. The results of the operator appear as the image in the Combustion map. The operator does not have to be the last operator in the pipe.

While Combustion is active, you can also adjust the operator. The Combustion map updates to show the results.

Unwrap Selected—Takes the current UVW mapping coordinates of the currently selected 3D object (or the current Face sub-object selection), and displays them in Combustion. This can help you coordinate the map and the mesh as you paint. The Unwrap display is only an overlay. It is displayed in Combustion but is not a part of the composite or the map.

UV List—Changes the mapping coordinate system (the direction in which the map is projected) from UV to VW or UW. UV projects onto the surface of geometry like a slide projector. VW and WU project the map at right angles to the geometry. With a Combustion map, UV is almost always the most useful choice.

UV—Specifies which mapping channel to unwrap and paint. Range=1 to 99.

Track Time—Links the Timeline in Combustion to the time slider in 3ds Max. When Track Time is enabled, you can use the time slider in 3ds Max to navigate between frames in Combustion.

Note: This control is not bidirectional; changing the frame in Combustion does *not* change the frame in 3ds Max.

Paint—When enabled, displays a paint cursor in 3ds Max viewports. You can then paint directly on the 3D geometry. Dragging the cursor in the viewport over the geometry in 3ds Max creates paint strokes inside Combustion.

Constrain To UV—When enabled, constrains paint strokes to remain within the edges of the UV mapping coordinates. When paint strokes are unconstrained on an object such as a box, they can jump to the other side of the map when you cross a map's edge. This can give erratic results. To prevent this, enable Constrain To UV.

In general, use the Constrain To UV option when you paint on boxes and other objects with planar maps. Disable this option when you want to paint on spherical maps or anywhere else the mapping has a singularity (where the edges of the map converge to a single point).

Selected Faces—Constrains the Combustion image to only the faces selected. This gives additional control or masking based on faces rather than UV mapping.

Project Info group

These readouts display the format of the Combustion Paint or composite operator. They are active when a Combustion workspace is loaded or Edit mode is active.

Width—Sets horizontal resolution of the frame in pixels.

Height—Sets vertical resolution of the frame in pixels.

Frames—Sets number of frames in the Combustion workspace.

Rate—Sets playback speed in frames per second.

Custom Resolution group

With these controls, you can customize the resolution of the Combustion map.

Enable—Enables the Width and Height controls.

Width and Height—Width changes the horizontal resolution of the map. Height changes the vertical resolution of the map.

Time group

These controls relate frames in the Combustion workspace to frames in the Combustion map. See the controls under "End Condition Group" for how to handle the map when it contains fewer frames than the 3ds Max scene.

Start Frame—Determines which frame of the Combustion sequence is used as the first frame of the Combustion map in 3ds Max.

Duration—Sets how many frames of the Combustion file sequence are used by the Combustion map in 3ds Max.

Filtering group

These controls determine the method for calculating *antialiasing (page 3–907)*:

Pyramidal—Sets the default antialiasing method. This method is faster than Summed Area filtering.

Summed Area—Implements a better method of antialiasing. Summed Area filtering uses more memory than Pyramidal. If it has to use virtual memory, it can dramatically increase rendering time.

None—Performs no antialiasing. This option takes the least time to render, but yields the lowest quality results.

End Condition group

These controls define what the 3ds Max renderer should do when the duration of the Combustion project (or the range of frames used in the Combustion map) is shorter than the rendering sequence in 3ds Max.

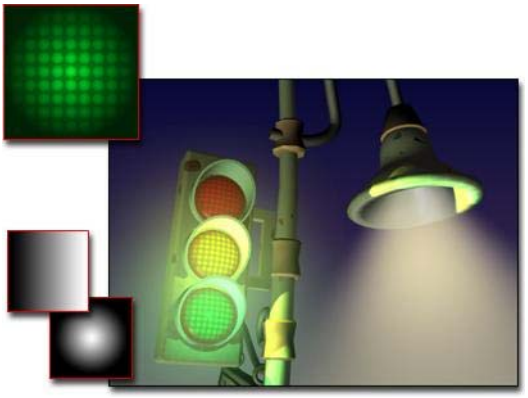
Loop—Plays the Combustion project animation repeatedly until the rendering sequence ends.

Ping Pong—Plays the animation forward, then backward, and repeatedly plays forward and backward until the rendering sequence is completed.

Hold—Plays the animation once, then repeatedly displays the last frame of the project until the rendering sequence is completed.

Gradient Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Gradient



Gradient maps used for the stoplight lamps, and for the background of the scene

Gradients shade from one color to another. You specify two or three colors for the gradient; the software interpolates intermediate values. Gradient maps are 2D maps.

Tip: You can swap colors by dragging one color swatch over another, then clicking Swap in the Copy or Swap Colors dialog. To reverse the overall direction of the gradient, swap the first and third colors.

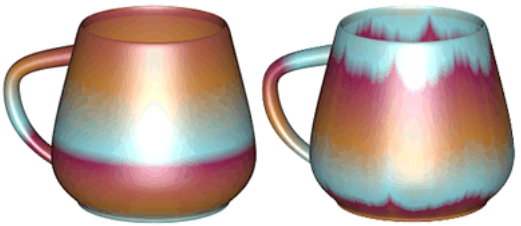
See also

For functionality shared with other 2D maps, see:

Coordinates Rollout (2D) (page 2–1625)

Noise Rollout (2D) (page 2–1630)

Output Rollout (page 2–1621)



Gradient-mapped material tiled and with noise

Procedures

To create a Gradient map:

1. Click a map button to assign a map.
2. Choose Gradient in the Material/Map Browser, and then click OK.

To change a gradient color:

1. In the Gradient Parameters rollout, click a color swatch to display the *Color Selector (page 1–161)*.
2. Adjust the color.
3. Click one of the other color swatches.
4. Adjust the color.

To choose the kind of gradient:

- Choose either Linear or Radial.

  A linear gradient shades from one color to another along a line. A radial gradient has one color on the inside and another on the outside, shading in a circular pattern.

To use a map for a color in the gradient:

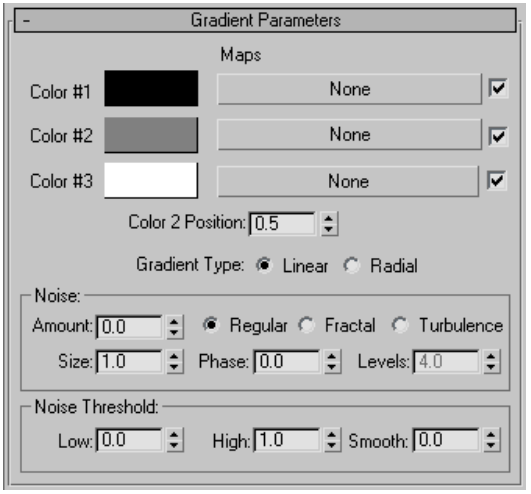
- In the Gradient Parameters rollout, click a map button to assign a map to a color.

To adjust the position of the second color:

- Change the Color 2 Position value.

  At the default value of 0.5, the second color is between the first and third colors. For a linear gradient, the second color's position ranges from the bottom at 0.0 to the top at 1.0. For a radial gradient, the second color's position ranges from the inside at 0.0 to the outside at 1.0.

Interface



Color #1-3—Sets the three colors that the gradient interpolates between. Displays the *Color Selector (page 1–161)*. You can drag and drop the colors from one swatch to another.

Maps—Displays a *map (page 3–968)* instead of the color. Maps are blended into the gradient in the same way that the gradient colors are blended. You can add nested procedural gradients in each window to make 5-, 7-, 9-color gradients, or more.

The check boxes enable or disable their associated maps.

Color 2 Position—Controls the center point of the middle color. The position ranges from 0 to 1. When it is 0, color 2 replaces color 3. When it is 1, color 2 replaces color 1.

Gradient Type—Linear interpolates the color based on the vertical position (V coordinate) while radial interpolates based on the distance from the center of the map (center is: U=0.5,V=0.5). With

both of these, you can rotate the gradient using the angle parameter under Coordinates, which is animatable.

Noise group

Amount—When nonzero (ranges from 0 to 1), applies a noise effect. This perturbs the color interpolation parameter using a 3D noise function based on U, V, and Phase. For example, a given pixel is halfway between the first and second color (the interpolation parameter is 0.5). If noise is added, the interpolation parameter would be perturbed by some amount so that it may become less or more than 0.5.

Regular—Generates plain noise. This is the same as Fractal noise with the Levels setting at 1. When the noise type is set to Regular, the Levels spinner becomes disabled (because Regular is not a fractal function).

Fractal—Generates noise using a fractal algorithm. The Levels option sets the number of iterations for the fractal noise.

Turbulence—Generates fractal noise with an absolute value function applied to it to make fault lines. The noise amount must be greater than 0 to see any effects of turbulence.

Size—Scales the noise function. Smaller values give smaller chunks of noise.

Phase—Controls the speed of the animation of the noise function. A 3D noise function is used for the noise. The first two parameters are U and V and the third is phase.

Levels—Sets the number of fractal iterations or turbulence (as a continuous function).

Noise Threshold group

When the noise value is above the Low threshold and below the High threshold, the dynamic range is stretched to fill 0–1. This produces a smaller discontinuity at the threshold transition and thus causes less potential aliasing.

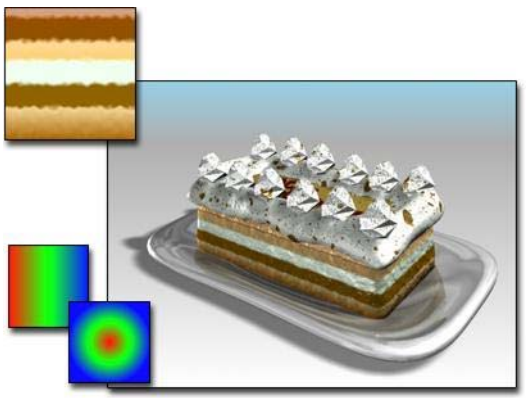
Low—Sets the low threshold.

High—Sets the high threshold.

Smooth—Helps make a smoother transition from the threshold value to the noise value. When smooth is 0, no smoothing is applied. When it is 1, the maximum amount of smoothing is applied.

Gradient Ramp Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Gradient Ramp



Gradient ramp used for the layers of the cake

Gradient Ramp is a 2D map similar to the Gradient map. It shades from one color to another. In this map, however, you can specify any number of colors or maps for the gradient. There are a variety of controls, making highly customized gradients possible. Almost any parameter of Gradient Ramp can be animated.

See also

For functionality shared with other 2D maps, see:

Coordinates Rollout (2D) (page 2–1625)

Noise Rollout (2D) (page 2–1630)

Output Rollout (page 2–1621)

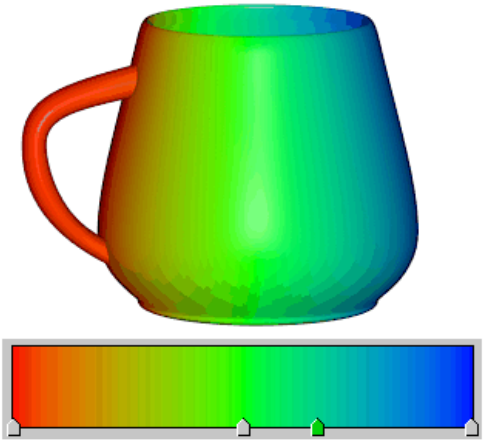
Procedure

To create a material with a Gradient Ramp map:

1. Open the Material Editor, and choose an unused sample slot.
2. Close the Basic Parameters rollout, and open the Maps rollout.
3. Click the Map button for Diffuse to display the Material/Map Browser.
4. In the Browser list, click Gradient Ramp. The map appears in the upper-left of the Browser window. Click OK.

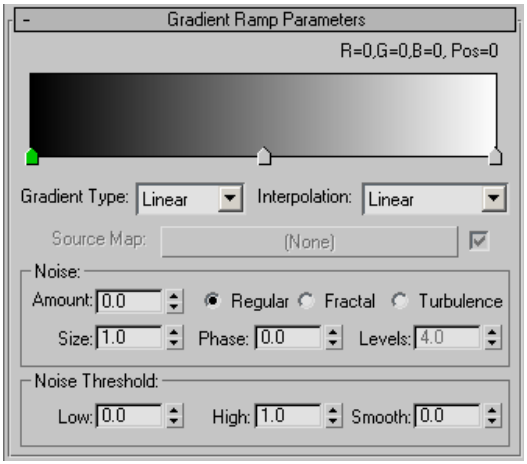
   Gradient Ramp map is applied to the sample slot as the Browser closes.
5. In the Material Editor, under its toolbar, give the material a name to identify its use in your scene.

   When you replace an existing map with a gradient ramp map, choosing to Keep Old Map as Submap in the Replace Map dialog, the old map becomes assigned to the first flag as a texture.



Gradient Ramp-mapped material with a colored gradient

Interface



Gradient bar shows default gradient and interpolation types

Gradient bar—Presents an editable representation of the gradient being created. The effect of the gradient moves from left (start point) to right (end point).

By default, three flags appear along the bottom edge of a red/green/blue gradient. Each flag controls a color (or map). The currently selected flag is green, and its RGB value and its position in the gradient (in the range 0 to 100) appear above the gradient bar. Each gradient can have any number of flags.

The gradient bar has the following features:

- Click anywhere along the bottom edge to create additional flags.
- Drag any flag to adjust the position of its color (or map) within the gradient. The start and end flags (Flag #1 at 0 and Flag #2 at 100) cannot be moved. However, other flags can occupy these positions and still be moved.
- More than one flag can occupy a given position. If two flags are at the same position, a slight edge appears between the colors. With three or more flags at the same position, the edge is a hard line.

Right-click options for gradient bar—Right-click in the gradient bar to display a menu with these options:

Reset—Returns gradient bar to defaults.

Load Gradient—Loads an existing gradient (*.dgr* file) into the gradient bar.

Save Gradient—Loads your current gradient bar as a *.dgr* file.

Copy, Paste—Copies a gradient and pastes it into another Gradient Ramp map.

Load UV Map—Selects a UV map.

Load Bitmap—Selects a bitmap.

Flag Mode—Toggles flag display.

Right-click options for flags—Right-click any flag to display a menu with the following options:

Copy and Paste—Lets you copy the current key and paste it to replace another key. The other key could be in another Gradient Ramp as well as the current one.

Edit Properties—Choose this option to display the *Flag Properties dialog (page 2–1655)*.

Delete—Deletes the flag.

Gradient Type—Chooses the type of gradient. The following Gradient types are available. These affect the entire gradient.

- **4 Corner—**An asymmetrical linear transition of colors.
- **Box—**A box.
- **Diagonal—**A linear diagonal transition of colors.
- **Lighting—**Based on the light intensity value. No light=far left; brightest light=far right.
- **Linear—**A smooth, linear transition of colors.
- **Mapped—**Lets you assign a map to use as the gradient. Enables the Source Map controls for specifying the map and turning it on and off.
- **Normal—**Based on the angle between the vector from the camera to the object and the surface normal vector at the sample point. The leftmost flag of the gradient is 0 degrees; the rightmost flag is 90 degrees.
- **Pong—**A diagonal sweep that repeats in the middle.
- **Radial—**A radial transition of colors.
- **Spiral—**A smooth, circular transition of colors.
- **Sweep—**A linear sweep transition of colors.
- **Tartan—**A plaid.

Interpolation—Chooses the type of interpolation. The following Interpolation types are available. These affect the entire gradient.

Note: Gradients are ordered from left to right. The "next" flag is to the right of the current flag; the "previous" flag is to the left.

Custom—Sets an individual interpolation type for each flag. Right-click the flag to display the *Flag Properties dialog (page 2–1655)* and set the interpolation.

Ease In—Weighted more toward the next flag than the current flag.

Ease In Out—Weighted more toward the current flag than the next flag.

Ease Out—Weighted more toward the previous flag than the next flag.

Linear—Constant from one flag to the next. (Default.)

Solid—No interpolation. Transitions are a sharp line.

Source Map—Click to assign a map to a mapped gradient. The check box turns the map on or off.

The Source Map controls are available only when Mapped is the chosen gradient type.

Noise group

Amount—When nonzero, a random noise effect is applied to the gradient, based on the interaction of the gradient ramp colors (and maps, if present). The higher this value, the greater the effect. Range=0 to 1.

Regular—Generates plain noise. Basically the same as fractal noise with levels disabled (because Regular is not a fractal function).

Fractal—Generates noise using a fractal algorithm. The Levels option sets the number of iterations for the fractal noise.

Turbulence—Generates fractal noise with an absolute value function applied to it to make fault lines. Note that the noise amount must be greater than 0 to see any effects of turbulence.

Size—Sets the scale of the noise function. Smaller values give smaller chunks of noise.

Phase—Controls the speed of the animation of the noise function. A 3D noise function is used for the noise; the first two parameters are U and V and the third is phase.

Levels—Sets the number of fractal iterations or turbulence (as a continuous function).

Noise Threshold group

When the noise value is above the Low threshold and below the High threshold, the dynamic range is stretched to fill 0 to 1. This causes a smaller discontinuity at the threshold transition and produces less potential aliasing.

High—Sets the high threshold.

Low—Sets the low threshold.

Smooth—Helps make a smoother transition from the threshold value to the noise value. When Smooth is 0, no smoothing is applied. When Smooth is 1, the maximum amount of smoothing is applied.

Flag Properties Dialog

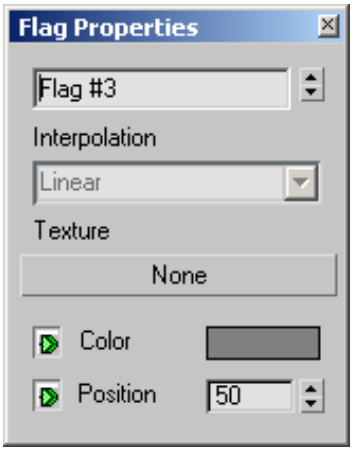
Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Gradient Ramp > Gradient Ramp Parameters rollout > Right-click a flag at the bottom of the gradient bar. > Edit Properties > Flag Properties dialog

By setting parameters on the Flag Properties dialog, you can customize the effect of each flag on the Gradient Ramp map. You have access to all flags from this dialog.

See also

Gradient Ramp Map (page 2–1652)

Interface



Name field—Select any flag with the spinners. The selected flag turns green on the gradient bar. Use this field to rename a flag to represent its use in the gradient.

Interpolation—Disabled unless the Interpolation type for the Gradient Ramp map is set to Custom. The available interpolation types for flags are

similar to the corresponding ones in the Gradient Ramp map:

Note: Gradients are ordered from left to right. The "next" flag is to the right of the current flag; the "previous" flag is to the left.

Ease In—Weighted more toward the next flag than the current flag.

Ease In Out—Weighted more toward the current flag than the next flag.

Ease Out—Weighted more toward the previous flag than the next flag.

Linear—Constant from one flag to the next. (Default.)

Solid—No interpolation. Transitions are a sharp line.

Texture—Assigns a map in place of a color. When unselected, the flag turns blue to indicate a map assignment.

Color—Click the color swatch to change the color controlled by the selected flag.

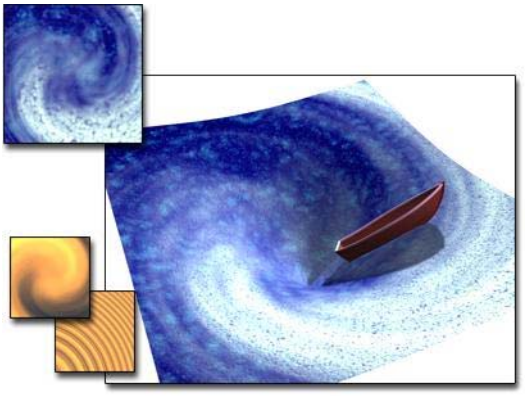
Position—Shows the current position of the selected flag. Use the spinners to reposition the flag, or enter a value directly. Position is not available for the start and end flags, because these flags can't be moved.



Animation keys—Animation keys are created by default for Color and Position, and the keys are active, indicated by the green triangles next to the Color and Position labels. You can turn these keys off if you don't intend to animate the gradient.

Swirl Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Swirl



Swirl used to create the whirlpool

Swirl is a 2D procedural map that generates patterns similar to the swirls in two-flavor ice creams. Like other two-color maps, either color can be replaced with other maps, so it's possible to swirl marble with wood, for example.

See also

For functionality shared with other 2D maps, see:

Coordinates Rollout (2D) (page 2–1625)

Noise Rollout (2D) (page 2–1630)

Procedure

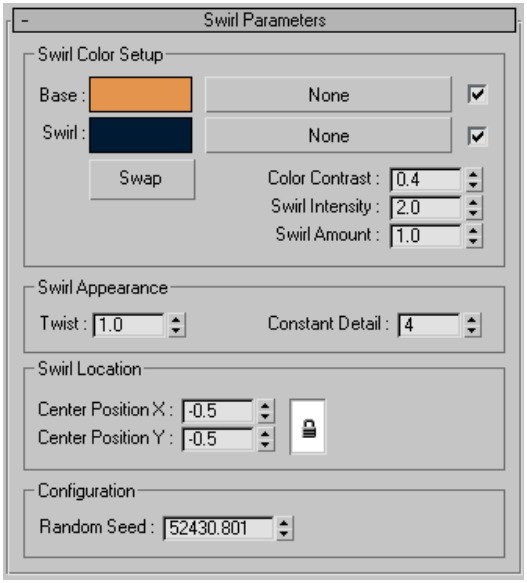
To create a Swirl-mapped material:

1. Open the Material Editor, and choose an unused sample slot.
2. Close the Basic Parameters rollout, and open the Maps rollout.
3. Turn on Diffuse Color. Click its Map button to display the Material/Map Browser.
4. In the Browser list, click Swirl. The map appears in the upper-left of the Browser window. Click OK.

Swirl map is applied to the sample slot as the Material/Map broswer closes.

5. In the Material Editor, under its toolbar, give the material a name to identify its use in your scene.

Interface



Swirl Color Setup group

Base—The underlying layer for the swirl effect. Click the color swatch to change. Click None to assign a map in place of a color. The check box enables or disables the map.

Swirl—Mixed with the Base color or map, produces the swirl effect. Click the color swatch to change this color. Click None to assign a map in place of a color.

Swap—Reverses the color or map assignments for Base and Swirl.

Color Contrast—Controls the contrast between Base and Swirl. At 0, the swirl is blurred. Higher values increase the contrast until all colors become black and white, even if Swirl Intensity and Swirl Amount are very high. Range=0 to 4.0; Default=0.4.

Swirl Intensity—Controls the intensity of the swirl color. Higher values create a more vibrant mix of colors. At 0, the swirl effect disappears. Range=-10 to 10.0; Default=2.0.

Swirl Amount—Controls the quantity of the Swirl color that gets mixed into the Base color. If set to 0, only the base color is used. Range=0 to 3.0; Default=1.0.

Swirl Appearance group

Twist—Changes the number of spirals in the swirl effect. Higher values increase the number of spirals. Negative values change the direction of the twist. At 0, the colors are randomly distributed, not swirled. Range=-20.0 to 20.0; Default=1.0.

Constant Detail—Changes the level of detail within a swirl. Lower values minimize the level of detail within the swirl. At 0, all detail is lost. Higher values increase detail until the swirl effect disappears. Values are in whole numbers. Range=0 to 10; Default=4.

Swirl Location group

Center Position X and Y—Adjusts the location of the swirl's center on the object.

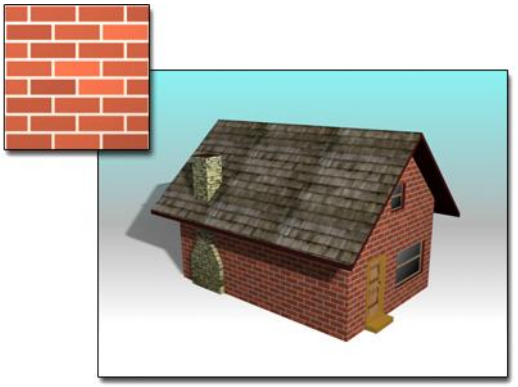
Lock—X and Y values remain identical as you adjust them. By turning off Lock and adjusting either the X or Y position, you can "slide" the swirl effect across the object. Default=on.

Configuration group

Random Seed—Sets a new starting point for the swirl effect. Changes the swirl pattern while maintaining other parameters. Range=0 to 65,535; No default.

Tiles Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Tiles



Tiles used for the walls of a house

Using the Tiles procedural map, you can create brick or stacked tiling of colors or material mappings. There are commonly defined architectural brick patterns available, or you can design custom patterns. With the Tiles map, you can:

- Assign many of the maps available through the Material Editor.
- Load textures and use colors in the pattern.
- Control the number of tiles in columns and rows.
- Control the size of the grout gap and its roughness.
- Apply random variance in the pattern.
- Control the stacking layout by shifting how the tiles line up.

See also

For functionality shared with other 2D maps, see:

Coordinates Rollout (2D) (page 2–1625)

Noise Rollout (2D) (page 2–1630)

Procedures

Example: To create a brick wall:

1. Create a wall using a Box primitive, or use an existing surface in one of your scenes.
2.  Open the *Material Editor (page 2–1409)*. Select an unused sample slot.
3. Click the Maps rollout to open it. Click the Map button for Diffuse to display the Material/Map Browser.
4. In the map list, select Tiles, then click OK.

   The Tiles map is now assigned to the sample slot. On the Material Editor, you see new rollouts appear for this map.
5. On the Standard Controls rollout, use Preset Type to select the type of tiles for the wall. Stack Bond is the default.
6.  On the Material Editor toolbar, click Assign Material To Selection to apply the tile map to the wall. Then click Show Map In Viewport to see the applied map.
7. Open the Advanced Controls rollout. Under Tiles Setup, adjust Horizontal and Vertical Count. The default is eight rows high, with three repeats of the pattern in each row. Visually scale the size of the tiles to your scene. Also adjust Texture as well as Color and Fade Variance to fine-tune the appearance of the tiles.
8. Under Grout Setup, adjust parameters for the texture of the grout, gap spacing between tiles, and roughness of the grout. You can also create missing bricks in the map by setting % Holes to a value above 0.
9. Under Miscellaneous, you can vary the color of the tiles by using the Random Seed option.

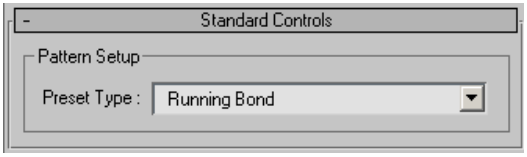
Example: To match the tiles on the top and side of a wall:

1. Select a wall mapped with tile.

2. On the Modify panel, choose Editable Mesh.
3. Turn on Sub-Object > Face.
4. Select the top face of the wall.
5. Open the Material Editor. Drag the wall's tile material to an unused sample slot to duplicate the material.
6. Under Tiles Setup, adjust the horizontal and vertical count of the new material to match the side of the wall.
7. Apply the new tile material to the selected faces on the top of the wall.
8. Under Stacking Layout, align the tiles by using the Line Shift option.

Interface

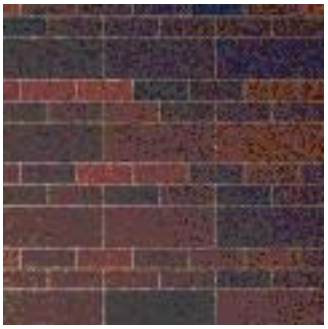
Standard Controls rollout



Preset Type—Lists the commonly defined architectural tile bonds, or patterns, plus a custom pattern, which you design by selecting options under the Advanced Controls and Stacking Layout rollouts. The following illustrations show some of the different bonds:



Common Flemish



Fine Running



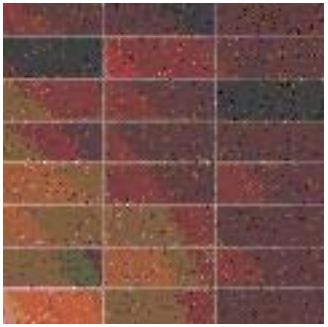
Fine Stack



1/2 Running

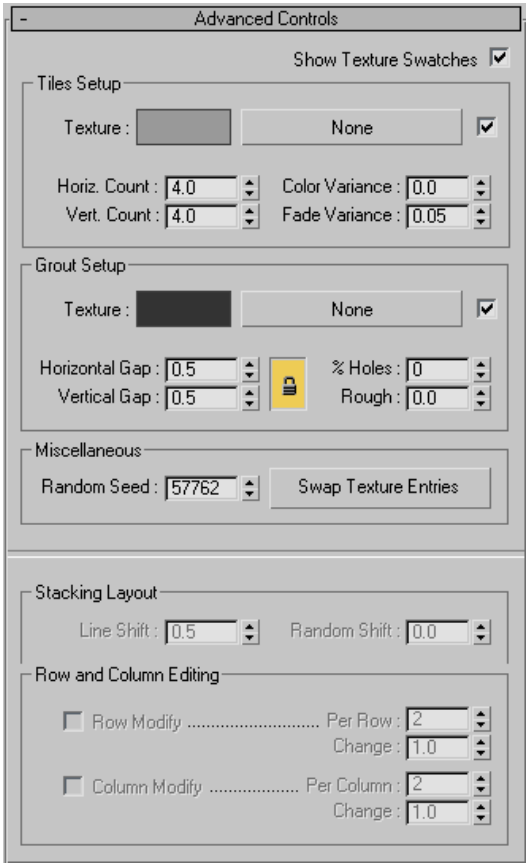


Running



Stack

Advanced Controls rollout



Show Texture Swatches—Updates to show the texture assigned by a map for Tiles or Grout.

Tiles Setup group

Texture—Controls the display of the current texture map for the tiles. When on, the texture is used as the tile pattern instead of the color swatch. When turned off, the color of the tiles is displayed; clicking the color swatch displays the *Color Selector (page 1–161)*.

None—Acts as a target where you drag and drop maps for the tiles. When you click this button with a map assigned, the software displays the rollout for the map. You return this button to

None (removing the assigned map) by dragging and dropping a None map from the Map/Material Browser.

Horiz. Count—Controls the number of tiles in a row.

Vert. Count—Controls the number of tiles in a column.

Color Variance—Controls the color variation among the tiles.

Fade Variance—Controls the fading variation among the tiles.

Grout Setup group

Texture—Controls the display of the current texture map for the grout. When on, the texture is used as the grout pattern instead of the color swatch. When off, the color of the grout is displayed, and clicking the color swatch displays the *Color Selector (page 1–161)*.

None—Acts as a target where you drag and drop maps for the grout. When you click this button with a map assigned, the software displays the rollout for the map. You return this button to None (removing the assigned map) by dragging and dropping a None map from the Map/Material Browser.

Horizontal Gap—Controls the horizontal size of the grout between the tiles. This value is locked by default to the vertical gap, so that both values change as you edit one or the other. To unlock them, click the lock icon.



Vertical Gap—Controls the vertical size of the grout between the tiles. This value is locked by default to the horizontal gap, so that both values change as you edit one or the other. To unlock them, click the lock icon.

% Holes—Sets the percentage of holes in the tiled surface caused by missing tiles. The grout shows through the holes.

Rough—Controls the roughness of the edges of the grout.

Miscellaneous group

Random Seed—Randomly applies patterns of color variation to the tiles. Does not require any other setting to generate completely different patterns.

Swap Texture Entries—Swaps the texture maps or colors between the tiles and the grout.

Stacking Layout group

Note: This group of controls is active only when Custom Tiles is selected in Standard Controls rollout > Pattern Setup > Preset Type.

Line Shift—Shifts every second row of tiles a distance of one unit.

Random Shift—Randomly shifts all rows of tiles a distance of one unit.

Row and Column Editing group

Note: This group of controls is enabled only when Custom Tiles is selected in Standard Controls rollout > Pattern Setup > Preset Type.

Row Modify—When on, creates a custom pattern for rows, based on the values of Per Row and Change. Default=off.

- **Per Row**—Specifies which rows to change. When Per Row equals 0, no rows change. When Per Row equals 1, every row changes. When Per Row is a value greater than 1, the change appears every *N* rows: a value of 2 changes every second row, a value of three changes every third row, and so on. Default=2.
- **Change**—Changes the width of tiles in the affected rows. A value of 1.0 is the default tile width. Values greater than 1.0 increase the

width of tiles, and values less than 1.0 decrease it. Range=0.0 to 5.0. Default=1.0.

A value of 0.0 is a special case: When the Change value is 0.0, no tiles appears in that row, and the underlying material shows through.

Column Modify—When on, creates a custom pattern for columns, based on the values of Per Column and Change. Default=off.

- **Per Column—**Specifies which columns to change. When Per Column equals 0, no columns change. When Per Column equals 1, every column changes. When Per Column is a value greater than 1, the change appears every *N* columns: a value of 2 changes every second column, a value of three changes every third column and so on. Default=2.
- **Change—**Changes the height of tiles in the affected columns. A value of 1.0 is the default tile height. Values greater than 1.0 increase the height of tiles, and values less than 1.0 decrease it. Range=0.0 to 5.0. Default=1.0.

  A value of 0.0 is a special case: When the Change value is 0.0, no tile appears in that column, and the underlying material shows through.

3D

3D maps are patterns generated procedurally in three dimensions. For example, Marble has a grain that goes through the assigned geometry. If you cut away part of an object with marble assigned as its texture, the grain in the cutaway portion matches the grain on the object's exterior. The following 3D maps are available in 3ds Max:

Cellular (page 2–1664): Generates a cellular pattern that's useful for a variety of visual effects, including mosaic tiling, pebbled surfaces, and ocean surfaces.

Dent (page 2–1667): Generates three-dimensional bumps over a surface.

Falloff (page 2–1670): Generates a value from white to black based on the angular falloff of the face normals on the surface of the geometry. The Falloff map provides greater flexibility when creating opacity falloff effects. Other effects include Shadow/Light, Distance Blend, and Fresnel.

Marble (page 2–1673): Simulates the grain of marble with two explicit colors and a third intermediate color.

Noise (page 2–1674): Noise is a turbulence pattern in three dimensions. Like Checker in 2D, it is based on two colors, either of which can be mapped.

Particle Age (page 2–1675): Alters the color (or map) of a particle based on the particle's life.

Particle MBlur (page 2–1676): (MBlur is short for Motion Blur.) Alters the opacity of the leading and trailing ends of particles based on their rate of movement.

Perlin Marble (page 2–1677): An alternative, procedural marble map with a turbulence pattern.

Planet (page 2–1678): Simulates the contours of a planet as seen from space.

Smoke (page 2–1679): Generates fractal-based turbulence patterns to simulate the effects of smoke in a beam of light, or other cloudy, flowing mapping effects.

Speckle (page 2–1680): Generates a speckled surface for creating patterned surfaces that can simulate granite and similar materials.

Splat (page 2–1681): Generates a fractal pattern similar to splattered paint.

Stucco (page 2–1682): Generates a fractal pattern similar to stucco.

Waves (page 2–1683): Creates watery or wavy effects by generating a number of spherical wave centers and randomly distributing them.

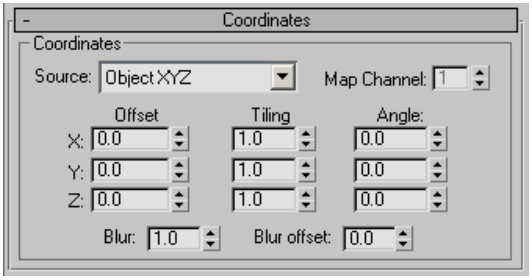
Wood (page 2–1684): Creates a 3D wood grain pattern.

Coordinates Rollout (3D)

Material Editor > Select sample slot. > Get Material > Material/Map Browser > Turn on 3D Maps. > Double-click a map type to apply it to sample slot. > Coordinates rollout is displayed in Material Editor.

By adjusting coordinate parameters, you can move a map relative to the volume of the object to which it is applied.

Interface



Source—Chooses the coordinate system to use. There are four options:

- **Object XYZ—**Uses the object's local coordinate system.
- **World XYZ—**Uses the scene's world coordinate system.
- **Explicit Map Channel—**Activates the Map Channel field. You can choose any channel from 1 to 99.
- **Vertex Color Channel—**Assigns vertex colors as a channel. See *Editable Mesh (page 1–996)* for details on assigning vertex color.

See also *Vertex Color Map (page 2–1693)*.

When one of the map channels is set, it locks the map into position on the vertices of the object so that the map "sticks" to the object as it deforms during animation.

When an object is deforming through its own local space (for example, when it is bending or twisting), the object appears to move through the map, because it passes through the XYZ coordinates of the 3D texture.

Map Channel—Unavailable unless the source is Explicit Map Channel. When available, you can choose any channel from 1 to 99.

Offset—Moves the map pattern along the specified axis.

Tiling—*Tiles (page 3–1022)* the map pattern along the specified axis and makes the pattern narrower.

Angle—Rotates the map pattern along the specified axis.

Blur—Affects the sharpness or blurriness of the map based on its distance from the view. The farther away the map is, the greater the blurring. The Blur value blurs maps in world space. Blur is primarily used to avoid *aliasing (page 3–907)*.

Blur Offset—Affects the sharpness or blurriness of the map without regard to its distance from the view. Blur Offset blurs the image itself in object space. Use when you want to soften or defocus the details in a map to achieve the effect of a blurred image.

Cellular Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Cellular



Cellular maps create the goblet textures.

The Cellular procedural map generates a pattern that's useful for a variety of visual effects, including mosaic tiling, pebbled surfaces, and even ocean surfaces.

Tip: The Material Editor sample slot doesn't show the cellular effect very clearly. For a better visual aid to getting the effect you want, assign the map to geometry and render the scene.

See also

For functionality shared with other 3D maps, see the following topics:

Coordinates Rollout (3D) (page 2–1663)

Output Rollout (page 2–1621)

Procedures

Example: To create confetti:

1. Assign the Cellular map as a Diffuse map.
2. Set the parameters as follows:

   Cell Color (use RGB values after clicking color swatch): **202, 75, 171**

   Variation: **55**

   Division Colors (use RGB values after clicking the color swatch):

   - First: **127, 150, 197**
   - Second: **0, 119, 163**

   Cell Characteristics:

   - Circular
   - Size: **2.4**
   - Spread: **0.43**
   - Fractal: off

   Thresholds:

   - Low: **0.19**
   - Mid: **0.65**
   - High: **0.86**

Example: To create alien skin:

1. Assign the Cellular map as a Diffuse map.
2. Set the parameters as follows:

   Cell Color (use RGB values after clicking the color swatch): **52, 107, 58**

   Variation: 0.0

   Division Colors (use RGB values after clicking the color swatch):

   - First: **112, 119, 64**
   - Second: **143, 137, 112**

   Cell Characteristics:

   - Circular
   - Size: **7.9**
   - Spread: **0.61**
   - Fractal: on
   - Iterations: **2.0**

   Thresholds:

   - Low: **0.17**
   - Mid: **0.64**

- High: **1.0**

3. Copy the Diffuse map to the Bump map.

   If you increase the Bump amount, increase the Cellular bump map's Bump Smoothing value as well.

Example: To create a tile mosaic:

1. Assign the Cellular map as a Diffuse map.
2. Set the parameters as follows:

   Cell Color (use RGB values after clicking the color swatch): **141, 120, 87**

   Variation: **54**

   Division Colors (use RGB values after clicking the color swatch):

   - First: **128, 128, 128**
   - Second: **221, 221, 221**

   Cell Characteristics:

   - Chips
   - Size: **7.0**
   - Spread: **0.35**
   - Fractal: off

   Thresholds:

   - Low: **0.42**
   - Mid: **0.76**
   - High: **1.0**

3. Assign a Mix map as the Bump map.
4. Click Material/Map Navigator to display the Navigator. Copy the Cellular Diffuse map by dragging it from the Navigator to the Color #2 map window of the Mix map.
5. A dialog is displayed. You are asked if this should be an instance or a copy. Select Copy and click OK.
6. Assign a Noise map to the Color #1 map window of the Mix map.
7. Set the Noise parameters as follows:

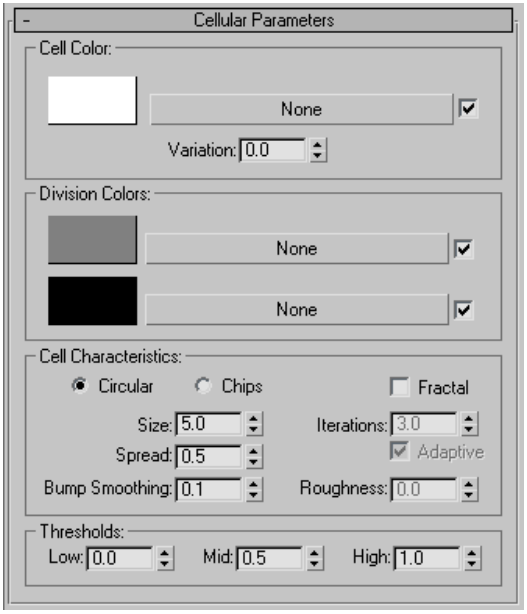
   Noise Type: Fractal

   Levels: **6.0**

   Size: **9.3**

8. In the Mix Parameters rollout of the Mix map, set the Mix Amount to **0.5**.
9. Go to top level of the material. In the Maps rollout, set Bump Amount to **82**.

Interface



Cell Color group

These controls specify the color of the cells.

Color swatch—Displays the *Color Selector (page 1–161)*. Choose a color for the cells.

Map button—Assigns a map to the cells, rather than a solid color.

Check box—When on, enables the map. When off, disables the map (cell color reverts to the color swatch).

Variation—Varies the color of the cells by randomly altering RGB values. The higher the variation, the

greater the random effect. This percentage value can range from 0 to 100. At 0, the color swatch or the map completely determines the cell color. Default=0.

Division Colors group

These controls specify the color of the divisions between cells. Cell divisions are a ramp between two colors or two maps.

Color swatches—Display the Color Selector for choosing a cell division color.

Map buttons—Assigns a map to one of the cell division colors.

Check boxes—When on, enables the associated map. When off, disables the associated map (the division color reverts to the color swatch).

Cell Characteristics group

These controls change the shape and size of the cells.

Circular / Chips—With Circular, the cells are circular. This gives a more organic, or bubbly look. With Chips, the cells have linear edges. This gives a more chipped or mosaic appearance. Default=Circular.

Size—Alters the overall scale of the map. Adjust this value to fit the map to your geometry. Default=5.0.

Spread—Alters the size of individual cells. Default=0.5.

Bump Smoothing—When you use a cellular map as a *bump map (page 2–1506)*, you might encounter aliasing or jagginess at the boundaries of the cells. If this occurs, increase this value. Default=0.1.

Fractal—Defines the cellular pattern as a fractal, thus enabling the three following additional parameters. Default=off.

Iterations—Sets the number of times the fractal function is applied. Caution: Increasing this value increases rendering time. Default=3.0.

Adaptive—When on, the number of fractal iterations is set adaptively. That is, the number of iterations increases the closer the geometry is to the scene's point of view, and decreases in the distance. This reduces aliasing and also saves time while rendering. Default=on.

Roughness—When you use the Cellular map as a *bump map (page 2–1506)*, this parameter controls how rough the bumps are. When Roughness is zero, each iteration is half the strength of the previous iteration, and half the size. As Roughness increases, each iteration is closer in strength and size to the previous iteration. When Roughness is at its maximum value of 1.0, each iteration is the same size and strength as the previous. In effect, this turns off the fractalization. Roughness has no effect unless Iterations is greater than 1.0. Default=0.0.

Thresholds group

These controls affect the relative size of cells and divisions. They are expressed as normalized percentages (0 to 1) of the sizes specified by the default algorithm.

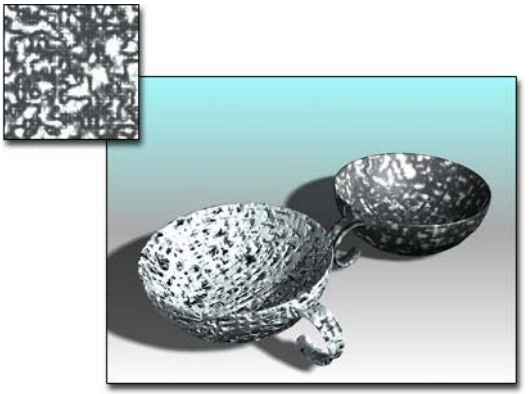
Low—Adjusts the size of the cells. Default=0.0.

Mid—Adjusts the size of the first division color, relative to the second. Default=0.5.

High—Adjusts the overall size of divisions. Default=1.0.

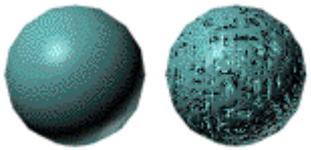
Dent Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Dent



Dent map gives texture to the cup on the left; cup on the right has same pattern, but without dents.

Dent is a 3D procedural map. During scanline rendering, Dent creates a random pattern based on fractal noise. The effect of the pattern depends on the map type.



Effect of default parameters

Defaults: Iterations=2, Size=200, Strength=20

Dent was designed to be used primarily as a bump map, and its default parameters are optimized for this usage. As a bump map, Dent renders three-dimensional dents over the surface of an object. Editable parameters control the size, depth, and complexity of the denting effect.

Dent can also be used with other maps. With a diffuse color map, for example, the two colors assigned to Dent mix in random swirls over the surface of the object. Either of the colors can be replaced by other maps.



Dent bump maps

At left, default parameters produce fairly uniform dents over surface.

At right, increased Strength parameter creates a deeper and more irregular pitting effect.



Dent diffuse maps

At left, Dent as a two-color diffuse map.

At right, colors replaced with Dent and Marble maps.

See also

For functionality shared with other 3D maps, see the following topic:

Coordinates Rollout (3D) (page 2–1663)

Procedures

To make a Dent map:

1. In the Material Editor, click a sample slot to make it active.
2. Open the Maps rollout.
3. Click Bump or other map button to display the Material/Map Browser.
4. Double-click Dent in the list of map types.

   The Material Editor displays the Coordinates and Dent Parameters rollouts.
5. Set parameters on the Dent Parameters rollout.

   The active sample slot updates to show the Dent effect.

To replace a color:

1. Click a color swatch labeled Color #1 or Color #2.
2. In the standard *Color Selector (page 1–161)*, choose a replacement color.

   The color updates in the color box and sample slot.

To swap a color:

- Click Swap.

  The position of the two colors is reversed in the color boxes and sample slot.

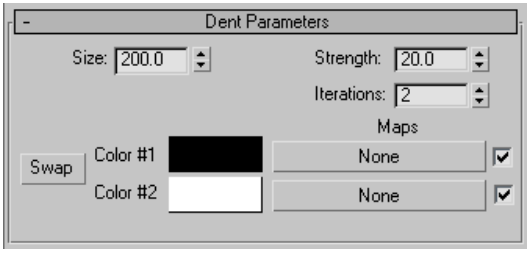
To replace a color with a map:

1. Click a Map bar marked None next to one of the color swatches.

   The Material/Map Browser is displayed.
2. Select a map from the list.

   The sample slot updates to show the map in place of the color.

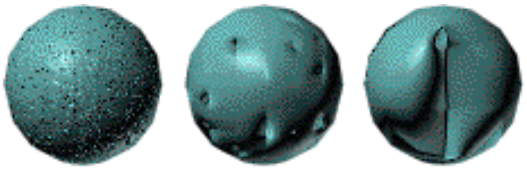
Interface



Size—Sets the relative size of dents. As the size increases, the number of dents decreases when other settings are the same. Default=200.

Decreasing Size creates the appearance of tiny dents spaced fairly evenly. The effect can resemble a "sand-covered" surface.

Increasing Size creates the appearance of distinct grooves and gouges on a surface. The effect sometimes has a "hardened lava" look.



Size=10, 500, and 1000

Iterations=1, Strength=20 (default)



Size=10, 500, and 1000

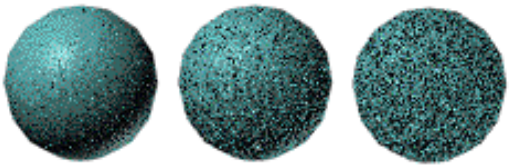
Iterations=3, Strength=20 (default)

Each set of three spheres uses the same size range, but varies the number of iterations. Strength is held constant in both sets.

Strength—Determines the relative coverage of the two colors. Higher values increase the coverage of Color #2, while lower values increase the coverage of Color #1. Default=20.

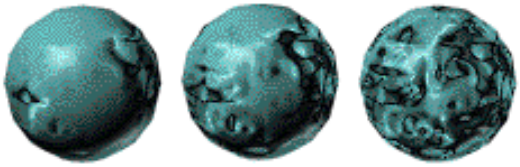
When using Dent as a bump map, increasing the Strength value typically makes the dents look deeper.

In the following illustrations, each set of three spheres uses the same Strength range, but varies the Size value between the two sets. The Iterations value is the same in both sets.



Size=10, Iterations=3

Strength=5, 20 (default), and 100



Size=1000, Iterations=3
Strength=5, 20 (default), and 100

Iterations—Sets the number of calculations used to create the dents. Default=2.

Dent is based on a fractal-noise equation. During rendering, a dented surface is calculated one or more times to produce the finished effect. Each calculation pass is an iteration.

As a surface is calculated, each iteration adds to the number of dents and the complexity and randomness of the final surface (dents become dented).

The Dent texture requires heavy calculation, especially at higher iterations. This can slow down rendering time considerably.



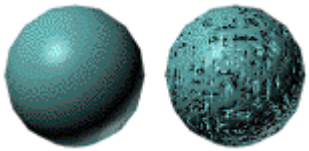
Iterations=1, 3, and 6
Size=500, Strength=20

The three spheres have uniform settings for size and strength. Only the number of iterations varies.

Swap—Reverses the position of colors or maps.

Colors—Allows choice of two colors where appropriate for a color component (such as Diffuse). Defaults=black for Color #1 and white for Color #2.

Dent can create patterns in an object's color as well as its surface. By using Dent as a diffuse color map, the entire surface is affected.



Size=500, Strength=60
Iterations=2

Two colors are mixed to produce a random pattern, governed by size, strength, and iteration settings. The default colors are black and white, but either can be replaced or swapped.

Left sphere: Sets the color to black and white.

Middle sphere: Replaces white with red.

Right sphere: Swap black and red.

By adjusting Strength, Size, and Iterations parameters, you vary the Dent patterns on a diffuse color map.

- Size sets the density of the dent pattern. At low settings, the pattern is dense. As Size increases with other settings held constant, the pattern becomes increasingly sparse.
- Strength sets the color strength in the dent pattern. At low settings, Color #1 (black) dominates the pattern. As Strength increases, Color #2 (white) replaces Color #1.
- Iterations sets the color iterations in the dent pattern. At low settings, Color #1 is dominant. As iterations increase, Color #2 gradually increases in the pattern.

Dent is applied as a diffuse map in the following examples. Colors are default black and white.



Size=100, 500, and 1000
Strength=20, Iterations=2
Color #1, #2=black, white (defaults)



Strength=20, 50, and 100

Size=500, Iterations=2



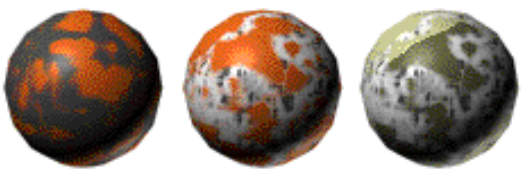
Iterations=2, 4, and 8

Size=500, Strength=20

Maps—Replaces colors with maps in the dent pattern. The check boxes enable or disable their associated map.

You can assign maps to one or both of the Dent color slots. Any kind of map can be used, including Dent. The map overrides the assigned color, which has no effect.

In the three spheres below, the assigned colors are progressively replaced with maps. Parameters of the original Dent map are the same for all spheres.



Size=500, Strength=60, Iterations=2

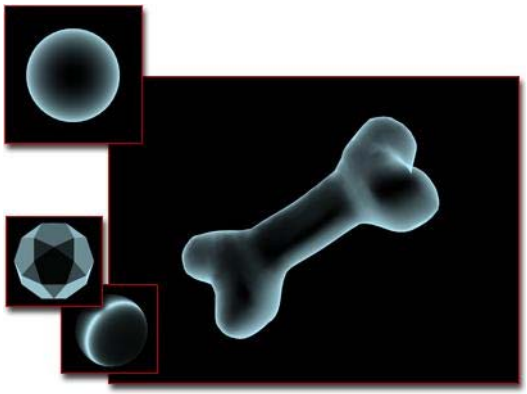
Left sphere: Applies Dent as a diffuse map. Color #1 is black; Color #2 is red.

Middle sphere: Replaces black with Dent map (all defaults).

Right sphere: Replaces red with Marble map (all defaults).

Falloff Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Falloff



Falloff map creates the appearance of translucency.

The Falloff map generates a value from white to black, based on the angular falloff of the face normals on the surface of the geometry. The direction used to specify the angular falloff varies, depending on the methods you choose. However, with the default settings, the map generates white on faces whose normals point outward from the current view, and black on faces whose normals are parallel to the current view.

Falloff map provides a greater variety of opacity falloff effects than those in the Falloff settings in a standard material's *Extended Parameters rollout (page 2–1471)*. You assign the Falloff map as an *opacity map (page 2–1503)*. However, you can also use Falloff for special effects, such as an iridescent look.

Note: When old files that use Falloff maps are brought into 3ds Max, the old Falloff interface is displayed, replacing the new Falloff interface.

See also

For functionality shared with other 3D maps, see *Output Rollout (page 2–1621)*

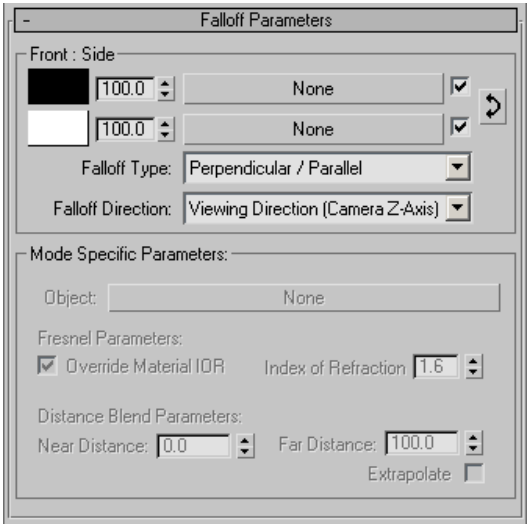
Procedure

To control opacity using a Falloff map:

1. Assign the Falloff map as an opacity map.
2. Render to see the effect.
3. Adjust the falloff parameters to vary the effect.

Interface

Falloff Parameters rollout



Front : Side—By default, "Front : Side" is the name of the group at the top of this rollout. Front : Side indicates Perpendicular/Parallel falloff. This name changes depending on the falloff type selected. In all cases, the name on the left refers to the top set of controls, and the name on the right to the bottom set.

The controls are as follows:

- Click the color swatches to assign colors.
- Use the numeric fields and spinners to adjust the relative strength of the colors.
- Click the buttons marked None to assign maps.
- Turn on the check boxes to activate the maps; otherwise the colors are used. These are on by default.
-  Click Swap Colors/Maps (the curved arrow) to exchange the assignments.

Falloff Type—Chooses the kind of falloff. Five options are available:

- **Perpendicular/Parallel**—Sets the angular falloff ranges between face normals that are perpendicular to the falloff direction and normals that are parallel to the falloff direction. The falloff range is based on a 90-degree change in face normal direction. (Default.)
- **Towards/Away**—Sets the angular falloff ranges between face normals that face toward (parallel to) the falloff direction and normals that face away from the falloff direction. The falloff range is based on a 180-degree change in face normal direction.
- **Fresnel**—Based on adjustments to the Index of Refraction (IOR). Results in dim reflections on surfaces facing the view, with much brighter reflections on angled faces, creating highlights like those on the sides of a glass.
- **Shadow/Light**—Adjusts between two subtextures based on how much light is falling on the object.
- **Distance Blend**—Adjusts between two subtextures based on Near Distance and Far Distance values. Uses include reducing aliasing on large terrain objects and controlling the shading in non-photorealistic environments.

Falloff Direction—Chooses the direction of falloff. Five options are available:

- **Viewing Direction (Camera Z-Axis)**—Sets the falloff direction relative to the camera (or screen). Changing object orientation doesn't affect the falloff map. (Default.)
- **Camera X/Y Axis**—Similar to Camera Z-Axis. For example, using Camera X-Axis with the Toward/Away falloff type runs the gradient from left (Toward) to right (Away).

- **Object**—Uses an object whose position determines the falloff direction. Click the wide button next to Object in the Mode Specific Parameters group, and then pick an object in the scene. The falloff direction is the direction from the point being shaded toward the object's center. Points on the side toward the object center get the Towards value, and those away from the object get the Away value.
- **Local X/Y/Z Axis**—Sets the falloff direction to one of the object's local axes. Changing the orientation of the object changes the falloff direction.
- **World X/Y/Z Axis**—Sets the falloff direction to one of the world coordinate system axes. Changing object orientation doesn't affect the falloff map.

When no object is chosen, the falloff direction uses the local X, Y, or Z axis of the object being shaded.

Mode Specific Parameters group

The first parameter applies and is available only when you set *Falloff Direction* to Object:

Object—Picks object from scene and puts its name on the button.

The following are parameters for the *Fresnel* falloff type:

Override Material IOR—Allows change to the Index of Refraction set by the material.

Index of Refraction—Sets a new Index of Refraction. This option is unavailable unless Override Material IOR is turned on.

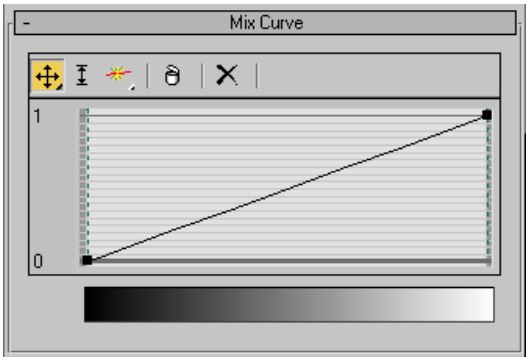
The following are parameters for the *Distance Blend* falloff type:

Near Distance—Sets the distance at which the blend effect begins.

Far Distance—Sets the distance at which the blend effect ends.

Extrapolate—When on, the effect continues beyond the Near and Far distances.

Mix Curve rollout



Using the graph on the Mix Curve rollout, you can precisely control the gradient produced by any falloff type. You see the resulting gradient in the bar below the graph.

Move flyout

 Moves a selected point in any direction, limited by the unselected points on either side.

 Constrains movement to the horizontal.

 Constrains movement to the vertical.

 **Scale Point**—Scales the selected point within the range of its gradient. On a Bezier corner point, this control is effectively the same as a vertical move. On a Bezier smooth point, you can scale the point itself or either handle. As with the move controls, scale is limited by the unselected points on either side.

Add Point flyout

 Adds a Bezier corner point anywhere on the graph line. The point makes a sharp angle when moved.

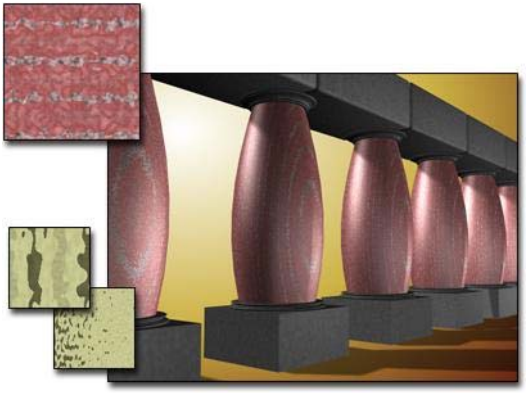
 Adds a Bezier smooth point anywhere on the graph line. Handles attached to the point create smooth curves when moved. On a Bezier smooth point, you can move the point or either handle.

 **Delete Point—**Removes selected points.

 **Reset Curves—**Returns graph to its default, a straight line between 0 and 1.

Marble Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Marble



Marble map used for the balusters

The Marble map produces a marbled surface with colored veins against a colored background. A third color is automatically generated.

Note: Another way to create marble is to use the *Perlin Marble Map (page 2–1677)*.

See also

For functionality shared with other 3D maps, see *Coordinates Rollout (3D) (page 2–1663)*.

Procedures

To create a Marble map:

1. Click a map button to assign a map.
2. Choose Marble in the Material/Map Browser, and then click OK.

To adjust the size of Marble veins:

- In the Marble Parameters rollout, change the Size value to change the overall scale of the vein pattern, and change Vein Width to change the width of veins relative to the overall scale.

  The larger the Size value, the wider the veins. The larger the Vein Width value, the more veins appear relative to the overall pattern.

To change vein color:

1. In the Marble Parameters rollout, click a color swatch to display the *Color Selector (page 1–161)*.
2. Adjust the color.
3. Click Close to dismiss the dialog.

To use a map for a vein:

- In the Marble Parameters rollout, click a map button to assign a map to a color.

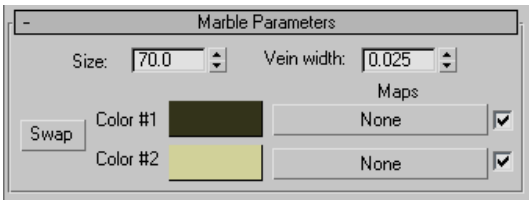
To swap the two vein colors:

- In the Marble Parameters rollout, click Swap.

To adjust mapping coordinates:

- In the Coordinates rollout, adjust Offset, Tiling, or Angle.

Interface



Size—Sets the spacing between the veins.

Vein Width—Sets the width of the veins.

Swap—Switches the position of the two colors or maps.

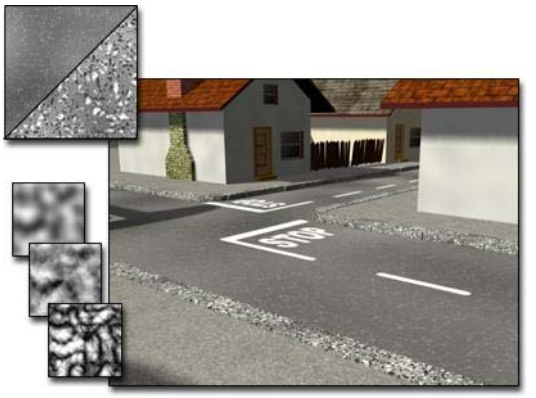
Color # 1 and Color # 2—Displays the *Color Selector (page 1–161)*. Select one color for the veins (color 1) and another for the background (color 2). A third color is generated from the two colors you select.

Maps—Selects the bitmaps or *procedural maps (page 3–997)* to appear in the veins or in the background color.

Turn on the check boxes to make the maps active.

Noise Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Noise



Noise map used for the edges of the street

The Noise map creates random perturbation of a surface based on the interaction of two colors or materials.

See also

For functionality shared with other 3D maps, see the following topics:

Coordinates Rollout (3D) (page 2–1663)

Output Rollout (page 2–1621)

Procedures

To change a component color:

1. In the Noise Parameters rollout, click a color swatch to display the *Color Selector (page 1–161)*.
2. Adjust the color.
3. Click OK to dismiss the dialog.

To use a map for a component:

- In the Noise Parameters rollout, click a map button to assign a map to a color.

To swap the two component colors:

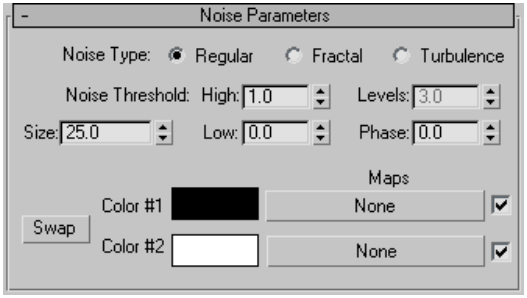
- In the Noise Parameters rollout, click Swap.

To adjust mapping coordinates:

- In the Coordinates rollout, adjust Offset, Tiling, or Angle.

Interface

Note: For the Noise map, the tiling and mirroring controls are disabled in the Texture Tiling And Output rollout.

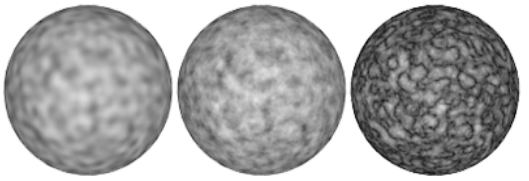


Noise Type—

- **Regular—**(The default.) Generates plain noise. Basically the same as fractal noise with the Levels setting at 1. When the noise type is set to

Regular, the Levels spinner is inactive (because Regular is not a fractal function).

- **Fractal**—Generates noise using a fractal algorithm. The Levels option sets the number of iterations for the fractal noise.
- **Turbulence**—Generates fractal noise with an absolute value function applied to it to make fault lines.



Regular, Fractal, Turbulence

Size—Sets the scale of the noise function, in 3ds Max units. Default=25.0.

Noise Threshold—When the noise value is above the Low threshold and below the High threshold, the dynamic range is stretched to fill 0 to 1. This creates a smaller discontinuity (technically, 1st order instead of 0 order) at the threshold transition and produces less potential *aliasing (page 3–907)*.

- **High**—Sets the high threshold. Default=1.0.
- **Low**—Sets the low threshold. Default=0.0.

Levels—Determines how much fractal energy is used for the Fractal and Turbulence noise functions. You can set the exact amount of turbulence you want, and also animate the number of fractal levels. Default=3.0.

Phase—Controls the speed of the animation of the noise function. Use this option to animate the noise function. Default=0.0.

Swap—Switches the position of the two colors or maps.

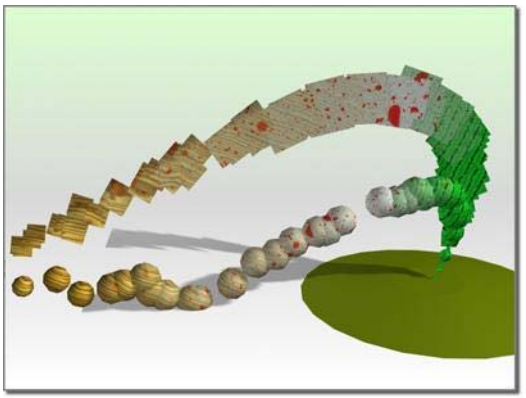
Color # 1 and Color # 2—Display the *Color Selector (page 1–161)* so you can choose one or the other of the two principal noise colors. Intermediate color values are generated from the two colors you select.

Maps—Select the bitmaps or procedural maps to appear in one or the other noise color.

Turn on the check boxes to make the maps active.

Particle Age Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Particle Age



Particle age changes the appearance of particles over time.

The Particle Age map is for use with *particle systems (page 3–990)*. Typically you assign the Particle Age map as a *Diffuse map (page 2–1498)*, or in Particle Flow with the *Material Dynamic operator (page 2–191)*. It alters the color (or map) of a particle based on the particle's life. The particles in a system begin as one color. At a specified age, they begin changing (by interpolation) to a second color, and then they change again to a third color before they die out.

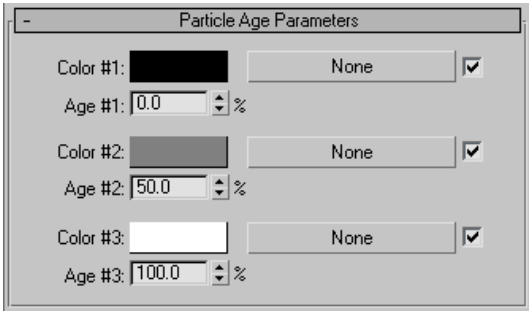
Tip: This map works well with the *Particle MBlur map (page 2–1676)*. For example, you could assign Particle Age as a *diffuse map (page 2–1498)* and Particle MBlur as an *Opacity map (page 2–1503)*. Another way to use Particle Age is in the mask channel of a *Blend (page 2–1588)* material. In this case, you could set up two of the colors to white, and one to black, which would make the particles change materials over their age. In addition, you could set up one of the two materials with an *effects*

channel (page 3–946) and use a *Glow (page 3–355)* filter on it through Video Post.

See also

For functionality shared with other 3D maps, see *Output Rollout (page 2–1621)*.

Interface



Color #1—Sets the color of a particle at its birth.

Click the button to the right of the color swatch to assign a map instead of a solid color. The check box turns the map on or off.

Age #1—Sets the age where a particle starts changing from Color #1 to Color #2, expressed as a percentage of the particle's entire life.

Color #2—Sets the color of a particle in mid-life. You can also assign a map to this color.

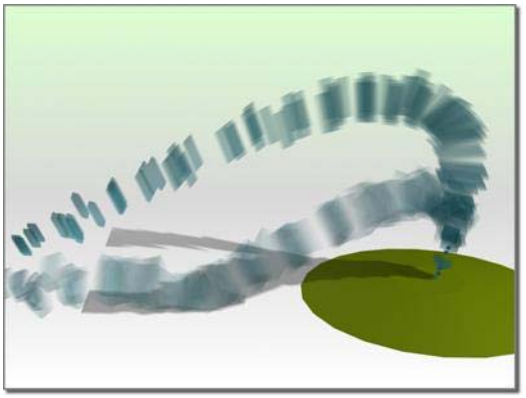
Age #2—Sets the age where a particle's color equals Color #2, expressed as a percentage of the particle's entire life.

Color #3—Sets the color of a particle at its death. You can also assign a map to this color.

Age #3—Sets the age where a particle changes to Color #3, expressed as a percentage of the particle's entire life.

Particle MBlur Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Particle MBlur



Particle MBlur makes particles blur as they move.

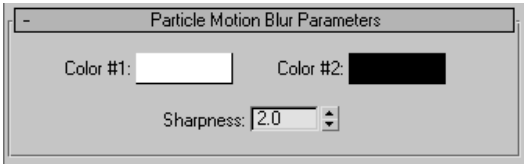
The Particle MBlur (Motion Blur) map is for use with *particle systems (page 3–990)*. The map alters the opacity of the leading and trailing ends of particles based on their rate of motion. The map is usually applied as an opacity map, but you can use it as a diffuse map for special effects.

The following conditions must be in effect to achieve particle motion blur:

- The Particle MBlur map must be in the same material that is assigned to the particles. For best results, it should be assigned as an opacity map
- The particle system must support the Particle MBlur map. Particle systems that support Particle MBlur include PArray, PCloud, Super Spray, and Spray.
- In the particle system's Particle Rotation rollout, in the Spin Axis Controls group, the Direction of Travel/MBlur option must be on.
- In this same group, the Stretch spinner must be greater than 0 to stretch the particles as a percent of their length based on the particle Speed setting.

- The correct type of particle must be used. MBlur works on all particle types except Constant, Facing, Metaparticles, and PArray Object Fragments. Also, in the Standard Particles category, MBlur does not support the Triangle and SixPoint particle types.
- The material assigned to the particle system must not be a Multi/Sub-Object material.

Interface



Color #1—A particle approaches this color as it reaches its slowest speed. By default, this color is white to provide the opaque end of the range for an opacity map.

Color #2—A particle approaches this color as it speeds up. As a default, this color is black to provide transparency in an opacity map.

Typically, you don't need to change either of these two colors.

Sharpness—Controls the transparency, relative to the speed. If Sharpness is set to 0, the entire particle is blurry and transparent, no matter how slow it is traveling. The default works well in many cases. Default=2.0.

Perlin Marble Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Perlin Marble



Perlin marble used for the texture of the goblet

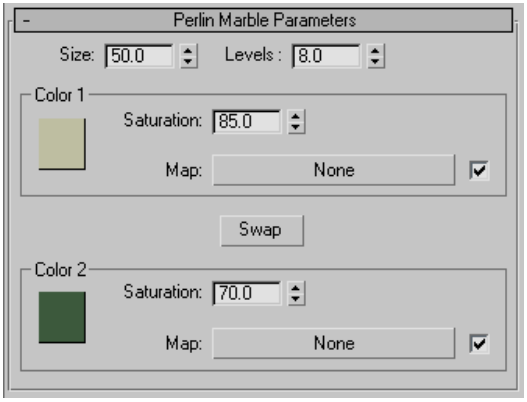
The Perlin Marble map generates a marble pattern using the Perline Turbulence algorithm. This map is an alternative to *Marble (page 2–1673)*, which is also a 3D material.

See also

For functionality shared with other 3D maps, see the following topic:

Coordinates Rollout (3D) (page 2–1663)

Interface



Size—Sets the size of the marble pattern. Change this to change the scale of marble, relative to the object's geometry. Default=50.

Levels—Sets the number of times the turbulence algorithm is applied. Can range from 1.0 to 10.0. The higher the value, the more complicated the marble pattern. Default=8.0.

Color 1 and Color 2 groups

The controls in these groups are identical. They determine the two main colors of the marble.

Color swatch—Click to display the *Color Selector (page 1–161)*. and change the color.

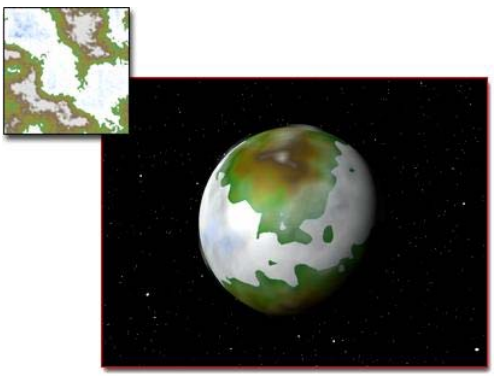
Map—Click to assign a map instead of a solid color. The check box turns the map on or off.

Saturation—Controls the saturation of the color in the map, without altering the color displayed in the color swatch. Lower values darken the color, and higher values lighten it. Range=1 to 100; Default=85 for Color 1, 70 for Color 2.

Swap—Click to swap Color 1 and Color 2.

Planet Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Planet



Planet map used to create another world

Planet is a 3D map that uses fractal math to simulate the colors on the surfaces of a planet. You can control the size of continents, percent of ocean coverage, and so on. This map is meant to be used as a diffuse map. It does not work well as a bump map.

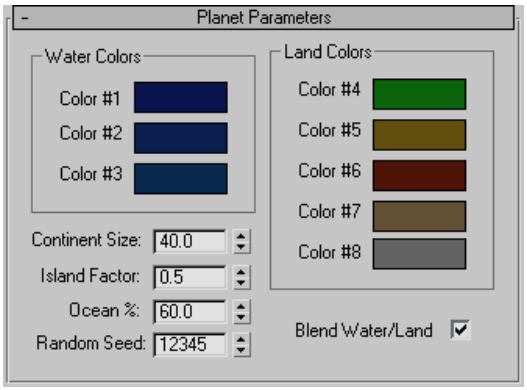
Tip: The Material Editor's sample slot doesn't show the planet effect very clearly. To help get the effect you want, double-click the sample slot to get a larger sample slot, or assign the map to geometry and render the scene. Another way to preview the planet map is to use the *Material Editor Options dialog (page 2–1436)* to set the 3D Map Sample Scale to equal a main dimension of the object you are applying the map to. For example, if you want to use the planet map on a sphere with a radius of 20, change the map scale from 100 (the default) to 20.

See also

For functionality shared with other 3D maps, see the following topic:

Coordinates Rollout (3D) (page 2–1663)

Interface



Continent Size—Sets the size of the fractal noise pattern used to generate the continents. The higher the value, the larger the continents. Default=40.

Island Factor—Sets the size of the fractal noise pattern used to generate islands and mountains. Can range from 0 to 100. At 0, the geography is very low. Higher settings create a more rugged landscape. Default=0.5.

Ocean %—Sets the percentage of the planet's surface that is covered by water. Default=60.0.

Random Seed—Sets the seed for pseudo-random generation of the pattern. Changing this number can change the pattern completely, even if other settings remain the same. On the other hand, a different Planet map with the same settings including the same Random Seed will appear the same.

Water Colors group

The colors in these three swatches are applied to the water areas of the planet surface.

Water colors—Click a swatch to display the *Color Selector (page 1–161)* and change the color. Color #1 is the "center" area of the water mass, Color #2 surrounds Color #1, and Color #3 surrounds Color #2, meeting the land.

Land Colors group

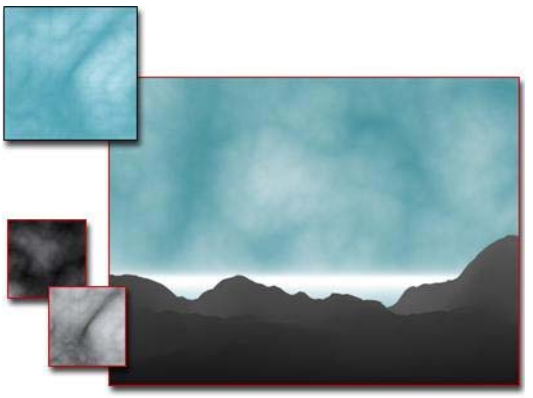
The colors in these five swatches are applied to the land areas of the planet surface. Their arrangement continues that of the water colors.

Land colors—Click a swatch to display the *Color Selector (page 1–161)* and change the color. Color #4 is the shoreline of the land, meeting the water; Color #5 comes next, working toward the center of the land mass. Color # 8 is at the center of the land mass.

Blend Water/Land—When on, the boundary between water and land is blended, giving a hazy appearance. When off, the boundary between water and land is sharp.

Smoke Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Smoke



Smoke map used to create clouds in the sky

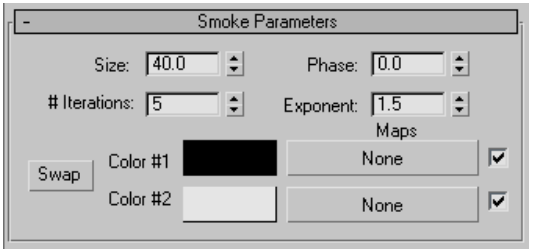
Smoke is a 3D map that generates amorphous, fractal-based turbulent patterns. It's primarily designed for animated *opacity mapping (page 2–1503)* to simulate the effects of smoke in a beam of light, or other cloudy, flowing mapping effects.

See also

For functionality shared with other 3D maps, see the following topic:

Coordinates Rollout (3D) (page 2–1663)

Interface



Size—Changes the scale of the smoke "clumps." Default=40.

 # Iterations—Sets the number of times the fractal function is applied. The higher the value, the more detail within the smoke, but the longer the calculation time. Default=5.

Phase—Shifts the turbulence within the smoke pattern. Animate this parameter to animate the movement of the smoke. Default=0.0.

Exponent—Makes color #2, representing the smoke, sharper and more wispy. As this value increases, the smoke "tendrils" become smaller within the pattern. Default=1.5.

Swap—Exchanges the colors.

Color #1—Represents the smokeless portion of the effect.

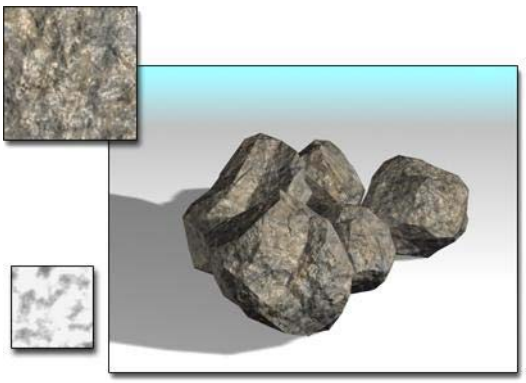
Color #2—Represents the smoke.

Because this map is usually used as an opacity map, you can adjust the luminance of the color values to alter the contrast of the smoke effect.

- Click a color swatch to change the color. Usually you only need to do this to adjust luminance.
- Click a map button to assign a map instead of a solid color. Turn on the check box to activate the map.

Speckle Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Speckle



Speckle map used for rocks

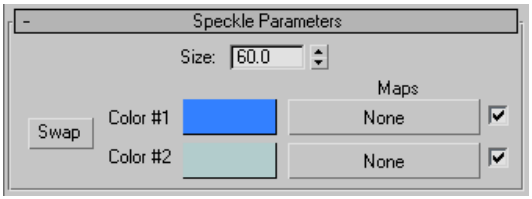
Speckle is a 3D map that generates a speckled surface pattern that's useful for *diffuse mapping (page 2–1498)* and *bump mapping (page 2–1506)* to create granite-like and other patterned surfaces.

See also

For functionality shared with other 3D maps, see the following topic:

Coordinates Rollout (3D) (page 2–1663)

Interface



Size—Adjusts the size of the speckles. Use this to make the speckles match your geometry. Default=60.

Swap—Exchanges the two color components.

Color #1—Represents the color of the speckles.

Color #2—Represents the color of the background.

Click one of the swatches to display the *Color Selector (page 1–161)* and change one of these component colors.

Maps—Click a button to assign a map that replaces one of the color components. Turning off the check box turns off the associated map (the Speckle map reverts to the associated color component).

Splat Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Splat



Splat map used for the patterns in ice cream

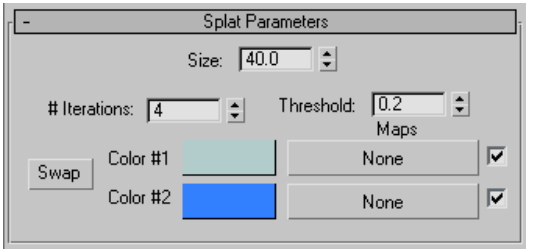
Splat is a 3D map that generates a fractal surface pattern that is useful as a *Diffuse map (page 2–1498)* for creating a pattern similar to splattered paint.

See also

For functionality shared with other 3D maps, see the following topic:

Coordinates Rollout (3D) (page 2–1663)

Interface



Size—Adjusts the size of the splats. Use this to make the splats match your geometry. Default=40.

 # Iterations—Sets the number of times the fractal function is evaluated. The higher the number, the more detailed the splats, but the longer the calculation time. Default=4.

Threshold—Determines how much of Color #1 is mixed with Color #2. At 0, only Color #1 is displayed; at 1, only Color #2 is displayed. Default=0.2.

Swap—Exchanges the two color components.

Color #1—Represents the color of the background.

Color #2—Represents the color of the splats.

Click one of the swatches to display the *Color Selector (page 1–161)* and change one of these colors.

Maps—Assigns a map to replace one of the color components. Turning off the check box turns off the associated map (the Splat map reverts to the associated color component).

Stucco Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Stucco



Stucco map used for a plaster wall

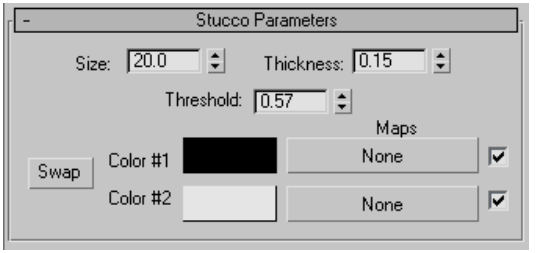
Stucco is a 3D map that generates a surface pattern that is useful for *bump mapping (page 2–1506)* for creating the effect of a stuccoed surface.

See also

For functionality shared with other 3D maps, see the following topic:

Coordinates Rollout (3D) (page 2–1663)

Interface



Size—Adjusts the size of the indentations. Use this to make the scale of the stucco match your geometry. Default=20.

Thickness—Blurs the border between the two colors. At 0, the borders are sharp. The higher the Thickness, the more the borders are blurred and the less distinct the indentations are. When you use Stucco as a bump map, the indentations are very faint at 0.5 and disappear at values not much greater. Default=0.15.

Threshold—Determines how much of Color #1 is mixed with Color #2. At 0, only Color #2 is displayed; at 1, only Color #1 is displayed. Default=0.57.

Swap—Exchanges the two color components.

Color #1—Represents the color of the indentations.

Color #2—Represents the background stucco color.

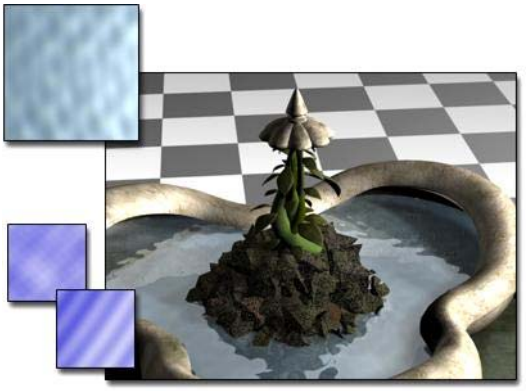
Click one of the swatches to display the *Color Selector (page 1–161)* and change one of these component colors.

Tip: Because the Stucco map is meant to be used as a bump map, usually you don't need to adjust the default colors.

Maps—Assigns a map to replace one of the color components. Turning off the check box turns off the associated map (the Stucco map reverts to the associated color component).

Waves Map

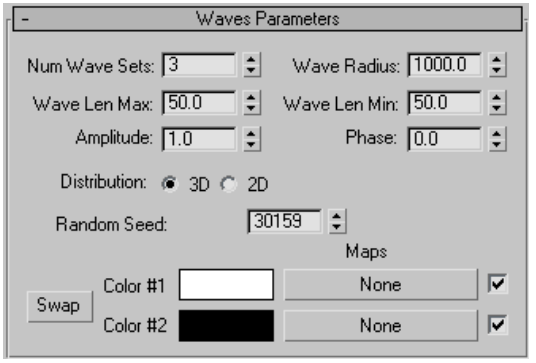
Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Waves



Waves map used for the pool in the fountain

Waves is a 3D map that creates watery or wavy effects. It generates a number of spherical wave centers and randomly distributes them over a sphere. You can control the number of wave sets, the amplitude, and the speed of the waves. This map works effectively as both a diffuse and bump map at the same time. It can also be useful in combination with an opacity map.

Interface



Num Wave Sets—Specifies how many wave sets are used in the pattern. Wave sets are groups of radially symmetrical waves that originate from randomly computed points along the surface of an imaginary sphere inside the object (a circle, in the case of 2D wave distribution). For calm water, set this to a low number. Use a high number for choppy water. Range= 1 to 50; Default=10.

Wave Radius—Specifies the radius, in 3ds Max units, of the imaginary sphere (3D distribution) or circle (2D distribution) whose surface is the origin of each wave set. A large radius produces large circular wave patterns, while a small radius produces dense, smaller waves. Default=800.

Wave Len Max and Wave Len Min—Define the interval used to randomly chose each wave center. If these two values are close together, the water appears more regular. If they're farther apart, the water is less regular. Default Max=50.0; Default Min=5.0.

Amplitude—Adjusts the strength and the depth of the waves by increasing the contrast between the two colors. Default=1.0.

Phase—Shifts the wave pattern. Animate this parameter to animate the motion of the pattern.

Distribution 3D/2D—3D distributes the wave centers on the surface of an imaginary sphere, affecting all sides of a 3D object. 2D distributes the wave in circles centered on the XY plane, which is more appropriate for flat water surfaces such as oceans and lakes.

Random Number Seed—Provides a seed number to generate the water pattern. The pattern changes with each seed, but all other settings are maintained.

Swap—Exchanges the colors.

Color #1 and #2—Click the color swatches to change the colors used in the pattern. Use one color for the wave troughs and the other for the wave peaks.

You can also click the map buttons to replace one or both colors with a mapped pattern. The check box enables or disables the map.

Wood Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Wood

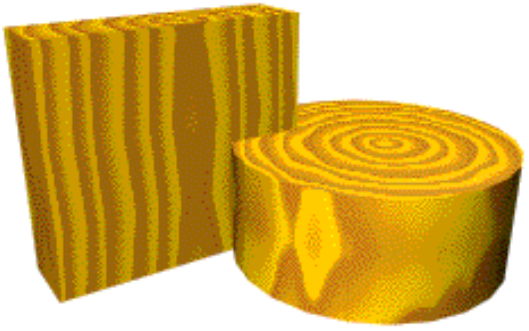


Wood map used for the seat of the bench

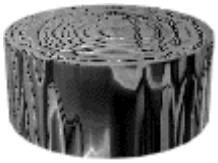
Wood is a 3D procedural map that renders a wavy grain-like pattern throughout the volume of an object. You can control the direction, thickness, and complexity of the grain.

Wood is primarily intended as a diffuse color map. Two colors assigned to Wood mix to form the grain pattern. Either of the colors can be replaced by other maps.

Wood can also be applied to other map types. When used as a bump map, Wood renders the grain pattern as a three-dimensional engraving of the surface.



Wood mapped to a box and cylinder



Wood used as a bump map

Procedures

To replace a color:

1. Click a color swatch labeled Color #1 or Color #2.
2. In the standard *Color Selector (page 1–161)*, choose a replacement color.

   The color updates in the color box and sample slot.

To exchange colors:

- Click Swap.

  The position of the two colors is reversed in the color boxes and sample slot.

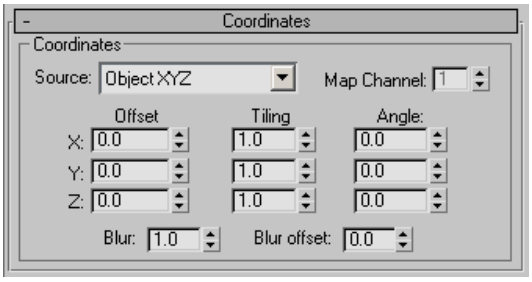
To replace a color with a map:

1. Click a Map bar marked None next to one of the color swatches. This displays the Material/Map Browser.
2. Select a map from the list.

   The sample slot updates to show the map in place of the color.

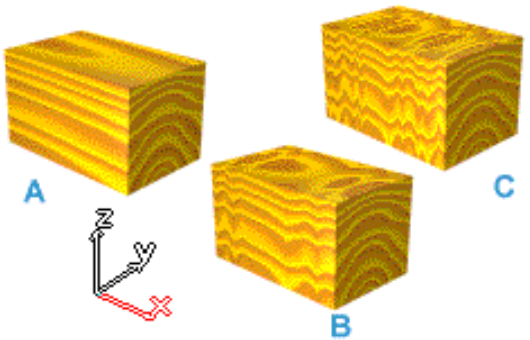
Interface

Coordinates rollout

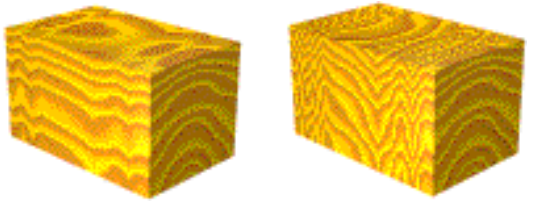


Tiling—Controls grain complexity or "distortion." By increasing this parameter along a given axis, the grain becomes increasingly compressed and wavy along that axis. Default=1.0 (X, Y, and Z axes).

Box A shows the default on all three axes. Boxes B and C show progressively higher Tile settings for the X axis. Increasing Tile on other axes produces similar effects.



A: Tile, X axis=1.0 (default)
B: Tile, X axis=4.0
C: Tile, X axis=8.0
Grain Thickness=3



Tile, X axis=4.0 for both
Axial Noise=1 (left), 4 (right)
Grain Thickness=3

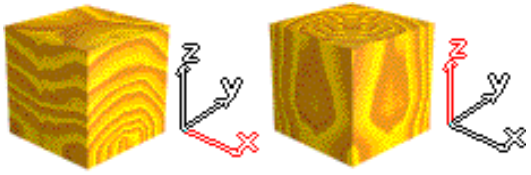
By combining Tile with Axial Noise, even greater distortion is possible. The bottom left box is the same as B, with Tile at 4.0 on the X axis. The bottom right box adds Axial Noise at 4.0.

Angle—Controls grain direction.

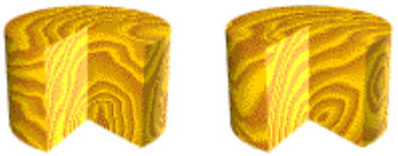
Default Grain: Wood is rendered with the grain running along the X axis of the object. This is illustrated by the left cube below. The default angle is 0 for X, Y, and Z axes.

Rotated Grain: By rotating the direction of the grain around an axis, you change the rendered effect. In the right cube, the Y axis is set to 90. This rotates the grain direction 90 degrees around the Y axis so that the grain is now running along the Z axis.

Other Angle Effects: Varying a single angle can create distorted grain effects (cylinder at left). Varying the three angles by a uniform amount keeps the grain running parallel through an object (cylinder at right).



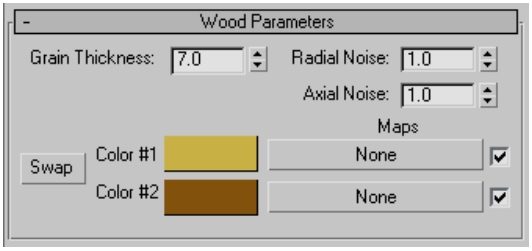
Angle=0,0,0 (defaults) and 0,90,0
Radial, Axial Noise=1 (defaults)
Grain Thickness=3

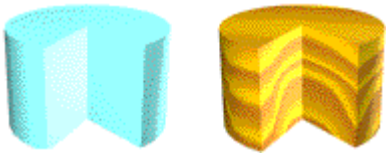


Angle=0,30,0 and 30,30,30
Radial, Axial Noise=1 (defaults)
Grain Thickness=3

(See *Coordinates Rollout (3D) (page 2–1663)* for a description of the other parameters in this rollout.)

Wood Parameters rollout





Defaults: Grain Thickness=7, Radial Noise=1.0, Axial Noise=1.0

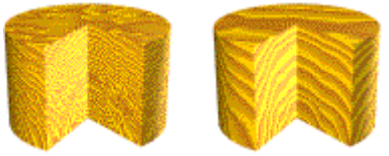
Grain Thickness—Sets the relative thickness of the color bands that make up the grain. Default=7.

The effect of thickness is relative to the context of the object. A grain that appears unrealistically wide on a small table might be acceptable on a large overhead beam.

Decreasing Grain Thickness creates grain lines that are closer together. The effect can resemble the fine grain of slow-growth hardwoods. At 0, grain disappears, resulting in what can look like composition board made from sawdust.

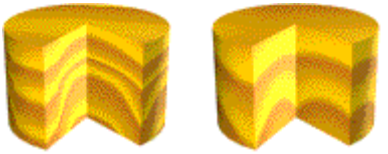
Increasing Grain Thickness creates grain lines that are farther apart. The effect can resemble tropical woods that grow continuously.

Grain thickness is shown increasing with identical noise settings.



Grain Thickness=1, 3

Radial, Axial Noise=1 (defaults)



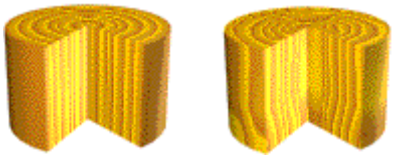
Grain Thickness=7 (default), 14

Radial, Axial Noise=1 (defaults)

Radial Noise—Sets the relative randomness of the pattern on a plane perpendicular to the grain, the circular ring structure (cylinder B). Default=1.0.

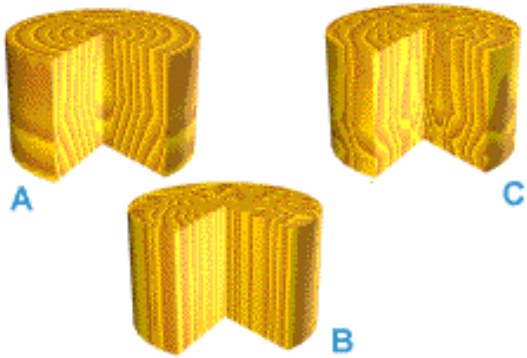
Axial Noise—Sets the relative randomness of the pattern on a plane parallel with the grain, along the length of the grain (cylinder A). Default=1.0.

Noise settings let you set the randomness or "irregularity" of the grain pattern in two directions. Without any noise, rings and grain are uniform and look inorganic (upper-left cylinder). The defaults for both parameters produce moderate irregularities (upper-right cylinder).



Radial, Axial Noise=0,0 and 1,1

Angle Y=90, Grain Thickness=3



A: Radial, Axial Noise=0,2

B: Radial, Axial Noise=2,0

C: Radial, Axial Noise=2,2

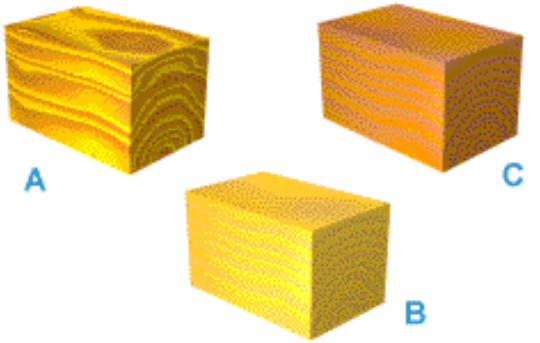
Angle Y=90, Grain Thickness=3

Cylinders A and B show each parameter acting alone. Cylinder C shows the combination of the same settings.

Swap—Exchanges the position of the colors.

Colors—Selects any two colors for the grain pattern. Defaults=brownish yellow for Color #1 and dark brown for Color #2. Either color can be replaced or swapped.

The choice of colors, along with grain pattern, is the primary way to represent different types of wood. In fairly uniform woods like yellow pine or redwood, the two colors are often near the same settings (examples B and C). Lighting also makes a difference in the apparent colors.

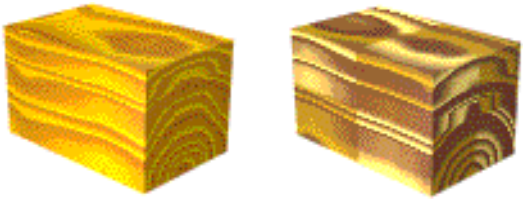


A: Default colors
B: Color #1=RGB 160,125,50
Color #2=RGB 170,135,25
C: Color #1=RGB 140,90,0
Color #2=RGB 130,80,50
Grain Thickness=3
Other settings at defaults

Maps—Replace colors with maps. The check boxes enable or disable their associated maps.

You can assign maps to one or both of the Wood color slots. Any kind of map can be used, including Wood. The map overrides the assigned color, which has no effect.

The left box uses defaults. The right box is the same, except that Color #1 has been replaced with a Checker map, whose colors have been changed to wood tones.



Left: Default
Right: Color #1 replaced with Checker map
Grain Thickness=3
Other settings at defaults

Compositor

Compositors are meant specifically for compositing other colors or maps. In image processing, compositing images refers to superimposing two or more images to combine them.

The following compositor maps are supplied with the software:

Composite (page 2–1688): Composites multiple maps. Unlike Mix, Composite doesn't have explicit controls for the amount of mixing. Instead, it bases the mix amount on the maps' alpha channel.

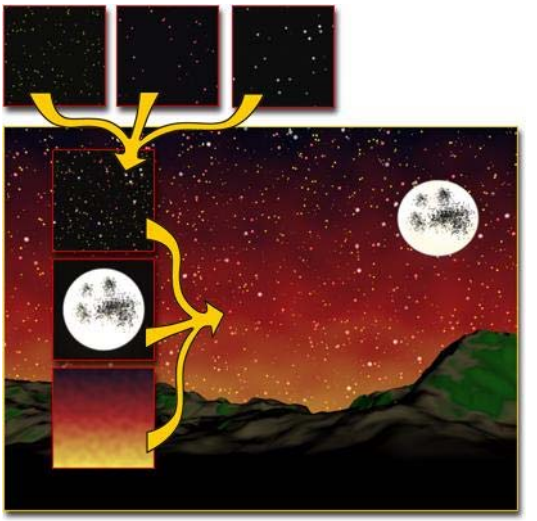
Mask (page 2–1689): A mask is a map itself, used in this case to control where a second map is applied to the surface.

Mix (page 2–1689): Mix mixes two colors or two maps. You can adjust the amount of mixing using a blend level you specify. The blend level can be mapped.

RGB Multiply (page 2–1691): Combines two maps by multiplying their RGB and alpha values.

Composite Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Composite



Composite map combines stars, moon, and a glow into the sky.

The Composite map type is made up of other *maps (page 3–968)*, which are overlaid on top of one another using the *alpha channel (page 3–907)*. For this type of map, use overlay images that already contain an alpha channel.

The controls for a Composite map are essentially a list of the maps it combines.

Viewports can display the multiple maps in a composite map. For multiple map display, the display driver must be *OpenGL (page 3–841)* or *Direct3D (page 3–844)*. The *software display driver (page 3–840)* does not support multiple map display.

Procedures

To assign a map:

1. In the Composite Parameters rollout, click a map button.

   The Browser is displayed.

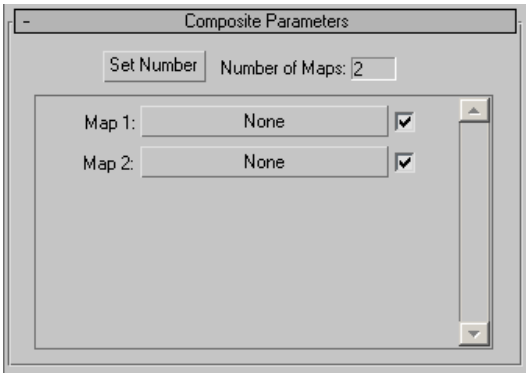
2. Choose a map type.

To change the number of maps to composite:

1. In the Composite Parameters rollout, click Set Number.
2. In the subdialog, choose the number of maps you want, and then click OK.

Warning: If you reduce the number of maps, previously assigned maps can be lost.

Interface



Set Number—Sets the number of maps to be composited.

Map 1, Map 2, and so on—Choose the button to add a map. Use the check boxes to enable or disable the associated map.

If there are more maps than the Material Editor can display in the Composite Parameters rollout, use the scroll bar on the right to move among them.

Mask Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Mask

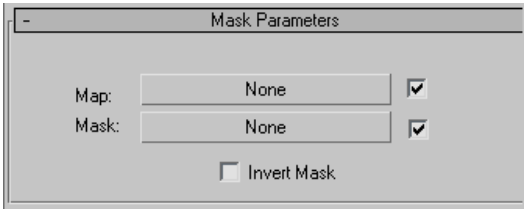


Mask map applies labels to the fire extinguisher.

With the Mask map, you can view one material through another on the surface. The mask controls where a second map is applied to the surface.

By default, lighter (whiter) areas of the mask are opaque, showing the map. Darker (blacker) areas of the mask are transparent, showing the underlying material. You can use Invert Mask to reverse the mask's effect.

Interface



These are the controls for the Mask map:

Map—Selects or creates the map to be viewed through the mask.

Mask—Selects or creates the map to use as a mask.

Invert Mask—Inverts the effect of the mask.

Mix Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Mix



Mix map blends skull and crossbones with the reflected scene.

With the Mix map, you can combine two colors or materials on a single side of the surface. You can also animate the Mix Amount parameter, and draw map *morphing (page 3–974)* function curves to control how the two maps are blended over time.

Viewports can display both maps in a mix map. For multiple map display, the display driver must be *OpenGL (page 3–841)* or *Direct3D (page 3–844)*. The *software display driver (page 3–840)* does not support multiple map display.

See also

For functionality shared with 3D maps, see the following topic:

Output Rollout (page 2–1621)

Procedures

To change a component color:

1. In the Mix Parameters rollout, click one of the two color swatches to display the *Color Selector (page 1–161)*.
2. Adjust the color.

To use a map as a component:

1. In the Mix Parameters rollout, click a map button next to one of the two color swatches.

   The Material/Map Browser is displayed.

2. Select a map type.

To exchange the two component colors:

- In the Mix Parameters rollout, click Swap.

To control the mix amount:

- In the Mix Parameters rollout, adjust the Mix Amount value.

  Mix Amount is the percentage of Color #2 used in the mix.

  You can also control the mix amount by using a map.

To control the mix amount using a map:

1. In the Mix Parameters rollout, click the map button next to Mix Amount.

   The Browser appears so you can

2. Select a map type.

   The intensity of pixels in this mixing map controls the mix. When the intensity is close to 0, one of the component colors or maps is visible; when it is close to full intensity, the other component is visible.

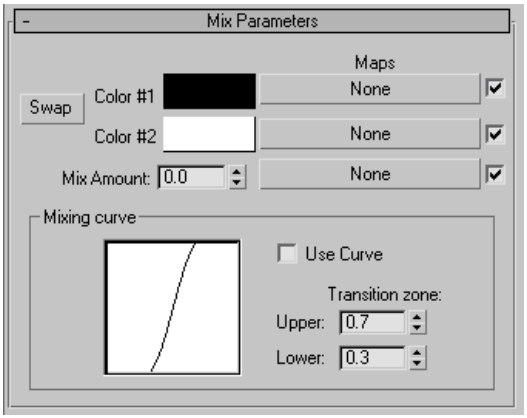
   Tip: Using a Noise map for the mixing map can give good effects that have a natural appearance.

   In the Mix Parameters rollout, Mix Amount is inactive while a map is assigned to this parameter. If Use Curve is off, the mixing map is used as is. If Use Curve is on, you can shift the effect of the mixing map's gradient ramp, offsetting it one way or the other and revealing more or less of the mix components.

To control the mix amount using the mix curve:

1. In the Mixing Curve group, turn on Use Curve.
2. Change the shape of the curve by adjusting the Transition Zone values.

Interface



Swap—Exchanges the two colors or maps.

Color # 1, Color # 2—Displays the *Color Selector (page 1–161)* to select the two colors to be mixed.

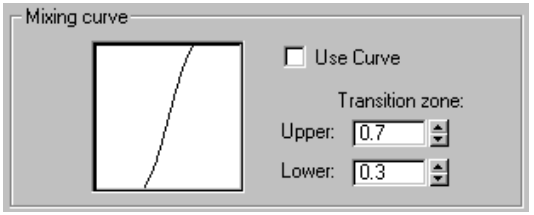
Maps—Select or create the bitmaps or procedural maps to be mixed instead of each color.

The check boxes enable or disable their associated maps.

Black areas of the map reveal color #1, and white areas of the map reveal color #2. Gray values reveal intermediate mixes.

Mix Amount—Determines the proportion of the mix. 0 means only Color 1 is visible on the surface, 1 means only Color 2 is visible. You can also use a map instead of the mix amount. The two colors will mix in greater or lesser degree according to the intensity of the map.

Mixing Curve group



These parameters control how gradual or how sharp the transition between the two colors being mixed will be. (This really only has meaning when you have a map applied to Mix Amount.)

Tip: Try mixing two standard materials using a noise map as a mask for some interesting mottled effects.

Use Curve—Determines whether the Mixing Curve effects the mix.

Transition Zone—Adjusts the level of the upper and lower limits. If the two values are the same, the two materials will meet at a definite edge. Wider ranges give more gradual mixing.

RGB Multiply Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > RGB Multiply



RGB Multiply enhances the texture of the ashtray.

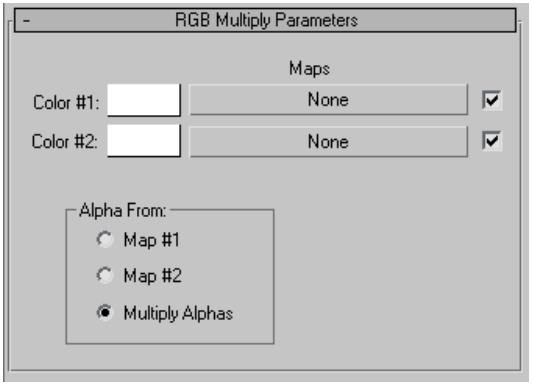
The RGB Multiply map is typically used for *bump maps (page 2–1506)*, where you might want to combine two maps to achieve the correct result.

This map combines two maps by multiplying their RGB values. For each pixel, the red of one map is multiplied times the red of the second map, the blue times the blue, and the green times the green.

If the maps have alpha channels, RGB Multiply can output either map's alpha channel, or a new alpha channel created by multiplying the alpha channel values of the two maps.

You can also make one of the maps a solid color. This tints the other map.

Interface



Color #1, Color #2—Click a map button to assign one of the maps. The check box disables or enables the map. To tint one of the maps, turn off the other map and click its color swatch to choose the tint color, using the *Color Selector (page 1–161)*.

Alpha From group

The buttons in this group let you determine how to generate alpha for the map. If neither map has an alpha channel, these options have no effect.

Map #1—Uses the first map's alpha channel.

Map #2—Uses the second map's alpha channel.

Multiply Alphas—Generates a new alpha channel by multiplying the alpha channels of the two maps.

Color Modifier

Color Modifier maps alter the color of pixels in a material. Each of the following maps uses a specific method to modify the color:

Output (page 2–1692): Applies bitmap output functions to parametric maps, such as Checker, that don't have these settings. These functions adjust the colors of the map.

RGB Tint (page 2–1693): Tints the color of a map based on red, green, and blue values.

Vertex Color (page 2–1693): Displays the effects of assigned vertex colors in the rendered scene. You assign vertex colors from the editable mesh.

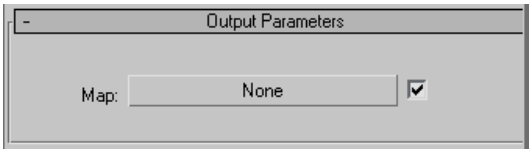
Output Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Output

With Output map, you can apply output settings to procedural maps, such as Checker or Marble, that don't have these settings.

Interface

Output Parameters rollout



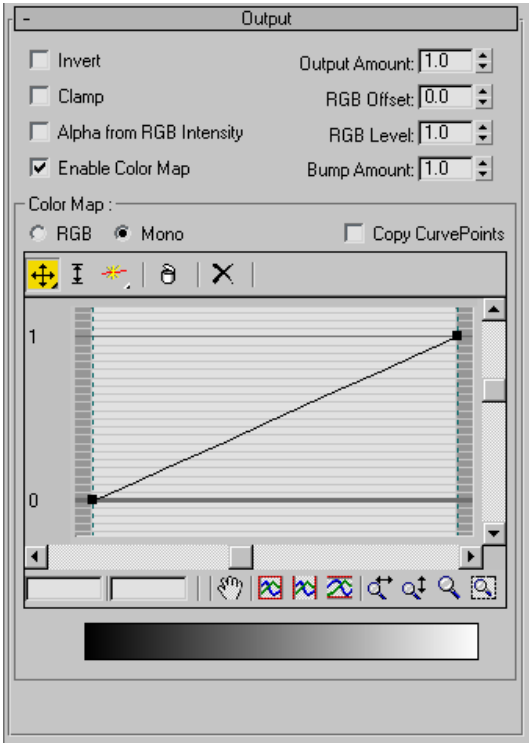
In this rollout, you choose the map to apply the output controls to.

Map—Displays a modal version of the Material/Map Browser so you can choose the map type.

The check box turns the map on or off.

Output rollout

The controls in this rollout are the same as for maps with a built-in output option. See *Output Rollout (page 2–1621)*.



RGB Tint Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > RGB Tint



RGB Tint map

RGB Tint adjusts the value of the three color channels in an image. Three color swatches represent these channels. Changing a color swatch adjusts the value of its associated color channel.

The channels are named Red, Green, and Blue for their default colors, but you can assign them any color. You are not limited to variations of red, green, and blue.

Procedures

To tint a map:

1. In the RGB Tint Parameters rollout, click the Map button marked None.

   The Material/Map Browser is displayed.
2. Select the map you want to tint
3. Click the R, G, or B color swatch.

   The *Color Selector (page 1–161)* is displayed.
4. Choose a new color.

   The red, green, or blue value of each pixel in the underlying map changes accordingly.

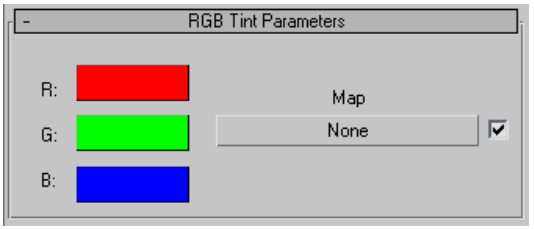
To change the saturation of one color in an image:

1. Click the R, G, or B color swatch.
2. On the Color Selector, increase or decrease Value to vary the color from light to dark.

To replace one color with another:

1. Click the R, G, or B color swatch.
2. On the Color Selector, increase or decrease Hue to change the color.

Interface



R/G/B—The *red, green, and blue (page 3–1001)* color swatches display the *Color Selector (page 1–161)* to adjust the value of the specific channel.

Map—Displays the Material/Map Browser to select the map to be tinted.

The check box turns the effect of the map on or off.

Vertex Color Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Vertex Color



Mapping vertex colors

Vertex Color map makes any vertex coloring applied to an object available for rendering. You

can assign vertex colors using the *VertexPaint Modifier (page 1–936)*, the *Assign Vertex Colors utility (page 2–1734)*, or the vertex controls for an *editable mesh (page 1–1003)*, *editable patch (page 1–975)*, or *editable poly (page 1–1029)*.

While vertex color assignment is primarily used for special applications, such as game engines or radiosity renderers, you can also use it to create colorful, gradient surface effects. You can also use it in design visualization: Use the *VertexPaint Modifier (page 1–936)* to paint your landscape different colors to represent grass, shrubbery, parking areas, etc., and then use Vertex Color map to use the vertex coloring in your rendered images. Incidentally, when you use the *Terrain object's (page 1–347)* Color By Elevation function, the software assigns a material that uses a Vertex Color map as the diffuse component.

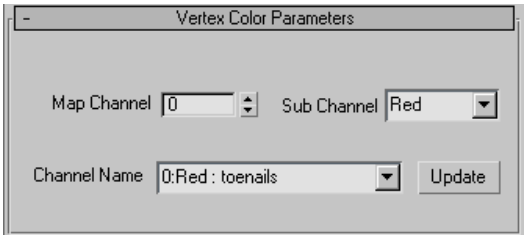
Tip: To view vertex colors in a viewport, right-click the object, choose Properties from the quad menu, and then turn on Vertex Channel Display in the Display Properties group.

Procedure

To use the vertex color map:

1. Assign vertex colors to an object.
2. Assign a material to the object, then assign a Vertex Color map to the material's diffuse component.
3. Optionally, if manipulating the map channels with the *Channel Info utility (page 2–1738)*, choose a map channel or sub-channel to render.
4. Render the scene.

Interface



These parameters let you define which map channel or sub-channel is to be rendered. One application is to support usage of the Vertex Color map in conjunction with the *Channel Info utility (page 2–1738)*.

The settings are interlinked; changing one parameter will change the other two, as appropriate.

Map Channel—Lets you specify which map channel to use. Range=0 to 99. Default=0.

Notes regarding this setting:

- If you set Map Channel to a channel that doesn't contain any vertex coloring data, attempting to render will generate a Missing Map Coordinates error message. To resolve this, apply vertex coloring to that channel.
- By default, the vertex coloring in map channel 1 is a color gradient derived from the UVW texture coordinates by converting UVW values to RGB values. Thus, at UV=0,0 (the lower-left corner of the map), the coloring is black; at UV=1,0, the coloring is red, and at UV=1,1 (the upper-right corner), the coloring is yellow (red + green=yellow). You can change these colors with a tool such as *VertexPaint modifier (page 1–936)*.
- Map Channel cannot be set to a negative value, thus the map doesn't support rendering of the vertex illumination (-1) or vertex alpha (-2) channel.

Sub Channel—Lets you can specify that the map will use either the Red, Green, or Blue sub-channel of the specified map channel, or all sub-channels.

Channel Name—After assigning the material with the Vertex Color map to an object with named map or vertex-color channels (see *Channel Info Utility (page 2–1738)*), you can click Update, and then, from this drop-down list, choose a named map channel from the object.

Update—Refreshes the contents of the Channel Name drop-down list. Use Update after applying the material to an object, or after adding channels to the object.

Note: There could be conflicts if one material with a Vertex Color map is assigned to objects with different named Map Channels, where one channel's name may be displayed in preference to another's.

Reflection and Refraction

These maps, grouped as "Other" in the *Material/Map Browser (page 2–1412)*, are maps that create reflections and refractions. Each of the following maps has specific uses:

Flat Mirror (page 2–1695): Generates reflections for flat surfaces. You assign it to faces rather than to the object as a whole.

Raytrace (page 2–1698): Creates accurate, fully raytraced reflections and refractions.

Reflect/Refract (page 2–1699): Generates reflections or refractions automatically, based on surrounding objects and the environment.

Thin Wall Refraction (page 2–1703): Generates refractions automatically, simulating objects and the environment seen through a refractive material such as glass or water.

Flat Mirror Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Flat Mirror



Flat mirror map reflects the ice-cream shop's interior.

The *Flat Mirror (page 3–942)* map produces a material that reflects surrounding objects when it is applied to a collection of coplanar faces. You assign it as a material's *reflection map (page 2–1508)*.

Reflect/Refract maps don't work well for flat surfaces because each face reflects part of the environment based on where its surface normal points. Using this technique, a large flat face can reflect only a small part of the environment. Flat Mirror automatically generates a reflection that encompasses a larger part of the environment, to better simulate a mirror-like surface.

Rules for Using Flat Mirror

Flat Mirror cannot generate reflections correctly unless you observe these rules:

- Assign Flat Mirror to selected faces only.

  There are two ways to do this. You can make the Flat Mirror material a sub-material of a *Multi/Sub-Object (page 2–1594)* material, or you can use the Apply To Faces With ID control.

- If you assign Flat Mirror to multiple faces, the faces must lie in a plane.
- Non-coplanar faces in the same object cannot have the same Flat Mirror material.

  In other words, if you want two different planes of an object to have flat reflections, you must use a Multi/Sub-Object material. Assign Flat Mirror to two different sub-materials, and assign different material IDs to the to different planar faces.
- The material ID used by the Flat Mirror sub-material must be unique to the coplanar faces in the object.

  If you assign Flat Mirror using Apply to Faces with ID, faces without that ID display the nonreflective components (diffuse color, and so on) of the material with the Flat Mirror reflection map.

Procedures

To assign a flat mirror to one face of an object:

1. In the Material Editor, create a standard material.
2. Assign a Flat Mirror map as the material's reflection map.
3. In the Flat Mirror Parameters rollout > Render group, turn on Apply To Faces With ID, and choose the material ID number the mirrored face will have.
4. Follow the next set of steps for assigning the material to the object.

To assign the mirror to a flat surface:

1. Select an object.
2. In the Modify command panel, apply Edit Mesh to the object.
3. Make sure Sub-Object is selected, and choose Face as the sub-object level.
4. Select a single face or multiple faces that lie in a single plane.
5. Assign the faces the material ID you chose for the Flat Mirror map.
6. Assign the material to the object.

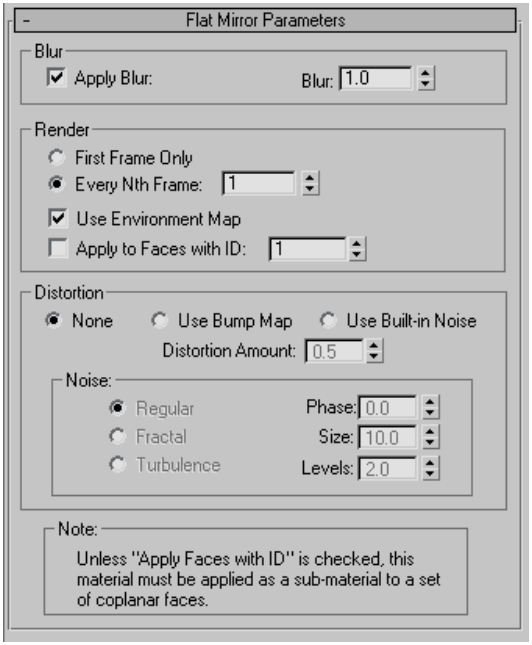
To assign a flat mirror using a Multi/Sub-Object material:

1. In the Material Editor, create a Multi/Sub-Object material.
2. Click one of the unused material buttons in the Multi/Sub-Object material's parameters.
3. In the new Standard sub-material, open the Maps rollout and click the map button for Reflection.
4. In the Material/Map Browser, choose Flat Mirror, and then click OK.

   Flat Mirror controls are similar to those for automatic reflection and refraction.
5. Apply Edit Mesh to the object, and then in the stack view area of the *modifier stack display (page 3–760)*, choose Face as the sub-object level.
6. Select a single face or multiple faces that lie in a single plane.
7. Assign the faces the material ID corresponding to the Flat Mirror sub-material slot.

   Using a Multi/Sub-Object material, you can apply Flat Mirror to different faces of the object that are not coplanar. However, faces that are not coplanar must use different sub-material slots, otherwise the software doesn't correctly generate the flat mirror reflections.

Interface



Blur group

Apply Blur—Turns on filtering to blur the maps.

Antialiasing is also applied to the Distortion effect, if any, when Apply Blur is turned on.

Blur—Affects the sharpness or blurriness of the generated map based on its distance from the object. The farther away the map is, the greater the blurring. Blur is primarily used to avoid *aliasing (page 3–907)*. It's a good idea to use a small amount of blurring for all maps in order to avoid the scintillation or aliasing that can occur when pixel details are reduced off in the distance. Default=1.0.

Render group

First Frame Only—The renderer creates the automatic flat mirror only on the first frame.

Every Nth Frame—The renderer creates the automatic flat mirror based on the *frame rate (page 3–944)* set by the spinner.

Use Environment Map—When off, environment maps are ignored by the mirror during rendering. It's useful to turn this off when you have mirrors in the scene and you're rotoscoping against a flat screen environment map. A screen environment map does not exist in 3D space the way the other environment-map types do, and will not render properly. Default=on.

Apply to Faces with ID—Specifies the material ID number where you want the mirror assigned.

You can assign a flat-mirror material to an object without having to make it a component of a *Multi/Sub-Object (page 2–1594)* material. The restriction is that the other faces on the object must be able to use the nonmirrored properties of the same material (its diffuse color, and so on). If the other faces need completely different material characteristics, you need to use a Multi/Sub-Object material.

For example, if you have an object, such as a box, with unique material IDs for each side, you can use Apply To Faces With ID to specify the side of the box that will show the mirror reflection. The remaining sides of the box will have the same material characteristics, but without the reflection.

Distortion group

To simulate irregular surfaces, you can distort the flat-mirror reflections. Distortion can be based on a bump map or on noise controls built into Flat Mirror material.

None—No distortion.

Use Bump Map—Distorts the reflection using the material's bump map.

A flat mirror surface that has a Bump map will appear bumpy, but its reflection won't be distorted by the bumps unless you use this option.

Use Built-In Noise—Distorts the reflection using the settings in the Noise group.

Distortion Amount—Adjusts the amount of distortion to the reflected image. This is the only value that affects the amount of distortion. No matter how high the Bump map's Amount spinner is set, or how extreme the Noise settings, if this Distortion Amount is set to 0, no distortion appears in the reflection itself. This control is inactive when None is chosen.

Noise group

The controls in this group are inactive unless you choose Use Built-In Noise as the active distortion type.

Regular—Generates plain noise. Basically the same as Fractal noise with the Levels setting at 1. When the noise type is set to Regular, the Levels spinner is inactive (because Regular is not a fractal function).

Fractal—Generates noise using a fractal algorithm. The Levels setting determines the number of iterations for the fractal noise.

Turbulence—Generates fractal noise with an absolute value function applied to it to make fault lines.

Phase—Controls the speed of the animation of the noise function. A 3D noise function is used for the noise, so that the first two parameters are U and V and the third is phase.

You can animate this parameter to animate the noise effect.

Size—Sets the scale of the noise function. Smaller values give smaller chunks of noise.

Levels—Sets the number of fractal iterations or turbulence (as a continuous function).

Raytrace Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Raytrace



Raytrace map creates highly reflective and refractive surfaces.

Raytrace map provides fully raytraced reflections and refractions. The reflections and refractions it generates are more accurate than those produced by the reflect/refract *map (page 2–1699)*. Rendering raytraced objects is slower than using Reflect/Refract. On the other hand, Raytrace is optimized for rendering 3ds Max scenes, and you can further optimize it for your scene by excluding specific objects or effects from raytracing.

You can also use the *Raytrace material (page 2–1512)*, which uses the same raytracer to generate accurate, raytraced reflections and refractions. The differences between Raytrace map and Raytrace material are:

- You use Raytrace map as you do other maps. This means you can add raytraced reflections or refractions to any kind of material.
- You can assign Raytrace map to material components other than reflect or refract,

although these are the main ways to use this map.

- Raytrace map has more extensive attenuation controls than Raytrace material.
- Raytrace map often renders more quickly than Raytrace material.

  Raytrace Map and Raytrace material have the same name because they use the same raytracer and share global parameters.

  Note: Raytracing does not always work correctly in orthogonal viewports (left, front, and so on). It works correctly in perspective viewports and camera viewports.

Interface

The interface for the Raytrace map is contained in four rollouts:

Raytracer Parameters Rollout (page 2–1704)

Raytrace: Attenuation Rollout (page 2–1706)

Raytrace: Basic Material Extensions Rollout (page 2–1707)

Raytrace: Refractive Material Extensions Rollout (page 2–1708)

Reflect/Refract Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Reflect/Refract



Reflect/Refract map used for the balloons

The Reflect/Refract map produces a reflective or refractive surface. To create reflection, assign this map type as the material's *reflection map (page 2–1508)*. To create refraction, assign it as the *refraction map (page 2–1509)*.

Note: A reflective object can reflect another reflective object. In real life, this creates a virtually infinite number of interreflections. In 3ds Max, you can set the number of interreflections in a range from 1 to 10. You set this Rendering Iterations parameter in the *Render Scene dialog (page 3–12)*.

This map works by using six renderings in the form of a cube that surrounds the objects. Reflect/Refract views the cubic maps from the perspective of the *pivot point (page 3–995)* of the object, mapping them onto the object's surface as a spherical reflection map.

You can choose to generate the cubic maps automatically, or to load previously created maps.

A reflective surface reflects the surrounding maps like a mirror. A refractive surface creates

the illusion that the surrounding maps are seen through the surface.

Note: Reflect/Refract is meant to be used with curved or irregularly shaped objects. For mirror-like flat surfaces that you want to reflect the environment accurately, use *Flat Mirror (page 2–1695)* material. For more accurate refractions, especially for an object in a refractive medium (such as a pencil in a glass of water), use *Thin Wall Refraction material (page 2–1703)*.

Automatic Cubic Maps

When you choose automatic cubic maps, 3ds Max generates the maps based on the perspective of the mapped object's pivot point. Automatic maps are based on the geometry of the scene.

The advantage to using automatic maps is that the six views are automatically generated at rendering time and can be easily updated for each frame of an animation. The disadvantage is that regenerating the maps increases rendering time. In addition, the maps exist only for the rendering and can't be edited or otherwise manipulated.

In the Material Editor, automatic reflections or refractions reflect or refract the background of the sample slot.

Assigning Cubic Maps

When you choose to load cubic maps from files, you use the controls in the From File group. The advantage to using From File is that the bitmap files already exist and take less rendering time. In addition, you can edit the bitmap images. The disadvantage is that it's more difficult to use bitmaps to render an animation because first you need to create the animated bitmaps.

Important: Assigned cubic maps must be square, and each of the six maps must be the same size.

You can automatically load six bitmaps at once if the six bitmaps have valid cubic-map file names. The first part of all six file names must be the same, and the last part is an underscore followed by a two-letter abbreviation of the map position, as shown in the following table:

Last Part of Cubic Map File Name	Meaning
_UP	Up
_DN	Down
_LF	Left
_RT	Right
_FR	Front
_BK	Back

For example, if you have six bitmaps on disk, labeled *view_up.bmp, view_fr.bmp, view_lf.bmp,* and so on, when you click one of the file buttons and choose the bitmap for that position, all six views are loaded.

If you select a bitmap whose name doesn't follow the convention, or you select one with a valid cubic-map name but assign it to the wrong button, only that bitmap will be loaded.

Since all assigned bitmaps must be the same size, when you assign a new bitmap to one of the windows, the sample slot doesn't update automatically. This avoids generating an error message each time you assign a bitmap. Once you've assigned all six maps and are sure they are square and the same size, click the Reload button to update all of the maps and redisplay the sample slot.

You can also use the Reload button to see your changes after you've edited one of the cubic maps by using a paint program.

Rendering Cubic Maps

The controls in the Render Cubic Map Files group let you generate the maps automatically and save them to disk. Use the To File button to specify the

folder and file name of the Up (_UP) bitmaps. Click Pick Object and Render Maps, and then click the object to map. The software creates the files and also assigns them to the six From File map buttons.

Rendering cubic maps has the same effect as automatic with the advantage that map rendering doesn't have to take place at scene rendering time. The disadvantage is that you can't create an automatically animated reflections or refractions this way.

Using Animated Cubic Maps

The cubic maps can be animations instead of bitmaps. Be sure that each is square and all have the same resolution. If you want the animated reflections to match animated changes in the scene, use automatic maps and set them to render every Nth frame.

Procedure

To generate and save cubic maps:

1. In the Source group, choose From File.
2. In the Render Cubic Map Files group, click the button next to To File.

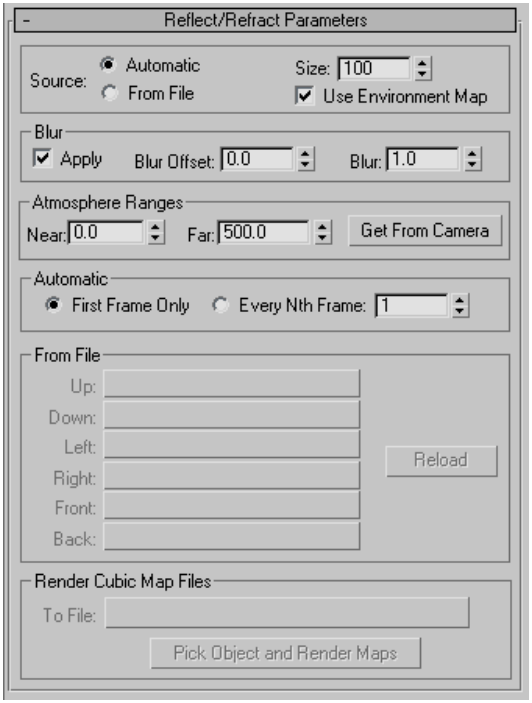
   A file dialog is displayed. Type a name for one of the six cubic bitmap files. You're specifying the name of the Up bitmap. There are two ways to do this:

   - Specify the entire file name; for example, *myview_up.bmp*.
   - Specify the file prefix and extension only; for example, *myview.bmp*.

     Important: You must specify at least a prefix and extension. The "_UP" is added automatically.
3. Click Pick Object and Render Maps.
4. Click an object in your scene where you want the six views rendered. This is usually the same object where you will apply the material, but it doesn't have to be.

   After selecting the object, a rendered frame window is displayed while the six views are rendered and saved to disk. Their file names then appear in the six buttons in the From File group. Each of the six file names is identical, except for a two-letter suffix indicating the direction of the rendered view.

Interface



Source—Chooses the source of the six cubic maps.

Automatic—Automatically generates by looking out in six directions from the pivot of the object with the material, then mapped onto the surface during rendering. When on, the options in the Automatic group are active, letting you choose whether the maps will be generated only once, or regenerated at specified frames in the animation.

From File—When on, you can specify the bitmaps to use.

When From File is active, the controls in the Render Cubic Map Files group are also available. You can generate the six cubic reflection maps automatically and save them to files, where you can load them with the From Files controls.

Size—Sets the size of the Reflect/Refract maps. The default value of 100 produces distinct images. Lower values lose progressively more detail.

Use Environment Map—When off, environment maps are ignored by Reflect/Refract map during rendering. It's useful to turn this off when you have mirrors in the scene and you're rotoscoping against a flat screen environment map. A screen environment map does not exist in 3D space the way the other environment-map types do, and will not render properly.

Blur group

Apply—Turns on filtering to blur the maps.

Blur Offset—Affects the sharpness or blurriness of the map without regard to its distance from the object. Use Blur Offset when you want to soften or defocus the details in a map to achieve the effect of a blurred image.

Blur—Affects the sharpness or blurriness of the generated map based on its distance from the object. The farther away the map is, the greater the blurring. Blur is primarily used to avoid *aliasing (page 3–907)*. It's a good idea to use a small amount of blurring for all maps in order to avoid the scintillation or aliasing that can occur when pixel details are reduced off in the distance. Default=1.

Atmosphere Ranges group

If your scene contains environmental *Fog (page 3–282)*, the cubic maps must have near and far range settings to properly render the fog from the point of view of the object assigned the material. The Near and Far spinners in this group let you specify a fog range relative to the object.

Near—Sets the near range for fog.

Far—Sets the far range for fog.

Get From Camera—Uses the Near and Far atmosphere range settings of a camera in the scene. Click this option, and then select the camera.

These values aren't dynamically linked to the camera object. They are simply copied from the camera's range values at the time you click the camera. If the camera's range values later change, the map's Near and Far values remain the same.

Automatic group

This controls are active only when Automatic is the active source for the Reflect/Refract maps.

First Frame Only—Tells the renderer to create automatic maps only on the first frame.

Every Nth Frame—Tells the renderer to create animated auto maps based on the *frame rate (page 3–944)* set by the spinner.

From File group

These controls are active when From File is active as the Reflect/Refract source. Here you assign the six bitmaps to be used as the cubic maps.

Up / Down / Left / Right / Front / Back—Assigns one of the six cubic maps. If the map is one of a set of six with the correct file name, all six are loaded.

If the map doesn't follow file naming conventions, or you assign it to a button of a different position (_UP to Front, for example), only that map is assigned. You must assign the others manually.

Reload—Reloads the assigned maps and updates the sample slot.

You can edit one or more of the cubic maps using a paint program, then click Reload to update the material and the scene.

Render Cubic Map Files group

To File—Choose a file name for the Up map (_UP).

Pick Object and Render Maps—Active when you choose a file. Click to turn on, then select the mapped object to render the six cubic maps. Assign the cubic maps to the six From File buttons.

Thin Wall Refraction Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Thin Wall Refraction



Thin wall refraction

Thin Wall Refraction simulates the "jog," or offset effect, you find when you view part of an image through a plate of glass. For objects that model glass, such as a Box in the shape of a window pane, this map is faster, uses less memory, and provides a much better visual effect than the Reflect/Refract map.

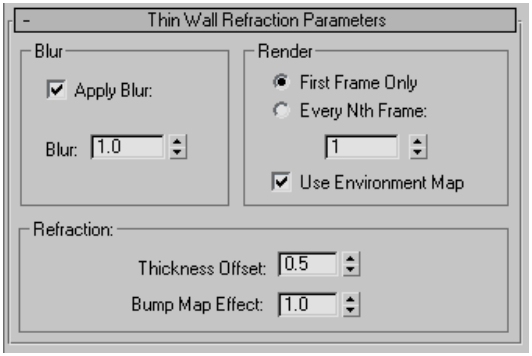
Tip: At 100% refraction and opacity, you can see no diffuse color or mapping, and there is not much illusion of a refractive material. The effect is invisible. In the Maps rollout of the parent material, set Refraction Amount to 50%, and in the Basic Parameters rollout, set Opacity to a value greater than 0.

Procedure

To assign the Thin Wall Refraction map to a material:

1. Click the Map button for Refraction in the material's Maps rollout.
2. In the Material/Map Browser, choose Thin Wall Refraction.
3. Adjust the map's parameters.
4. In the parent material, set the Refraction Map Amount to 50%.
5. Assign the material to an object.

Interface



Blur group

These controls are for antialiasing.

Apply Blur—Turns on filtering to blur the maps.

Blur—Affects the sharpness or blurriness of the generated map based on its distance from the object. The farther away the map is, the greater the blurring. Blur is primarily used to avoid *aliasing (page 3–907)*. It's a good idea to use a small amount of blurring for all maps in order to avoid the scintillation or aliasing that can occur when pixel details are reduced off in the distance. Default=1.0.

Render group

These controls affect how the refraction should behave in animations.

First Frame Only—Tells the renderer to create the refracted image only on the first frame.

This is the fastest option. You can use it if the camera and refractive object don't move.

Every Nth Frame—Tells the renderer to regenerate the refracted image based on the *frame rate (page 3–944)* set by the spinner.

Every single frame provides the most accurate result, but takes longest to render.

Use Environment Map—When off, environment maps are ignored by the refraction during rendering. It's useful to turn it this off when you have refractions in the scene and you're rotoscoping against a flat screen environment map. A screen environment map does not exist in 3D space the way the other environment map types do, and will not render properly. Default=on.

Refraction group

These controls are specific to the Thin Wall Refraction effect:

Thickness Offset—Affects the size of the refractive offset, or jog effect. At 0, there's no offset, and the object can appear invisible in the rendered scene. At 10.0, the offset is at its greatest. Range from 0.0 to 10.0; Default=0.5.

Note: The IOR (index of refraction) spinner in the parent material's Extended Parameters rollout also affects the offset effect.

Bump Map Effect—Affects the magnitude of refraction due to the presence of a bump map. This parameter multiplies the current bump map Amount in the parent material. Reduce this value to reduce the effect of the secondary refraction; increase this value to increase the effect. If there is no bump map assigned, this value has no effect. Default=1.0.

If there is unevenness in the surface of the glass, there is a secondary refraction. Thin Wall Refraction generates this secondary refraction if the material also has a bump map present. The algorithm guesses at the scaling of the secondary refraction, and can create too large an effect. If this happens, scale the effect down by reducing this value to less than one.

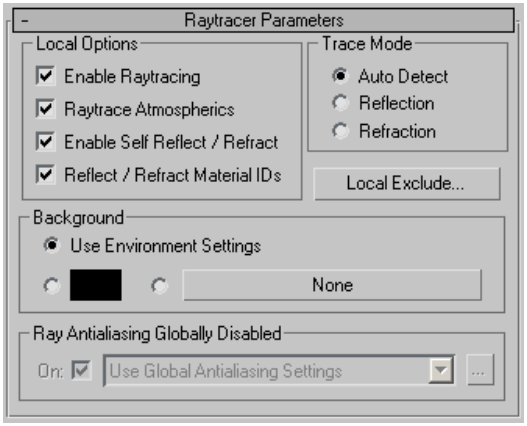
Raytrace Map Rollouts

Raytracer Parameters Rollout

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Raytrace > Raytracer Parameters rollout

This rollout contains the main controls for the *Raytrace map (page 2–1698)*.

Interface



Local Options group

Enable Raytracing—Turns the raytracer on or off. Default=on.

Even with raytracing off, Raytrace material and Raytrace map still reflect and refract the environment, including both the environment map for the scene, and the environment map assigned to the Raytrace material.

Raytrace Atmospherics—Turns the raytracing of atmospheric effects on or off. Atmospheric effects include fire, fog, volume light, and so on. Default=on.

Enable Self Reflect/Refract—Turns self reflection/refraction on or off. Default=on.

Can an object reflect itself? For example, a teapot's body reflects the teapot's handle, but a sphere will never reflect itself. If you don't need this effect, you can improve render time by turning off this toggle.

Tip: If you have a transparent object such as glass, and have self reflect/refract turned on, you don't have to make the object *2-sided (page 3–901)*. The raytracer sees back faces when exiting refractive objects.

Reflect/Refract Material IDs—When on, the material reflects effects assigned to material IDs in the renderer's *G-buffer (page 3–946)* on or off. Default=on.

By default, Raytrace material and Raytrace map reflect effects assigned to a material's ID, so that G-buffer effects are not lost. For example, if a raytraced object reflects a lamp made to glow with the Video Post Glow filter (Lens Effects Glow), the reflection glows as well.

Trace Mode group

With options in this group, you select whether to cast reflected or refracted rays.

Auto Detect—If assigned to the material's Reflection component, the raytracer will reflect. If assigned to Refraction, it will refract. If you assign Raytrace to any other component, you have to manually specify whether you want reflected rays or refracted rays. (Default.)

Note: Auto Detect might fail when you use Raytrace map in a material with a strong bump map. When you use a strong bump map, choose one of the explicit options.

Reflection—Casts reflected rays off the object's surface.

Refraction—Casts refracted rays into or through the object's surface.

Note: Raytrace reflects and transmits the IDs in *material effects channel (page 2–1443)* (*G-buffer (page 3–946)*), so it can create glowing reflections, and so on.

Local Exclude—Click to display the local *Include/Exclude dialog (page 2–1531)*.

An object that is excluded locally is excluded from this map only.

Tip: Using exclusion lists is one of the best and simplest ways to speed up the raytracer.

Background group

Use Environment Settings—Respects the environment settings of the current scene.

Color Swatch—Overrides the environment settings with the specified color.

Map Button—Overrides the environment settings with the specified map.

By specifying an environment map, you override the environment map for the scene as a whole. Both reflection and refraction use the scene-wide environment map unless you use this option to specify another map. With this control, you can use different environment maps on a per-object

basis, or provide an environment to specified objects when the scene as a whole has none.

Raytraced Reflection and Refraction Antialiaser group

Controls in this group let you override the global antialiasing settings for raytraced maps and materials. They are unavailable if antialiasing is turned off globally. To turn on antialiasing globally, choose Rendering > Raytrace Globals to display the *Global Raytracer Settings dialog (page 2–1528)*.

On—When on, uses antialiasing. Default=unavailable unless global antialiasing is turned on; on if global antialiasing is turned on.

Drop-down list—Chooses which antialiasing settings to use. There are three alternatives:

- **Use Global Antialiasing Settings—**(The default.) Uses the global antialiasing settings.
- **Fast Adaptive Antialiaser—**Uses the Fast Adaptive antialiaser, regardless of the global setting.
- **Multiresolution Adaptive Antialiaser—**Uses the Multiresolution Adaptive antialiaser, regardless of the global setting.

...—The button with the ellipsis to the right of the drop-down list displays another dialog to let you set antialiasing controls locally. The dialog displayed depends on which alternative you chose in the drop-down list, as follows:

- **Use Global Antialiasing Settings—**Clicking ... displays the *Global Raytracer Settings dialog (page 2–1528)*.
- **Fast Adaptive Antialiaser—**Clicking ... displays the Fast Adaptive Antialiaser *dialog (page 2–1533)*.
- **Multiresolution Adaptive Antialiaser—**Clicking ... displays the *Multiresolution Adaptive Antialiaser dialog (page 2–1534)*.

When you change settings for an antialiaser locally, you don't affect the global settings for that antialiaser.

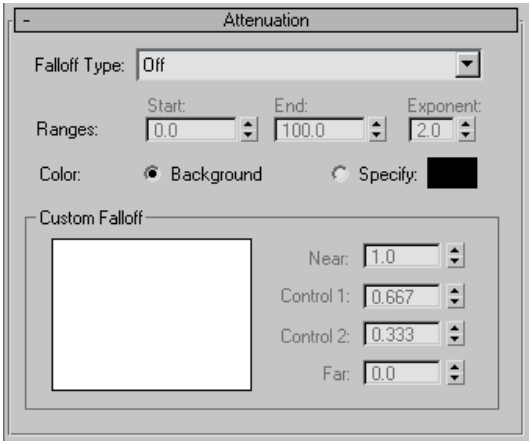
Raytrace: Attenuation Rollout

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Raytrace > Attenuation Rollout

When a ray is reflected off an object or refracted through it, by default the ray travels through space forever, with no *attenuation (page 3–912)*. The controls in this rollout allow you to attenuate rays, so their strength diminishes over distance.

In the *Raytrace map (page 2–1698)*, attenuation is implemented by a clipping algorithm. Objects beyond the maximum attenuation range aren't even considered by the raytracer. Because of this, assigning attenuation can speed up rendering time.

Interface



Falloff Type—Choose the falloff to use.

- **Off—**Turns off attenuation. (Default.)
- **Linear—**Sets linear attenuation. Linear attenuation is calculated between the start and end range values.

- **Inverse Square**—Sets inverse square attenuation. Inverse square attenuation is calculated beginning at the start range, and doesn't use the end range. Inverse square is the actual attenuation rate for light in the real world. However, it doesn't always give the effect you want in a rendered scene.
- **Exponential**—Sets exponential attenuation. Exponential attenuation is calculated between the start and end range values. You also specify the exponent to use.
- **Custom Falloff**—Specifies a custom curve to use for attenuation (falloff).

Start Range—The distance in world units where attenuation begins. Default=0.0.

End Range—Sets the distance in world units where the ray is fully attenuated. Default=100.0.

This is not used by inverse square attenuation.

Exponent—Sets the exponent used in exponential falloff. Default=2.0.

This is used only by exponential attenuation.

Color

These controls affect the behavior of light rays as they attenuate out. By default, as a ray fades out, it is rendered as the background color.

You can set a custom color instead.

Background—As the ray attenuates out, returns the background (either the scene's background or the background specified locally in the Raytracer Parameters rollout) rather than the actual color of what the reflected/refracted ray sees. (Default.)

Specify—Sets the color that is returned by the ray as it attenuates out.

If you choose not to use the background color, black or gray usually work best as the attenuation color,

Custom Falloff group

These controls are inactive unless the Falloff Type is set to Custom Falloff.

Custom Falloff—Uses the falloff curve to determine the falloff between the start and the end ranges.

These are the controls for custom attenuation. The custom attenuation curve is at the left. The gray scale bar below the curve shows how the curve will affect the falloff as light rays diminish in strength.

Near—Sets the strength of the reflected/refracted ray at the start range distance. This is a normalized percentage that can range from 0.0 to 1.0. Default=1.0.

Control 1—Controls the shape of the curve near the curve start. Default=0.667.

Control 2—Controls the shape of the curve near the curve end. Default=0.333.

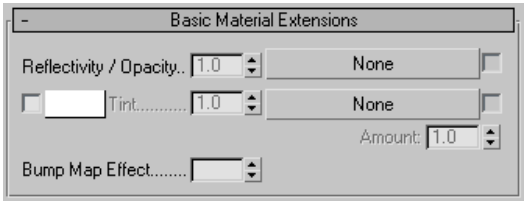
Far—Sets the strength of the reflected/refracted ray at the end range distance. This is a normalized percentage that can range from 0.0 to 1.0. Default=0.0.

Raytrace: Basic Material Extensions Rollout

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Raytrace > Basic Material Extensions Rollout

This rollout contains controls for fine-tuning the effect of the *Raytrace map (page 2–1698)*.

Interface



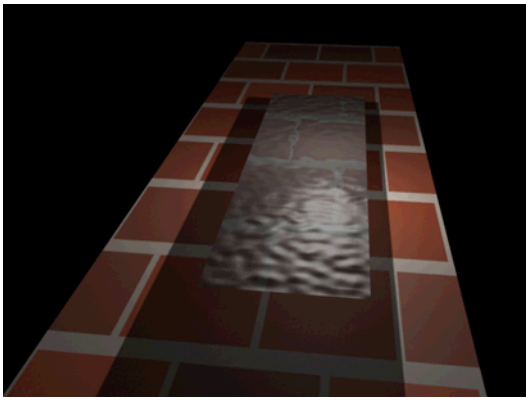
Reflectivity/Opacity—These controls affect the intensity of the raytracer's results.

- **Spinner—**Controls the amount of raytracing used by the material it is assigned to. Analogous to the Output Amount parameter in the Output rollout of the *Bitmap map type (page 2–1631)*.
- **Map button—**Assigns a map that controls the amount of raytracing. You can vary the amount of raytracing used over the surface of the object.
- **Check box—**Enables or disables the map.

Tint—With these controls, you can tint the colors returned by the raytracer. Tinting applies only to reflected colors; it doesn't affect the material's diffuse component.

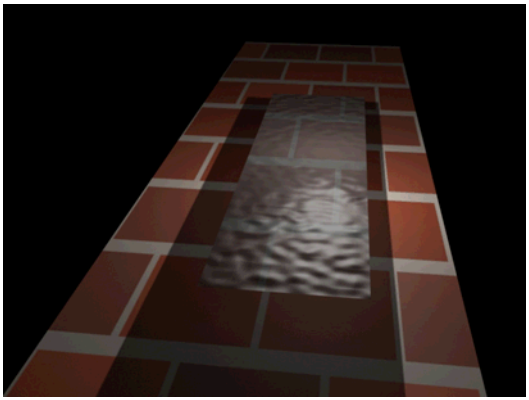
- **Check box—**Turns basic tinting on or off. Default=off.
- **Color swatch—**Assigns a tint color for reflections. Default=white.
- **Amount spinner—**Sets the amount of tinting used. Default=1.0.
- **Map button—**Assigns a map to use for tinting. You can vary the tint colors over the surface of the object.
- **Check box—**Enables or disables the map.

Bump Map Effect—Controls the effect of a bump map on rays that the surface reflects and refracts. For example, you might want to make a glass object highly bumpy, but reduce the bump effect on refracted parts of the scene.



Bump Map Effect equals 1.0.

The refraction of tiles behind the glass is distorted by the bump map.



Bump Map Effect equals 0.0.

The refraction of tiles behind the glass is not distorted.

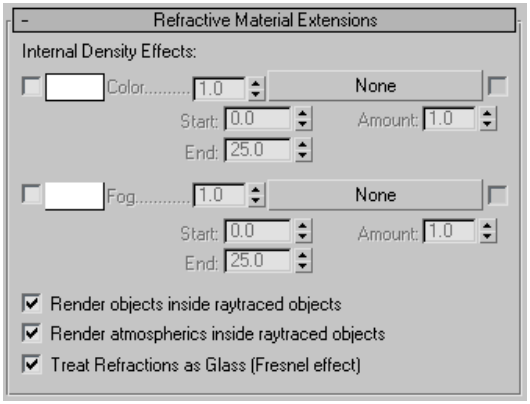
Bump Map Effect is active only when raytracing is enabled.

Raytrace: Refractive Material Extensions Rollout

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Raytrace > Refractive Material Extensions Rollout

With the controls in this rollout, you can fine-tune the effect of the *Raytrace map (page 2–1698)* on a material's refraction component.

Interface



Internal Density Effects

Color—With these controls, you can specify a transmission color based on thickness. The density color gives the appearance of color within the object itself, like tinted glass.

- **Enable—**Turns color density on or off.
- **Color swatch—**Displays a *Color Selector (page 1–161)*. Choose the transmission color.
- **Amount—**Controls the amount of density color. Reducing this value reduces the density color effect. Range=0 to 1.0. Default=1.0.
- **Color Map—**Assigns a map to the density color component. Use the check box to enable or disable the map.
- **Start and End—**A thin piece of tinted glass is mainly clear, while a thick piece of the same glass has more color. Start and End Distance, expressed in world units, controls help you simulate this effect. Start is the position in the object where the density color begins to appear (Default=0.0). End is the position in the object where the density color reaches its full Amount value. To have a lighter effect, increase the End value. To have a heavier effect, reduce the End value.

Fog—Density fog is also a thickness-based effect. It fills the object with a fog that is both opaque and self illuminated. The effect is like smoke trapped in a glass, or wax at the tip of a candle. Colored fog in tubular objects can resemble neon tubes.

- **Enable—**Turns fog on or off.
- **Color swatch—**Displays a *Color Selector (page 1–161)* for choosing the fog color.
- **Amount—**Controls the amount of density fog. Reducing this value reduces the density fog effect and makes the fog translucent. Range=0 to 1.0. Default=1.0.
- **Color Map—**Assigns a map to the fog component. Use the check box to enable or disable the map.
- **Start and End—**Start and End Distance controls, expressed in world units, adjust the fog effect based on the object's dimensions. Start is the position in the object where the density fog begins to appear (default=0.0). End is the position in the object where the density fog reaches its full Amount value. To have a lighter effect, increase the End value. To have a heavier effect, reduce the End value.

Render objects inside raytraced objects—Turns the rendering of objects inside raytraced objects on or off. Default=on.

Render atmospherics inside raytraced objects—Turns the rendering of atmospheric effects inside raytraced objects on or off. Atmospheric effects include fire, fog, volume light, and so on. Default=on.

Treat Refractions as Glass (Fresnel effect)— When on, applies a Fresnel effect to the refraction. This can add a bit of reflection to the refracting object, depending on the viewing angle of the object. When off, the object is refractive only. Default=on.

mental ray Shaders

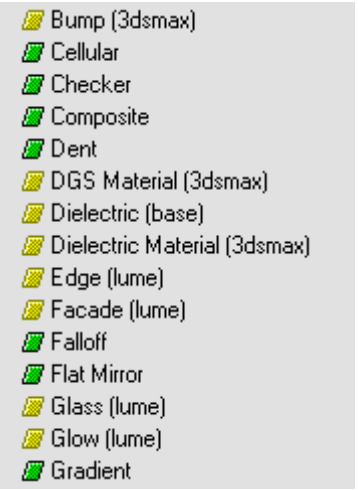
In mental ray, a shader is a function that calculates light effects. There can be shaders for lights, cameras (lens shaders), materials, shadows, and so on.

Note: In 3D modeling, the more common usage of "shader" signifies an algorithm that specifies how a surface responds to light. (The shaders for standard 3ds Max fall into this category.) With the mental ray renderer, "shader" has a more general sense of *any* algorithm used in rendering.

The *mental ray renderer (page 3–78)* can render most types of 3ds Max materials and maps. See *3ds Max Materials in mental ray Renderings (page 3–83)*. In addition, if you have enabled mental ray extensions (see *mental ray Preferences (page 3–837)*), you can apply a variety of shaders to materials. Materials designed for use with the mental ray renderer have specific components to which you can assign a shader. And for standard 3ds Max material types, the *mental ray Connection rollout (page 2–1461)* lets you add mental ray shading.

Warning: When you use the scanline renderer, mental ray shaders typically appear as black or white surfaces, or they are ignored entirely.

You assign a mental ray shader the same way you do a map. In the *Material/Map Browser (page 2–1412)*, mental ray shaders appear with a yellow icon, instead of the green icon used for maps.



mental ray maps in the browser's list are shown with yellow icons.

The shaders listed in the Browser depend on which type of shader component you have chosen in the Material Editor. For example, when you assign a Surface shader, the Browser lists a variety of shaders and standard 3ds Max maps. But when you assign a more special-purpose Contour shader, the Browser lists only those shaders that generate contour lines.

Note: Other kinds of special-purpose shaders include shaders for cameras and lights. Buttons to assign camera shaders are found on the Render Scene dialog's *Camera Effects rollout (page 3–101)*, and the buttons to assign light shaders are on a light object's *mental ray Light Shader rollout (page 2–1345)* (which appears only on the Modify panel, not the Create panel).

The shaders listed in the Browser come from several libraries that are provided with 3ds Max. Shaders created specifically for 3ds Max are described in this document. Shaders provided with the mental images or lume shader libraries have their own online documentation. The following topics link to the descriptions of specific shaders:

- *Custom Shaders for 3ds Max (page 2–1711)* links to descriptions of the shaders provided in the product-specific library, *3dsmax.mi*.
- *mental images Shader Libraries (page 2–1712)* links to descriptions of the shaders provided in the three standard mental ray libraries from mental images: *base.mi, contour.mi,* and *physics.mi*.
- *Shaders in the LumeTools Collection (page 2–1713)* includes links to descriptions of shaders in the lume library, *lume.mi*.
- The *Car Paint (page 2–1576)* material is also available as a shader, with the same set of parameters.
- The *mr Sun (page 2–1319)*, *mr Sky (page 2–1318)*, and *mr Physical Sky (page 2–1321)* shaders are components of the *mental ray Sun & Sky (page 2–1313)* solution.

Note: When you wire the parameters of an object whose material has mental ray shaders assigned, names of shader parameters might differ from those in the Material Editor interface. Also, parameters not supported by 3ds Max might appear as blanks in the wiring menu.

mental ray Shaders not Provided with 3ds Max

If your installation includes shader libraries other than those listed in this reference (whether obtained from a third-party source, or custom written), then the Browser might list those shaders as well. Documentation for third-party or custom shaders should come from the shader's provider.

Where Shaders Are Installed

Shaders provided with 3ds Max are installed in the subdirectory *\mentalray\shaders_standard*, below the 3ds Max root directory. The *\include* folder is for the MI include files, and the *\shaders* folder is for the DLLs.

Third-party shaders should *not* be installed in *\shaders_standard*. They should be installed either in *\shaders_3rdparty* or *\shaders_autoload*. If a third-party shader is present in *\shaders_autoload,* it is loaded automatically when you start 3ds Max. If a third-party shader is present in *\shaders_3rdparty,* you must update the file *3rdparty.mi* so it explicitly loads the shader. The MI file *\shaders_3rdparty\3rdparty.mi* contains comments that explain how to add load statements.

Note: Shaders listed in the MI file are loaded in reverse order: that is, from the bottom of the list to the top.

Custom Shaders for 3ds Max

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Pick a mental ray shader other than a mental images library shader or a lume shader.

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

For use with the *mental ray renderer (page 3–78)*, 3ds Max provides a number of custom shaders. These are the following:

3D Displacement Shader (mental ray) (page 2–1714)

Bump Shader (mental ray) (page 2–1716)

DGS Material Shader (mental ray) (page 2–1717)

Dielectric Material Shader (mental ray) (page 2–1719)

Environment Shader (mental ray) (page 2–1721)

Height Map Displacement Shader (mental ray) (page 2–1722)

Material to Shader (mental ray) (page 2–1723)

Shader List (mental ray) (page 2–1723)

UV Generator Shader (mental ray) (page 2–1724)

UV Coordinate Shader (mental ray) (page 2–1728)

XYZ Generator Shader (mental ray) (page 2–1729)

XYZ Coordinate Shader (mental ray) (page 2–1730)

mental images Shader Libraries

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Pick a mental ray shader other than a custom 3ds Max shader or a lume shader.

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The shaders provided with standard libraries from mental images are meant for use with the *mental ray renderer (page 3–78)*. There are three standard libraries: Base Shaders (*base.mi*), Physics Shaders (*physics.mi*), and Contour Shaders (*contour.mi*).

Note: In the mental image libraries, the names of base shaders have the prefix "mib_" and the names of contour shaders have the prefix "contour_". These prefixes don't appear in the 3ds Max user interface or in the table that follows. (Names of physics shaders have no conventional prefix.)

The following table lists the mental images library shaders provided with 3ds Max.

Tip: When you follow a link to the documentation for mental images library shaders, scroll up a bit in your browser. The links tend to go directly to the shader's declaration code, and often there are some introductory paragraphs directly above the code. If the link goes to the beginning of a section, scroll down instead.

Shader	Library
Ambient/Reflective Occlusion	base
Combi	contour
Contour Composite	contour
Contour Contrast Function Levels	contour
Contour Only	contour
Contour PS (PostScript)	contour
Contour Store Function	contour
Curvature	contour
Depth Fade	contour
DGS Material Photon	physics
Dielectric	base
Factor Color	contour
Layer Thinner	contour
Light Infinite	base
Light Point	base
Light Spot	base
Opacity	base
Parti Volume	physics
Photon Basic	base
Reflect	base
Refract	base
Shadow Transparency	base
Simple	contour
Texture Remap	base
Texture Rotate	base
Texture Wave	base
Transmat	physics
Transmat Photon	physics
Transparency	base
Two Sided	base

Shader	Library
Width From Color	contour
Width From Light	contour
Width From Light Dir	contour

Note: You can also access the mental images shader help by choosing Help > Additional Help, highlighting Mental Ray Standard Library in the list, and then clicking Display Help.

Shaders in the LumeTools Collection

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Pick a shader with "(lume)" in its name.

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The LumeTools Collection of shaders provide a variety of naturalistic effects when used with the *mental ray renderer (page 3–78)*. In the Material/Map Browser, the name of these shaders is followed by "(lume)." These are the lume shaders provided with 3ds Max:

- Beam
- Distortion
- Edge and Edge Shadow
- Facade
- Glass
- Glow
- Landscape
- Metal
- Mist
- Night
- Ocean
- Stain
- Submerge
- Translucency
- Water Surface
- Wrap Around
- Wet-Dry Mixer

Note: You can also access the lume shader help by choosing Help > Additional Help, highlighting Lume Tools in the list, and then clicking Display Help.

Connect Parameter to Shader Dialog (mental ray)

Material Editor > mental ray Connection rollout > Click a shader button. > Material/Map Browser > Pick a mental ray shader that returns multiple values.

Material Editor > Shader controls > Click a shader button. > Material/Map Browser > Pick a mental ray shader that returns multiple values.

Material Editor > DGS material > Click a shader button. > Material/Map Browser > Pick a mental ray shader that returns multiple values.

Material Editor > Shader controls > Click a shader return parameter button (to the right of the main shader button).

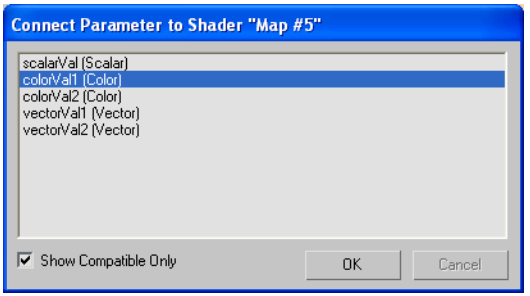
Note: The mental ray Connection rollout is available in the Material Editor only if you have enabled the mental ray extensions by using the mental ray Preferences panel. In addition, shaders don't appear unless the mental ray renderer is the currently active renderer.

Some mental ray shaders return multiple values. If you choose one of these, a Connect Parameter To Shader dialog appears. Choose one of the return values in the list, and then click OK.

Important: *UV Coordinates (page 2–1728)* and *XYZ Coordinates (page 2–1730)* are the only shaders with multiple return values provided with 3ds Max. You might encounter multiple return values in shaders provided with other shader libraries or custom shader code.

The components for some mental ray materials and shaders can be assigned other shaders. For these components, the main shader button is accompanied on the right by a small button. If no shader is assigned, or the shader assigned has only a single return value, the button is disabled and shows a dot in the middle. If the shader assigned can return multiple values, text appears in this button, and a tooltip shows the parameter name. Clicking the button displays the Connect Parameter To Shader dialog, allowing you to change the parameter being used.

Interface



List of return values—Lists the parameters that the shader returns. The type of each parameter is indicated in parentheses, following the parameter's name.

Show Compatible Only—When on the list shows only the return values whose type is compatible with the component the shader is assigned to. When off, shows all return values from the shader, whether they are compatible or not. Default=on.

3ds Max Custom Shaders

3D Displacement Shader (mental ray)

Material Editor > mental ray Connection rollout > For the Displacement component, turn off the lock button. > Click the button for the Displacement component. > Material/Map Browser > 3D Displacement (3dsmax)

Material Editor > mental ray material > Click the button for the Displacement component. > Material/Map Browser > 3D Displacement (3dsmax)

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

A 3D Displacement shader displaces the geometry of surfaces. The effect is similar to displacement mapping of a standard material. You can apply mental ray displacement to any kind of object, unlike standard displacement mapping, which is restricted to surface models (meshes, patches, polys, and NURBS surfaces).

Displaced surfaces are smooth if the displaced polygons share normals; otherwise, the displaced surfaces are faceted. Also, unless normals are shared, faces can become separated in the displaced mesh. To prevent this, make sure adjacent surfaces belong to the same shading group.

When the mental ray renderer is the active renderer, mental ray displacement is the only displacement method used, unless your scene includes a *Displace modifier (page 1–629)*, which always uses standard 3ds Max displacement.

Tip: Before you render, you can disable or enable displacement by using the Displacement toggle in the Options group on the *Common Parameters rollout (page 3–27)*.

Global settings for the mental ray displacement method are in the Displacement group on the

Shadows And Displacement rollout (page 3–114) of the Render Scene dialog's Renderer panel.

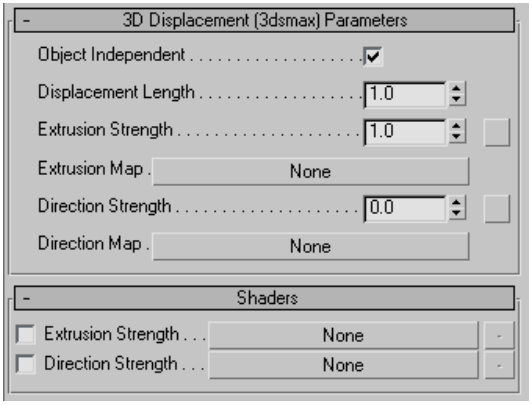
See also

mental ray Displacement (page 3–96)

mental ray Connection Rollout (page 2–1461)

mental ray Material (page 2–1544)

Interface



3D Displacement (3dsmax) Parameters rollout

Note: The button to the right of the Factor and Direction Strength controls is a shortcut shader button. Clicking one of these buttons displays the *Material/Map Browser (page 2–1412)* so you can assign a shader to this component. When a map or a shader has been assigned to a component, this button displays the letter "M," and the comparable button on the Shaders rollout displays the map or shader name.

Object Independent—When on, the displacement effect is independent of the size of the object's bounding box. When off, the displacement effect is scaled according to the size of the object. Default=on.

Scaling the displacement based on object size is the standard behavior for regular 3ds Max displacement mapping.

Displacement Length—This is the length of displacement when Object Independent is on, the extrusion map is at 100 per cent (white) and the Extrusion Strength equals 1.0. Lower gray levels in the extrusion map, or other values of Extrusion Strength, scale the amount of displacement. When Object Independent is off, this value is disregarded. Default=1.0.

Extrusion Strength—Controls the height of the displacement. This value is a multiplier: at the default value of 1.0, the map's effect is unchanged. Greater values increase the effect of the map, and lower values decrease it. Default=1.0.

Extrusion Map—Click to display the *Material/Map Browser (page 2–1412)* and choose a map to use for the displacement. Displacement maps apply the gray scale of the map to generate the displacement. Lighter colors in the 2D image push outward more strongly than darker colors, resulting in a 3D displacement of the geometry.

Direction Strength—Controls the strength of the direction shader. Default=0.0.

Important: Adding a direction shader has no visible effect unless you set Direction Strength to be greater than its default value of zero. (Direction Strength values less than zero have no effect.)

Direction Map—Click to display the *Material/Map Browser (page 2–1412)* and choose a shader to use for the map direction. The direction of the displacement is perturbed according to the RGB values of the shader output or map pixels. Red values offset in the U axis, Green values offset in V, and Blue values offset in W (using the object-local UVW coordinates).

Shaders rollout

The controls on this rollout let you assign a map or shader to the Factor or Direction Strength parameters. Click the button for a component to display the *Material/Map Browser (page 2–1412)*

and assign the map or shader. Use the toggle at the left to turn the effect of the map off or on.

The button to the right of each main shader button is for shaders that can return multiple parameters. If a shader that returns multiple parameters is assigned to the component, the button's tooltip shows the parameter name. Clicking the button displays a *Connect Parameter To Shader dialog (page 2–1713)*, which lets you change which parameter is being used.

Important: *UV Coordinate (page 2–1728)* and *XYZ Coordinate (page 2–1730)* are the only shaders with multiple return values provided with 3ds Max. You might encounter multiple return values in shaders provided with other shader libraries or custom shader code.

Bump Shader (mental ray)

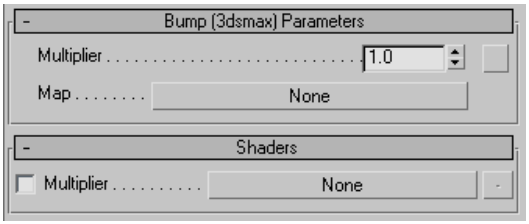
Material Editor > mental ray material > Click the button for the Surface or Bump component. > Material/Map Browser > Bump (3dsmax)

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The Bump shader provides bump mapping for the mental ray renderer. Bumps are created by perturbing face normals before the object is rendered, using the same method as *bump mapping (page 2–1506)* for the scanline renderer.

Warning: Although you can assign a Bump shader to the Surface component, if you assign *only* a Bump shader, the surface will render as black. For the Surface component, use the Bump shader in a *Shader List (page 2–1723)*, or for the *mental ray material (page 2–1544)*, use the Bump component itself.

Interface



Bump (3dsmax) Parameters rollout

Multiplier—Adjust the bump effect by multiplying the map values. Negative Multiplier values reverse the bump effect: hollow areas now protrude, and raised areas become hollow. Default=1.0.

Map—Click to display the *Material/Map Browser (page 2–1412)* and choose a map to use for generating bumps. Bump mapping uses the intensity of the map to affect the surface of the material. The intensity affects the apparent bumpiness of the surface: white areas protrude, and black areas recede.

Shaders rollout

The controls on this rollout let you assign a map or shader to the Multiplier parameter. Click the button for a component to display the *Material/Map Browser (page 2–1412)* and assign the map or shader. Use the toggle at the left to turn the effect of the map off or on.

The button to the right of the main shader button is for shaders that can return multiple parameters. If a shader that returns multiple parameters is assigned to the component, the button's tooltip shows the parameter name. Clicking the button displays a *Connect Parameter To Shader dialog (page 2–1713)*, which lets you change which parameter is being used.

Important: *UV Coordinate (page 2–1728)* and *XYZ Coordinate (page 2–1730)* are the only shaders with multiple return values provided with 3ds Max. You might

encounter multiple return values in shaders provided with other shader libraries or custom shader code.

DGS Material Shader (mental ray)

Material Editor > mental ray Connection rollout > Unlock the Surface or Photon component. > Click the shader button for the Surface or Photon component. > Material/Map Browser > DGS Material (3dsmax)

Material Editor > mental ray material > Click the button for the Surface or Photon component. > Material/Map Browser > DGS Material (3dsmax)

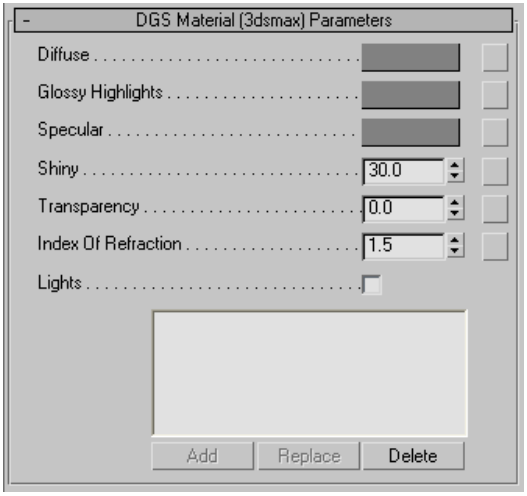
Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

DGS stands for Diffuse, Glossy, Specular. This shader is a mental ray phenomenon (a scripted shader tree) that provides a physically accurate simulation of a surface. With the mental ray Connection rollout of a basic 3ds Max material, or a mental ray material, you can assign the DGS Material shader to either the Surface or Photon component.

Despite the similarity in name, this shader is distinct from the *DGS material (page 2–1580)*. Both, in turn, provide a custom 3ds Max interface to the "DGS Material Photon" shader that is part of the mental images physics library.

Interface

Parameters rollout



Note: The button to the right of the first six controls is a shortcut shader button. Clicking one of these buttons displays the *Material/Map Browser (page 2–1412)* so you can assign a shader to this component. When a map or a shader has been assigned to a component, this button displays the letter "M," and the comparable button on the Shaders rollout displays the map or shader name.

Diffuse—Click the color swatch to display a *Color Selector (page 1–161)* and change the material's diffuse color.

Glossy Highlights—Click the color swatch to display a Color Selector and change the color of glossy highlights.

Specular—Click the color swatch to display a Color Selector and change the color of mirror reflections. When the specular color is white, the material is 100 percent reflective, like a mirror. When the specular color is black, the material does not reflect any of its surroundings.

Shiny—Sets the width of glossy highlights. The larger this value, the smaller the highlights. Default=30.0.

Transparency—Specifies the transparency. The effective range of Transparency is from 0.0 to 1.0. At 0.0 the material is fully opaque. At 1.0 it is fully transparent. Default=0.0.

Warning: You can set the value of Transparency to be greater than 1.0, but this has no effect. An anomaly of the user interface for shaders in the mental ray and lume libraries, is that spinner values are not "clamped" to lie within their effective ranges, as they are for controls in 3ds Max.

The value of Transparency also indirectly specifies the reflectivity of the material, which is calculated as 1.0 minus the Transparency value.

Index Of Refraction—Specifies the IOR. In the physical world, the IOR results from the relative speeds of light through the transparent material and the medium the eye or the camera is in. Typically this is related to the object's density. The higher the IOR, the denser the object. Default=1.5.

See *Extended Parameters Rollout (Standard Material) (page 2–1471)* for a list of IOR values for commonly encountered materials.

Note: When the IOR equals 1.0, there is no refraction, and calculating the transparency can take less time than when the material is refractive.

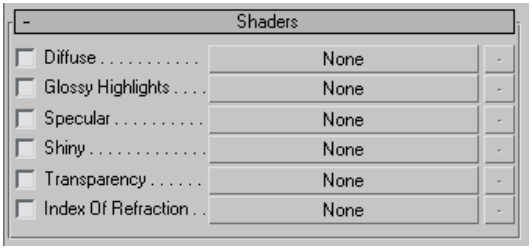
Lights—When on, the material is illuminated only by those lights specified in the list. When Lights is turned off, all lights in the scene affect the material. Default=off.

The remaining light controls are unavailable unless Lights is turned on.

- **List of lights—**Displays the lights you have chosen to illuminate this material.
- **Add—**Adds a light to the list. Click Add to turn it on, then click the light object in a viewport.
- **Replace—**Replaces a light in the list. Highlight a light's name in the list, click Replace to turn it on, then click the replacement light object in a viewport.
- **Delete—**Deletes a light from the list. Highlight a light's name in the list, then click Delete.

Shaders rollout

The controls on this rollout let you assign a map or shader to one of the basic parameters of the DGS Material shader. This is comparable to mapping a component of a standard material; by adding shaders, you can create a shader tree that generates complex effects.



Click the button for a component to display the *Material/Map Browser (page 2–1412)* and assign the map or shader. Use the toggle at the left to turn the effect of the map off or on.

 The button to the right of each main shader button is for shaders that can return multiple parameters. If a shader that returns multiple parameters is assigned to the component, the button's tooltip shows the parameter name. Clicking the button displays a *Connect Parameter To Shader dialog (page 2–1713)*, which lets you change which parameter is being used.

Important: *UV Coordinates (page 2–1728)* and *XYZ Coordinates (page 2–1730)* are the only shaders with multiple return values provided with 3ds Max. You might encounter multiple return values in shaders provided with other shader libraries or custom shader code.

For all the DGS Material shader components, the available mental ray shaders are the same:

Shader	Library
Bump (page 2–1716)	3ds Max
DGS Material (this shader)	3ds Max
Dielectric	base
Dielectric Material (page 2–1719)	3ds Max
Edge	lume
Facade	lume
Glass	lume
Glow	lume
Landscape	lume
Material to Shader (page 2–1723)	3ds Max
Metal	lume
Ocean	lume
Opacity	base
Reflect	base
Refract	base
Shader List (page 2–1723)	3ds Max
Stain	lume
Translucency	lume
Transmat	physics
Transparency	base
Two Sided	base
UV Generator (page 2–1724)	3ds Max
Water Surface	lume
Wet-Dry Mixer	lume
XYZ Generator (page 2–1729)	3ds Max

Dielectric Material Shader (mental ray)

Material Editor > mental ray Connection rollout > Unlock the Surface or Photon component. > Click the shader button for the Surface or Photon component. > Material/Map Browser > Dielectric Material (3dsmax) or Dielectric Material Photon (3dsmax)

Material Editor > mental ray material > Click the button for the Surface or Photon component. > Material/Map Browser > Dielectric Material (3dsmax) or Dielectric Material Photon (3dsmax)

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The Dielectric Material shader creates transparent, refractive materials that are physically accurate. A dielectric material, such as glass, is a material whose surface transmits most light that strikes it at angles close to perpendicular (90 degrees), but reflects most light that strikes at glancing angles (close to zero degrees).

When applied to the Surface component, this shader affects the surface's appearance. When applied to the Photon component, it affects its photon behavior for caustics and global illumination. (The Glass material is a mental ray phenomenon (a scripted shader tree) that is equivalent to a *mental ray material (page 2–1544)* with a Dielectric Material shader assigned to both its Surface and Photon components, with the parameter settings identical for both.)

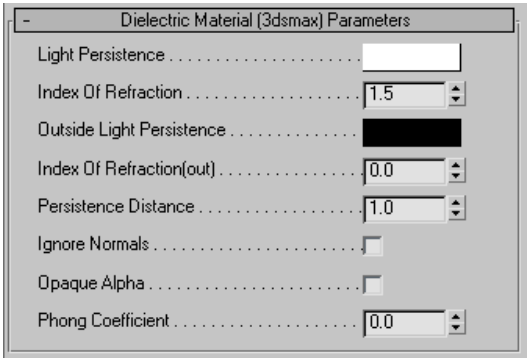
Note: This material does not use a shadow shader, so shadows will always be opaque unless you use a Dielectric Material shader for the Photon component, and generate caustics when you render.

Adjacent Refractive Materials

Two controls, Outside Light Persistence and Index Of Refraction (Out), are for situations where you are modeling two adjacent refractive materials. Consider a drink in a martini glass. The glass

has an index of refraction (IOR) of 1.5, while the alcohol in the glass has an IOR of about 1.3. To create a physically accurate model of this situation, use *three* glass materials: one for the glass itself, one for the alcohol, and a third material for the surfaces where they touch each other. For this third material, set the "inside" IOR to 1.3, and the outside IOR to 1.5.

Interface



Light Persistence—In conjunction with the Persistence Distance, controls the percentage of light that the volume transmits. For example, if the color is set to R=G=B=0.5 and the Persistence Distance is set to 2.0, then objects with a thickness of 2.0 units will appear 50 per cent transparent. Default=white (R=G=B=1.0).

Because transparency depends on the thickness of the object, objects with varying thickness show different transparency depending on the angle from which they are viewed.

Index Of Refraction—Specifies the Index Of Refraction (IOR). In the physical world, the IOR results from the relative speeds of light through the transparent material and the medium the eye or the camera is in. Typically this is related to the object's density. The higher the IOR, the denser the object. Default=1.5.

See *Extended Parameters Rollout (Standard Material) (page 2–1471)* for a list of IOR values for commonly encountered materials.

Outside Light Persistence—In conjunction with the Persistence Distance, controls the percentage of light transmitted on the other side of a surface. When set to the default of black, this control has no effect. See the section "Adjacent Refractive Materials," above. Default=black (R=G=B=0.0).

Index Of Refraction (out)—Sets the IOR on the other side of a surface. When set to the default of zero, this control has no effect. See the section "Adjacent Refractive Materials," above. Default=0.0.

Persistence Distance—In conjunction with the Light Persistence color, controls the percentage of light that the volume transmits. It is the distance at which light transmission is reduced to the percentage specified by the Light Persistence RGB values. Default=1.0.

If you specify an Outside Light Persistence color, that setting also uses the Persistence Distance.

Ignore Normals—When on, the renderer does not use normals to decide whether a light ray is entering or leaving the object. Normally, the shader uses normals to decide whether a ray is entering or leaving an object. (It is entering if the normal points toward the ray, leaving if the normal points away from the ray.) This can present a problem for rendering objects whose normals are not unified. When Ignore Normals is on, the shader decides whether a ray is entering or leaving the object by counting the number of times the ray has intersected the object. Default=off.

Opaque Alpha—When on, refracted rays that touch the environment don't generate a transparent alpha value. (This is how 3ds Max usually treats the environment.) When off, refracted rays that touch the environment render a transparent alpha value, which can help if you plan to use the rendering as part of a composite. Default=off.

Phong Coefficient—When greater than zero, generates Phong highlights on the surface. The highlights appear in the sample slot. In general this value must be greater than 10 for highlights to be apparent. Default=0.0.

Environment Shader (mental ray)

Material Editor > mental ray Connection rollout > Assign a shader to the Environment component. > Material/Map Browser > Environment (3dsmax)

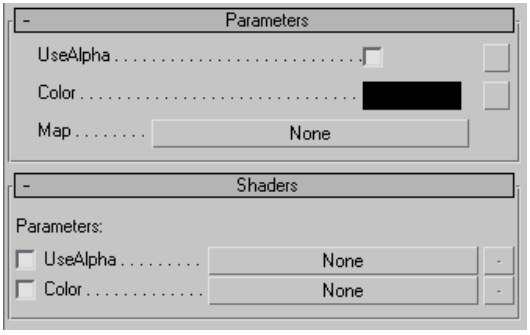
Material Editor > mental ray material > Assign a shader to the Environment component. > Material/Map Browser > Environment (3dsmax)

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The Environment shader lets you specify an environment that is local to the material. Controls for the Environment shader are similar to those for a scene's environment on the Render Scene dialog's Environment panel. However, the local Environment shader doesn't affect the scene background. Instead, it provides an environment that the material can reflect or refract.

If an environment map is present, it generates the reflections or refractions, and they are not ray traced.

Interface



Parameters rollout

Note: The button to the right of the UseAlpha and Color controls is a shortcut shader button. Clicking one of these buttons displays the *Material/Map Browser (page 2–1412)* so you can assign a shader to this component. When a map or a shader has been assigned to a component, this button displays the letter "M," and the comparable button on the Shaders rollout displays the map or shader name.

UseAlpha—When on, uses the map's alpha channel, if it has one. The alpha channel specifies those portions of the map that are transparent or translucent. Default=off.

Color—Click the color swatch to display a *Color Selector (page 1–161)* and choose a color to use as the environment.

Map—Click the button to display a *Material/Map Browser (page 2–1412)* and choose a map to use as the environment.

Shaders rollout

The controls on this rollout let you assign a map or shader to the UseAlpha and Color parameters. Click the button for a component to display the *Material/Map Browser (page 2–1412)* and assign the map or shader. Use the toggle at the left to turn the effect of the map off or on.

 The button to the right of each main shader button is for shaders that can return multiple parameters. If a shader that returns multiple parameters is assigned to the component, the button's tooltip shows the parameter name. Clicking the button displays a *Connect Parameter To Shader dialog (page 2–1713)*, which lets you change which parameter is being used.

Important: *UV Coordinates (page 2–1728)* and *XYZ Coordinates (page 2–1730)* are the only shaders with multiple return values provided with 3ds Max. You might

encounter multiple return values in shaders provided with other shader libraries or custom shader code.

Height Map Displacement Shader (mental ray)

Material Editor > mental ray Connection rollout > For the Displacement component, turn off the lock button. > Click the button for the Displacement component. > Material/Map Browser > Height Map Displacement (3dsmax)

Material Editor > mental ray material > Click the button for the Displacement component. > Material/Map Browser > Height Map Displacement (3dsmax)

Note: Shaders don't appear unless the mental ray renderer is the currently active renderer.

The Height Map Displacement shader displaces the geometry of surfaces, and is specifically intended for use with height maps generated by normal mapping; see *Creating and Using Normal Bump Maps (page 3–150)*.

Important: When applying a material containing this map to an object, the mental ray Displacement > Smoothing option must be off. If such materials are applied to *all* objects in the scene, you can *turn off Smoothing globally*. Otherwise, turn off smoothing for each object whose material uses a height map via the Object Properties > *mental ray panel (page 1–126)* (turn off Use Global Settings and then turn off Smoothing).

Tip: Avoid the temptation to apply MeshSmooth to a model when creating a height map for it. This changes the shape of the model so the height values will not be correct. The low-res model must have exactly the same shape when the map is created and when it is used for displacement. Also, MeshSmooth does not use the same algorithm as the mental ray displacement smoothing, so using both forms of smoothing won't work perfectly. The best results are obtained by not smoothing the low-res model when the map is created and also not using mental ray smoothing.

Also, avoid using a paint program to modify the height map. The values in the height map depend on the shape of both the low-res and high-res models, and it's easy to damage the mathematical accuracy. If you paint any changes onto the map, you must be careful to preserve the faceted look, and avoid the temptation to blur away the facets.

You might try painting in Additive and Subtractive mode to add or subtract to the displacement, because Normal mode will set a fixed displacement, making it difficult for an artist to control the result.

See also

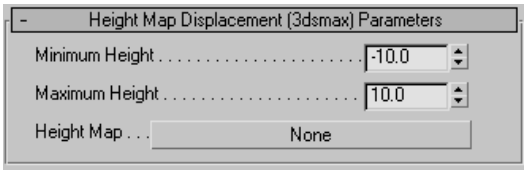
3D Displacement Shader (mental ray) (page 2–1714)

mental ray Displacement (page 3–96)

mental ray Connection Rollout (page 2–1461)

mental ray Material (page 2–1544)

Interface



Height Map Displacement (3dsmax) Parameters rollout

Be sure to enter the same values for Minimum and Maximum Height as the equivalents on the Projection Options dialog, as specified below.

Minimum Height—The "Min Height" value specified on the *Render to Texture: Projection Options dialog (page 3–165)* when creating the height map. Default=-10.0.

Maximum Height—The "Max Height" value specified on the *Render to Texture: Projection*

Options dialog (page 3–165) when creating the height map. Default=10.0.

Height Map—The height map itself (usually a bitmap).

Material to Shader (mental ray)

Material Editor > mental ray Connection rollout > Assign a shader. > Material/Map Browser > Material to Shader

Material Editor > mental ray or DGS material > Assign a shader. > Material/Map Browser > Material to Shader

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

Lets you use a regular 3ds Max material as a shader. Depending on the component to which this shader is assigned (Surface, Shadow, Displacement, Volume, and so on), the mental ray renderer uses the appropriate material component.

For example, if you want a mental ray material's Surface component to look like a standard material you have, assign Material To Shader as the Surface shader, and then assign it the standard material.

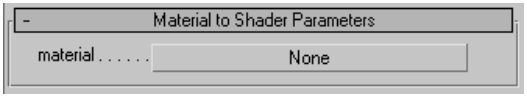
Note: Material To Shader doesn't work as an environment background. Use the original 3ds Max material, instead.

Tip: To edit the material assigned to Material To Shader, you can drag the button to an unused sample slot in the Material Editor (be sure to choose Instance when prompted). Or you can follow these steps:

1. In an unused sample slot, create the material and adjust its settings.
2. Save the material to a library.
3. Assign the Material To Shader to its component.
4. When you click the Material To Shader's shader button, browse from the library and load the material you prepared in advance.

If you need to further adjust the material, you can repeat these steps (without having to reassign the Material To Shader).

Interface



Material button—Click to display the *Material/Map Browser (page 2–1412)* and choose the material to use for shading.

Shader List (mental ray)

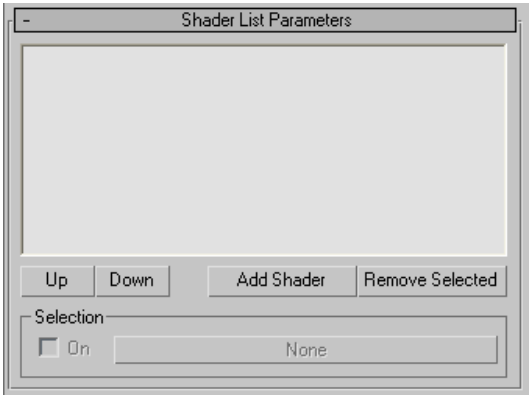
Material Editor > mental ray Connection rollout > Assign a shader. > Material/Map Browser > Shader List

Material Editor > mental ray, DGS, or Glass material > Assign a shader. > Material/Map Browser > Shader List

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The Shader List shader provides an interface for constructing a mental ray shader list. A shader list combines the effect of multiple shaders: each shader is called in turn, the first one's output being treated as input to the next, and so on.

Interface



List of shaders—Shows the names of the shaders in the list. Highlight a shader's name to alter its position in the list, or to access its parameters.

Each active shader in the list is called in order, from top to bottom.

Up—Moves the selected shader up in the list.

Down—Moves the selected shader down in the list.

Add Shader—Displays a *Material/Map Browser (page 2–1412)* so you can choose a shader to add to the list.

Remove Selected—Removes the selected shader from the list.

Selection group

On—When on, the shader is active. When off, the shader is inactive and isn't called. You can use this toggle to disable a shader without removing it from the list entirely.

Shader button—Shows the name of the currently selected shader. Click the button to view that shader's parameters in the Material Editor.

 When you are done adjusting an individual shader's parameters, you can click Go To Parent to return to the Shader List Parameters rollout.

UV Generator

Material Editor > Any shader with a Coords parameter (or other vector value). > Click the shader button. > Material/Map Browser > UV Generator (3dsmax)

Material Editor > mental ray Connection rollout > Unlock the Surface component and click the shader button. > Material/Map Browser > UV Generator (3dsmax)

Material Editor > DGS material > Assign a shader to any component. > Material/Map Browser > UV Generator (3dsmax)

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The UV Generator shader returns 2D mapping coordinates. You can use its settings to adjust a 2D map. The parameters for this shader are equivalent to parameters on the *Coordinates rollout for 2D maps (page 2–1625)*.

Interface

The interface for this shader consists of two rollouts:

UV Generator Parameters Rollout (page 2–1725)

Shaders Rollout (UV Generator) (page 2–1727)

UV Generator Parameters Rollout

Material Editor > Any shader with a Coords parameter (or other vector value). > Click the shader button. > Material/Map Browser > UV Generator (3dsmax) > UV Generator (3dsmax) Parameters rollout

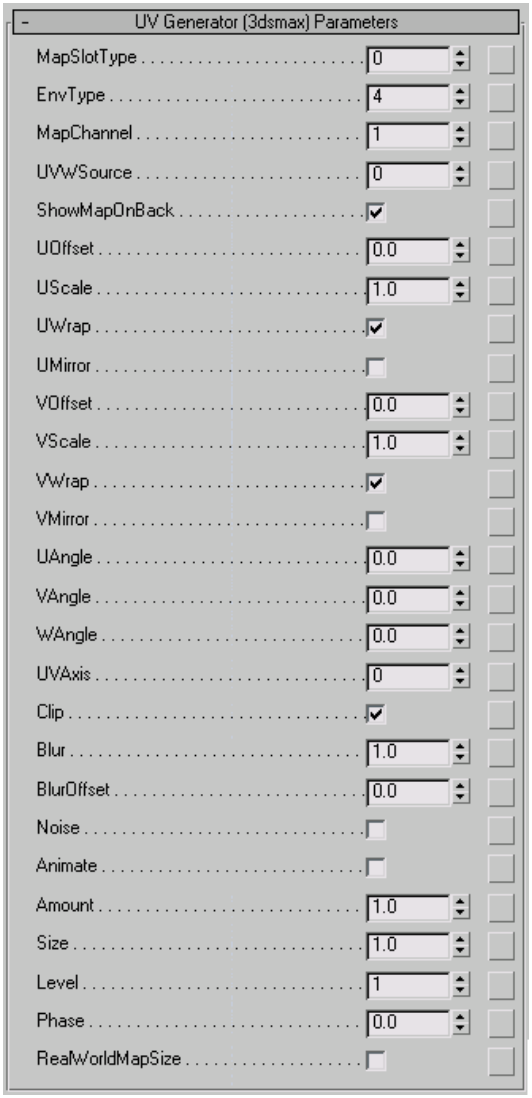
Material Editor > mental ray Connection rollout > Unlock the Surface component and click the shader button. > Material/Map Browser > UV Generator (3dsmax) > UV Generator (3dsmax) Parameters rollout

Material Editor > DGS material > Assign a shader to any component. > Material/Map Browser > UV Generator (3dsmax) > UV Generator (3dsmax) Parameters rollout

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The parameters for this shader are equivalent to parameters on the *Coordinates rollout for 2D maps (page 2–1625)*.

Interface



MapSlotType—Chooses whether the map is applied as an environment map or a texture map. Default=0 (Texture).

- 0 is for a Texture map.
- 1 is for an Environment map.

EnvType—If the map is applied as an environment map, this value chooses the mapping type. If

the map is applied as a texture map, this value is ignored. Default=4 (screen).

- 1 is for Spherical.
- 2 is for Cylindrical.
- 3 is for Shrink-Wrap.
- 4 is for Screen.

Screen projection projects as a flat backdrop in the scene. See *UVW Map Modifier (page 1–922)* for a description of spherical, cylindrical, and shrink-wrap projection.

MapChannel—If the map is applied as a texture map and UVWSource is set to 0 (Explicit), this value sets the map channel index; otherwise, this value is ignored. Range=1 to 99. Default=1.

UVWSource—If the map is applied as a texture map, this value chooses the source of UVW coordinates. If the map is applied as an environment map, this value is ignored. Default=0 (Explicit).

- 0 is for Explicit.

  When Explicit is the source, use MapChannel to set the specific map channel.

- 1 is for Object XYZ.

  Object XYZ uses planar mapping based on the object's local coordinates (disregarding the pivot point location). For rendering purposes, planar mapping doesn't project through to the back of the object unless you turn on ShowMapOnBack.

- 2 is for World XYZ.

  World XYZ uses planar mapping based on the scene's world coordinates (disregarding the object's bounding box). For rendering purposes, planar mapping doesn't project through to the back of the object unless you turn on ShowMapOnBack.

ShowMapOnBack—When on, planar mapping (Planar from Object XYZ, or with the UVW Map modifier) projects through to render on the back of the object. When off, planar mapping doesn't render on the object's back. Default=on.

This toggle is available only when Tiling is off in both dimensions. Its effect is visible only when you render the scene.

Note: In viewports, planar mapping always projects to the back of the object, whether Show Map On Back is turned on or not. To override this, turn off Tiling.

UOffset—Changes the U position of the map in *UV coordinates (page 3–1028)*. The map moves in relation to its size. Default=0.0.

For example, if you want to shift the map its full width to the left, and half its width downward from its original position, you enter -1 in the U Offset field and 0.5 in the V offset field.

UScale—Determines the number of times the map is *tiled (page 3–1022)* (repeated) along the U axis. Default=1.0.

UWrap—Turns tiling on or off in the U axis. Default=on.

UMirror—*Mirrors (page 3–1022)* the map left-to-right along the U axis. Default=off.

VOffset—Changes the V position of the map in UV coordinates. Default=0.0.

VScale—Determines the number of times the map is tiled (repeated) along the V axis. Default=1.0.

VWrap—Turns tiling on or off in the V axis. Default=on.

VMirror—Mirrors the map top-to-bottom along the V axis. Default=off.

UAngle, VAngle, and WAngle—Rotate the map about the U, V, or W axis (in degrees). Default=0.0.

UVAxis—Changes the mapping coordinate system used for the map. The default UV coordinates project the map onto the surface like a slide projector. The VW and WU coordinates rotate

the map so that it is perpendicular to the surface. Default=0 (UV).

- 0 is for UV.
- 1 is for VW.
- 2 is for WU.

Clip—When on, UVs are clipped. When off, UVs are wrapped. Default=on.

Blur—Affects the sharpness or blurriness of the map based on its distance from the view. The farther away the map is, the greater the blurring. The Blur value blurs maps in world space. Blur is primarily used to avoid *aliasing (page 3–907)*. Default=1.0.

BlurOffset—Affects the sharpness or blurriness of the map without regard to its distance from the view. Blur Offset blurs the image itself in object space. Use this option when you want to soften or defocus the details in a map to achieve the effect of a blurred image. See *Blur / Blur Offset (page 3–918)*. Default=0.0

Noise—When on, noise settings affect the map. When off, no noise is applied. Default=off.

Animate—Determines if the noise effect is animated. This parameter must be turned on if you intend to animate the noise. Default=off.

Amount—Sets the strength of the fractal function, expressed as a percentage. If the amount is 0 there is no noise. If the amount is 100 the map becomes pure noise. Default-1.0.

Size—Sets the scale of the noise function relative to geometry. At very small values, the noise effect becomes white noise. At large values, the scale can exceed the scale of the geometry, in which case it has little or no effect. Range=0.001 to 100. Default=1.0.

Level—Or iterations: the number of times the function is applied. The effect of the level is dependent on the Amount value. The stronger the amount, the greater the effect of increasing the Level value. Range=1 to 10. Default=1.

Phase—Controls the speed of the animation of the noise function. Default=0.0.

RealWorldMapSize—Controls the scaling method used for texture mapped materials that are applied to the object. The scaling values are controlled by the Use Real-World Scale settings found on the applied material's Coordinates rollout. Default=on.

Shaders Rollout (UV Generator)

Material Editor > Any shader with a Coords parameter (or other vector value). > Click the shader button. > Material/Map Browser > UV Generator (3dsmax) > Shaders rollout

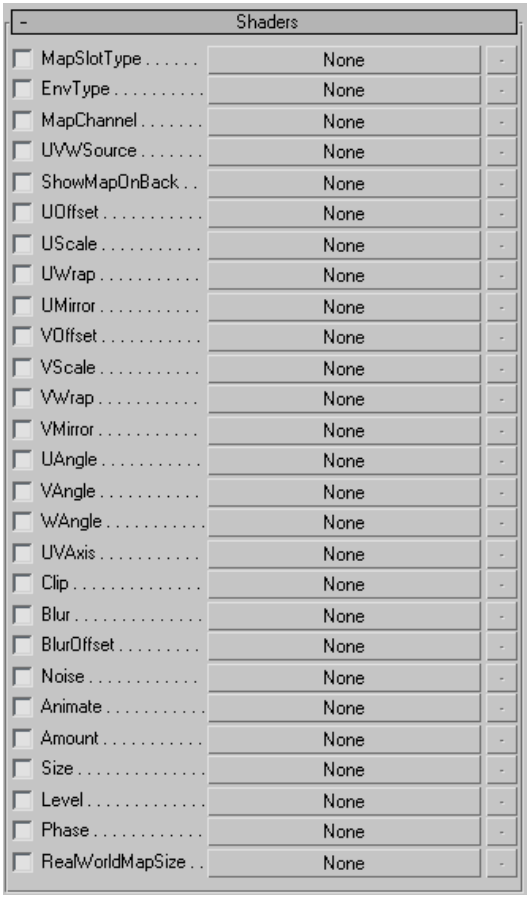
Material Editor > mental ray Connection rollout > Unlock the Surface component and click the shader button. > Material/Map Browser > UV Generator (3dsmax) > Shaders rollout

Material Editor > DGS material > Assign a shader to any component. > Material/Map Browser > UV Generator (3dsmax) > Shaders rollout

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The controls on this rollout let you assign a map or shader to one of the basic parameters of the UV Generator shader. This is comparable to mapping a component of a standard material; by adding shaders, you can create a shader tree that generates complex effects.

Interface



 The button to the right of each main shader button is for shaders that can return multiple parameters. If a shader that returns multiple parameters is assigned to the component, the button's tooltip shows the parameter name. Clicking the button displays a *Connect Parameter To Shader dialog (page 2–1713)*, which lets you change which parameter is being used.

Important: *UV Coordinates (page 2–1728)* and *XYZ Coordinates (page 2–1730)* are the only shaders with multiple return values provided with 3ds Max. You might encounter multiple return values in shaders provided with other shader libraries or custom shader code.

UV Coordinate Shader (mental ray)

Material Editor > Any shader with a Coords parameter or other vector value. > Click the shader button. > Material/Map Browser > UV Coordinate (3dsmax)

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

This shader is identical to the *UV Generator shader (page 2–1724)*, except that it returns two values instead of one.

- The UV coordinate (same value as UV Generator would return)

  This value is named UV.

- The UV coordinate's derivative

  This value, also a vector, is named dUV.

When you choose this shader, a *Connect Parameter To Shader dialog (page 2–1713)* is displayed, prompting you to choose which of the two values to use. You can later change the choice of value by using the "dot" button to the right of the main shader button.

Interface

The UV Coordinate shader has the same parameters as the UV Generator shader. See *UV Generator Parameters Rollout (page 2–1725)* for a description of the basic settings.

XYZ Generator

Material Editor > Any shader with a Coords parameter (or other vector value). > Click the shader button. > Material/Map Browser > XYZ Generator (3dsmax)

Material Editor > mental ray Connection rollout > Unlock the Surface component and click the shader button. > Material/Map Browser > XYZ Generator (3dsmax)

Material Editor > DGS material > Assign a shader to any component. > Material/Map Browser > XYZ Generator (3dsmax)

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The XYZ Generator shader returns 3D mapping coordinates. You can use its settings to adjust a 3D map. The parameters for this shader are equivalent to parameters on the *Coordinates rollout for 3D maps (page 2–1663)*.

Interface

The interface for this shader consists of two rollouts:

XYZ Generator Parameters Rollout (page 2–1729)

Shaders Rollout (XYZ Generator) (page 2–1730)

XYZ Generator Parameters Rollout

Material Editor > Any shader with a Coords parameter (or other vector value). > Click the shader button. > Material/Map Browser > XYZ Generator (3dsmax) > XYZ Generator (3dsmax) Parameters rollout

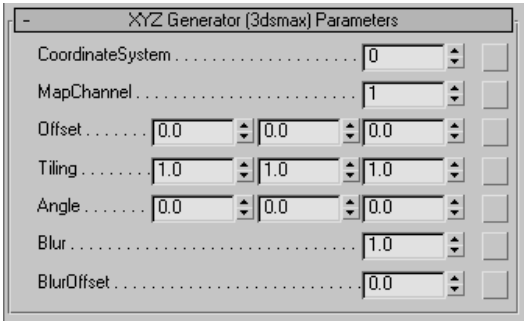
Material Editor > mental ray Connection rollout > Unlock the Surface component and click the shader button. > Material/Map Browser > XYZ Generator (3dsmax) > XYZ Generator (3dsmax) Parameters rollout

Material Editor > DGS material > Assign a shader to any component. > Material/Map Browser > XYZ Generator (3dsmax) > XYZ Generator (3dsmax) Parameters rollout

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The parameters for this shader are equivalent to parameters on the *Coordinates rollout for 3D maps (page 2–1663)*.

Interface



CoordinateSystem—Chooses the source coordinate system. Default=0 (Object XYZ).

- 0 is for Object XYZ.

  Object XYZ uses the object's local coordinate system.

- 1 is for Explicit Map Channel.

  Lets you use the MapChannel value to choose any channel from 1 to 99.

- 3 is for World XYZ.

MapChannel—When Explicit Map Channel is the coordinate system source, this value lets you set the

map channel; otherwise, it is ignored. Range=1 to 99. Default=1.

Offset—Moves the map pattern in X, Y, and Z. Default=(0.0, 0.0, 0.0).

Tiling—*Tiles (page 3–1022)* the map pattern in X, Y, and Z, and makes the pattern narrower. Default=(1.0, 1.0, 1.0).

Angle—Rotates the map pattern in X, Y, and Z. Default=(0.0, 0.0, 0.0).

Blur—Affects the sharpness or blurriness of the map based on its distance from the view. The farther away the map is, the greater the blurring. The Blur value blurs maps in world space. Blur is primarily used to avoid aliasing. Default=1.0.

BlurOffset—Affects the sharpness or blurriness of the map without regard to its distance from the view. Blur Offset blurs the image itself in object space. Use when you want to soften or defocus the details in a map to achieve the effect of a blurred image. Default=0.0.

Shaders Rollout (XYZ Generator)

Material Editor > Any shader with a Coords parameter (or other vector value). > Click the shader button. > Material/Map Browser > XYZ Generator (3dsmax) > Shaders rollout

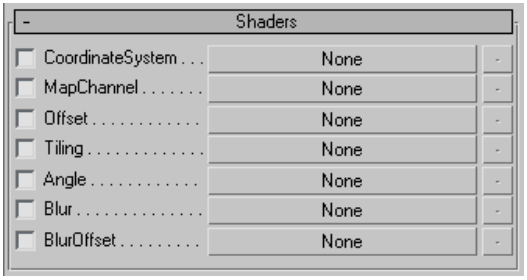
Material Editor > mental ray Connection rollout > Unlock the Surface component and click the shader button. > Material/Map Browser > XYZ Generator (3dsmax) > Shaders rollout

Material Editor > DGS material > Assign a shader to any component. > Material/Map Browser > XYZ Generator (3dsmax) > Shaders rollout

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

The controls on this rollout let you assign a map or shader to one of the basic parameters of the XYZ Generator shader. This is comparable to mapping a component of a standard material; by adding shaders, you can create a shader tree that generates complex effects.

Interface



 The button to the right of each main shader button is for shaders that can return multiple parameters. If a shader that returns multiple parameters is assigned to the component, the button's tooltip shows the parameter name. Clicking the button displays a *Connect Parameter To Shader dialog (page 2–1713)*, which lets you change which parameter is being used.

Important: *UV Coordinates (page 2–1728)* and *XYZ Coordinates (page 2–1730)* are the only shaders with multiple return values provided with 3ds Max. You might encounter multiple return values in shaders provided with other shader libraries or custom shader code.

XYZ Coordinate Shader (mental ray)

Material Editor > Any shader with a Coords parameter or other vector value. > Click the shader button. > Material/Map Browser > XYZ Coordinate (3dsmax)

Note: Shaders appear in the Browser only if the mental ray renderer is the currently active renderer.

This shader is identical to the *XYZ Generator shader (page 2–1729)*, except that it returns two values instead of one.

- The XYZ coordinate (same value as XYZ Generator would return)

This value is named XYZ.

- The XYZ coordinate's derivative

  This value, also a vector, is named dXYZ.

When you choose this shader, a *Connect Parameter To Shader dialog (page 2–1713)* is displayed, prompting you to choose which of the two values to use. You can later change the choice of value by using the "dot" button to the right of the main shader button.

Interface

The XYZ Coordinate shader has the same parameters as the XYZ Generator shader. See *XYZ Generator Parameters Rollout (page 2–1729)* for a description of the basic settings.

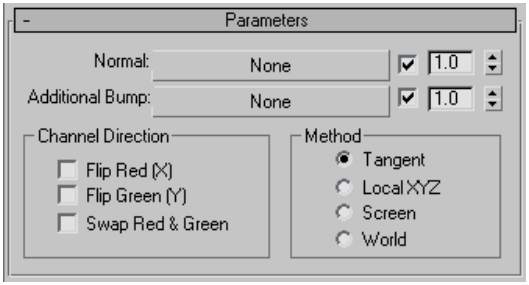
Normal Bump Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Normal Bump

The Normal Bump map lets you use a texture-baked Normals map (see *Baked Texture Elements (page 3–146)*). Typically you assign it to a material's Bump component, Displacement component, or both. Using the map for Displacement can correct edges that otherwise look unrealistically smooth; however, this adds faces to the geometry.

Tip: A Normals map for the indicated material component is generated automatically if you turn on Output Into Normal Bump in the Selected Elements Unique Settings group of the Render To Texture dialog's *Output rollout (page 3–160)*.

Interface



Normal—As a rule, contains a Normals map generated by *Render To Texture (page 3–144)*.

Use the toggle to enable or disable use of the map (default=on). Use the spinner to increase or decrease the map's effect.

Additional Bump—This optional component can contain an additional map to modify the bump or displacement effect. It is treated as a regular *bump map (page 2–1506)*.

Use the toggle to enable or disable use of the map (default=on). Use the spinner to increase or decrease the map's effect.

Channel Direction group

By default, the Normals map's red channel indicates left versus right, while green indicates up versus down (and blue indicates vertical distance). The controls in this group let you adjust that interpretation.

Flip Red (X)—Flips the red channel so that left and right are reversed.

Flip Green (Y)—Flips the green channel so that up and down are reversed.

Swap Red & Green—Swaps the red and green channels so that normal mapping rotates 90 degrees.

Method group

The Method group lets you choose which coordinate to use on the normals. These controls are the same as those in the *Projection Options dialog (page 3–165)*.

- **Tangent—**(The default.) Project at a tangent to the target object's surface.

  This is the method to use for objects that both move and deform, such as animated characters.

- **Local XYZ—**Project using the object's local coordinates.

  This method can be used for stationary or moving objects, but not for objects that deform: if the object deforms, the projection will appear incorrect at some frames.

- **Screen—**Project using screen coordinates; that is, flat projection in the Z axis. X is horizontal, increasing in a positive direction to the right; Y is vertical, increasing in a positive direction upward; and Z is perpendicular to the screen, increasing in a positive direction toward the viewer.

  This method is useful mainly for stationary objects seen only from a single angle; for example, a statue seen through a window.

- **World—**Project using world coordinates.

  This is useful mainly for objects that don't move or deform; otherwise, a moving object with world-projected normals will appear to "swim" through the texture.

Camera Map Per Pixel Map

Material Editor > Maps rollout > Click a Map button. > Material/Map Browser > Camera Map per Pixel

The Camera Map Per Pixel map lets you project a map from the direction of a particular camera. It is meant as an aid to 2D matte painting: You can render a scene, adjust the rendering using an image-editing application, then use this adjusted image as a matte that is projected back onto the 3D geometry.

Tip: Final rendering can be slow. Script-driven *network rendering (page 3–173)* can help improve performance.

Limitations

The Camera Map Per Pixel does not handle these situations:

- Animated objects.

  The projection does not use UVW mapping.

- Animated textures.
- Occlusion based on a Z-depth channel is handled in a limited way only.

Procedure

To use Camera Map Per Pixel:

1. Create the 3D model.
2. Set up a camera.
3. Set up the rendering resolution you want.

   To get good results, the plate should be at least 2K pixels; 3K to 6K, or higher is recommended.

4. Render the scene to an editable image format such as *TIFF (page 3–634)*.
5. Render the scene again, this time to a format such as *RPF (page 3–631)* that has a Z-depth component. Make sure the Z option is turned on.
6. Use an image-editing application to make changes you want to the editable image.
7. Apply Camera Map Per Pixel to the *diffuse component (page 2–1498)* of the geometry on which you want the matte to appear. Use these settings:

   - Set Camera to the same camera you used for the renderings.

- Set Texture to the matte image you edited.
- Set ZBuffer Mask to the Z-depth rendering (the RPF or RLA file).

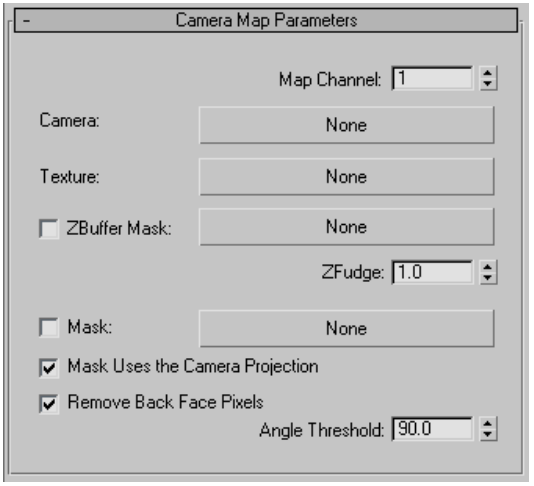
  Adjusting (which usually means increasing) the value of ZFudge can improve the quality of edges of the projection.

  Usually it is a good idea to have Remove Back Face Pixels turned on. Another way to adjust the projection edge is to adjust this control's Angle Threshold.

  Tip: If you have persistent problems with seams, try generating a mask with an alpha channel and using it to clean up the edge of the projection.

If you have multiple mattes to project, you might need to slice geometry to make each map's target a separate object.

Interface



Map Channel—Sets which map channel to use. Default=1.

Camera—Click to turn on, then select a camera in the scene by picking it in a viewport or using the *Select Objects dialog (page 1–78)*.

Once you have assigned a camera, its name appears on this button.

This camera should be the one used to render the map used in the Texture and ZBuffer Mask components.

Texture—Click to assign the texture to project. 3ds Max lets you assign any kind of map, but typically this will be a *Bitmap (page 2–1631)* that contains an image file that you first rendered from the same camera, then edited using a different application.

ZBuffer Mask—Click to assign a map that contains Z-depth data used to mask the projection from unwanted surfaces. Typically this will be an *RPF file (page 3–631)* or an *RLA file (page 3–630)* rendered using the same camera, with the Z channel option turned on.

Use the toggle to turn use of the ZBuffer Mask on or off. By default, it is off, and it is not turned on automatically when you assign the ZBuffer Mask.

- **ZFudge—**ZFudge values other than 1.0 add a margin of deviation to the use of the Z-depth data, letting you fine-tune the Z-Buffer masking. Default=1.0.

Mask—Behaves like the mask in the *Mask map (page 2–1689)*: it lets you view one map through another. Black areas of the mask are transparent, white areas are opaque, and gray areas are partially transparent, based on the percentage of the gray.

- **Mask Uses the Camera Projection—**When on, the mask uses the same camera projection as the Texture and ZBuffer Mask. When off, it uses the object's UVW coordinates. Default=on.

Remove Back Face Pixels—When on, sets the projection to exclude surfaces that face away from the camera, based on the value of Angle Threshold. Default=on.

- **Angle Threshold—**Specifies the angle to use as a cutoff when removing backface pixels. Default=90.0.

  At the default of 90 degrees, faces perpendicular to the camera, or at a greater angle, are not projected.

Material, Mapping, and Vertex Color Utilities

Assign Vertex Colors Utility

Utilities panel > Utilities rollout > More button > Utilities dialog > Assign Vertex Colors

The Assign Vertex Colors utility assigns vertex colors based on the material assigned to the object and the lighting in the scene. The utility applies a *VertexPaint modifier (page 1–936)* to the object when Assign To Selected is clicked. Once the VertexPaint modifier has been applied to the object, go to the Modify panel or click Edit to access the VertexPaint tools.

Tip: To render vertex colors, you must apply a material that has a *Vertex Color map (page 2–1693)* in its diffuse component. To view vertex colors in viewports, right-click the object, choose *Properties (page 1–117)* from the quad menu, and then turn on **Vertex Color** in the Display Properties group.

The Assign Vertex Colors utility supports light inclusion or exclusion when using the Scene Lights option.

All of the commands found within the Assign Vertex Color utility are also available from the Modify panel when a VertexPaint modifier has been applied to an object.

Vertex Colors and Radiosity

The Vertex Colors utility supports *radiosity (page 3–51)*. If you use radiosity with assigned vertex colors, be sure to turn on the option Re-Use Direct Illumination From Radiosity Solution. This option is in the Rendering Parameters rollout. See *Radiosity Controls (page 3–61)*. When this option is on, the renderer simply displays the vertex colors assigned by the radiosity solution: strictly speaking, it is not rendering at all.

The additional option Render Direct Illumination, also on the Rendering Parameters rollout, causes direct lighting not to be saved in the corresponding mesh. This corresponds to the options Radiosity, Render Direct Illumination, in which case Assign Vertex Colors gets indirect illumination from the radiosity mesh but renders direct illumination separately; or Radiosity, Indirect Illumination Only, in which case Assign Vertex Colors doesn't apply direct illumination to vertices at all.

See also

VertexPaint Modifier (page 1–936)

Vertex Color Map (page 2–1693)

Procedures

To use the Assign Vertex Colors utility:

1. Assign materials to the objects you want to affect. These can be mapped or unmapped materials.
2. Light the objects.
3. Select the objects you want to affect.
4. Expand the Display Properties rollout in the Display panel and turn on Vertex Colors.
5. Access the Assign Vertex Colors utility.
6. Choose one of the Light Model options.
7. Choose one of the Color Assignment choices.
8. Click Assign to Selected.

VertexPaint modifiers are applied to the selected objects, and the vertex colors for the objects are taken from their materials and from the lighting in the scene, depending on the options you choose under Light Model.

Changing the material or the lighting in the scene won't change the vertex colors. To do this, click the Update All button.

Note: The new vertex colors are stored in the Vertex Paint modifier. If you want to access them, go to the Modify panel and access the parameters in the rollouts there. You can also use the tools found in the Vertex Paint floating dialog to create layers, paint, blur or adjust color. The Vertex Paint floater launches when you go to the Modify panel and the object is selected.

Example: To use the Assign Vertex Colors utility on a specific object:

1. Create a sphere with 24 segments.
2. Apply a mapped material to the sphere, and turn on Show Map In Viewports.

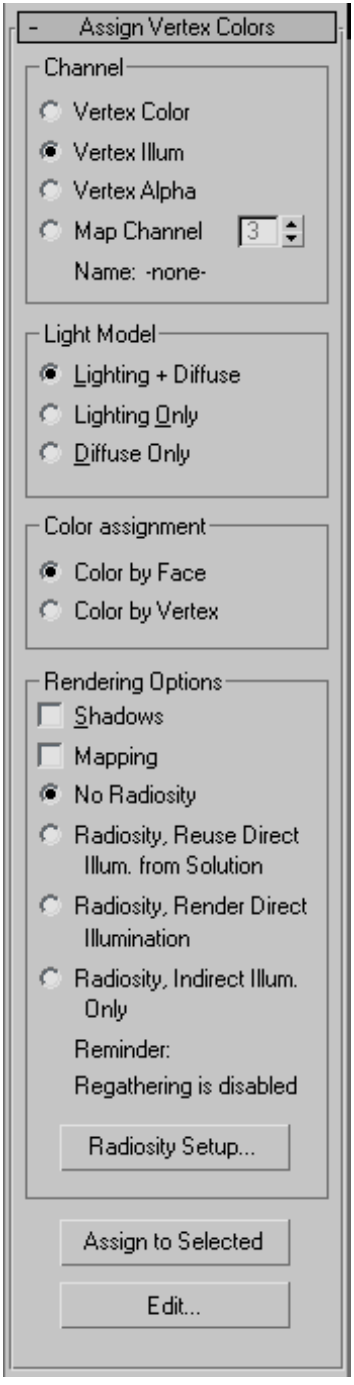
   The mapped sphere is displayed in the viewport.
3. Open the Object Properties dialog for the sphere, turn on Vertex Colors, and click OK.
4. With the sphere selected, open the Assign Vertex Colors utility.
5. Choose Shaded and turn on Use Maps.
6. Click Assign To Selected.

   A blurred version of the mapping appears on the sphere. The vertices are now colored based on the material and the lighting in the scene. The mapping is blurred because the resolution of the mesh at 24 segments is much lower than the pixel resolution of the map.
7. Go to the Modify panel and note the VertexPaint modifier.
8. Move down in the Stack to the creation parameters, click Yes at the warning prompt, and increase the Segments to **70**.

   The new vertices shift the already assigned vertices.
9. Return to the VertexPaint level of the stack, and click Assign in the Assign Vertex Colors rollout. Had we returned to the Utilities panel, we would have added another Vertex Paint modifier to the stack; clicking Assign in the VertexPaint Modifier only updates that modifier.

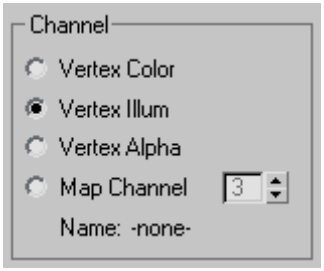
   Tip: Vertex colors will only show up in a rendered scene if you assign the Vertex Color map to the diffuse channel. However, if you do this, you can't properly update your vertex colors with the Assign Vertex Colors utility. The solution is to assign a Blend material to your object. Assign the straight diffuse bitmap to Material 1, and the Vertex Color map to Material 2 of the Blend. Switch to 100 percent of Material 2 when rendering, and 100 percent of Material 1 when updating the vertex colors.

Interface



Channel group

Here you'll find tools to choose which channel type the vertex color utility will assign. If you choose map channel, you can also specify the map channel ID number.



- **Vertex Color—**Choose this to assign a vertex color layer.
- **Vertex Illum—**Choose this to assign a vertex lighting layer.
- **Vertex Alpha—**Choose this to assign a vertex transparency layer.
- **Map Channel—**Choose this to assign a specifically numbered map channel.

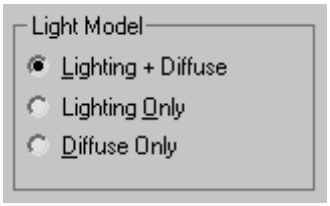
  Map channel spinner—Use this to define the channel number. Available only when Map Channel is chosen.

Name—If a channel has a name defined, it will appear here. Channels can be named using the *Channel Info Utility (page 2–1738)*.

Note: Although the Color, Illum, and Alpha channels have specific names, in fact 3ds Max does not enforce what kind of data is saved in them, and any of the three channels can contain four-channel (RGBA) vertex color data.

Light Model group

Provides options that let you specify how the surface of the object appears to be illuminated.

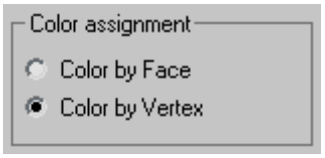


- **Lighting + Diffuse—**Uses the current scene lighting and materials to affect the vertex colors.
- **Lighting Only—**Uses only lighting to assign vertex colors, ignoring material properties.

  When this option is chosen, Shadows and Mapping are disabled in the Rendering Options rollout.
- **Diffuse—**Uses the material's diffuse color, ignoring the lighting.

Color Assignment group

Lets you specify how colors are interpolated across surfaces.



- **Color by Face—**(The default.) Colors are interpolated between the center of each face.

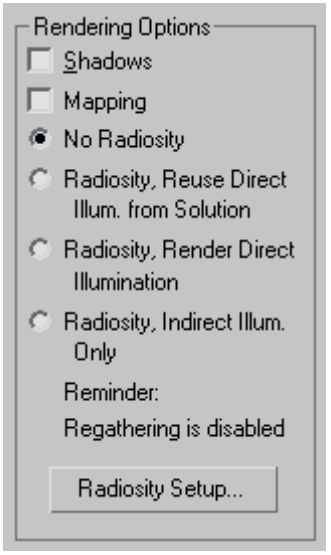
  Color By Face samples fewer points, so it is the quicker method. On the other hand, results are less accurate.
- **Color by Vertex—**Colors are interpolated between vertices.

  For each face, this method uses three points instead of one, so it is slower but usually more accurate. An exception can occur when an object's shadow falls between two vertices: in such a case, the object should occlude lighting, but because only vertices are taken into account, the shadow is not calculated and a "light leak" occurs.

Rendering Options group

The options in this group let you choose whether to include shadows, texture maps, or a radiosity solution in vertex colors.

Note: You can save a radiosity solution in vertex colors, but not *Light Tracer (page 3–44)* illumination, which is not stored in the scene's geometry.



Shadows—When on, shadows are used when the vertices are shaded. Default=off.

Tip: You can soften the shadow edge by using the VertexPaint modifier's Paint or Blur tools.

Mapping—When on, texture maps are used when the vertices are shaded. Default=off.

The radio buttons specify how to use radiosity data.

- **No Radiosity—**(The default.) Do not use the radiosity solution when assigning vertex colors.

  Note: This option is the only one available unless a *radiosity solution (page 3–51)* is present in the scene.

- **Radiosity, Reuse Direct Illum. from Solution**—Includes radiosity in the vertex color assignments, and uses the direct illumination from the solution.

  This is comparable to the choice Re-Use Direct Illumination From Radiosity Solution on the *Rendering Parameters rollout (page 3–71)*.

  This choice disables the Shadows toggle, because shadows don't need to be recomputed.

- **Radiosity, Render Direct Illumination**—Includes radiosity in the vertex color assignments, but uses a separate pass to render direct illumination.

  This is comparable to the choice Render Direct Illumination on the *Rendering Parameters rollout (page 3–71)*.

- **Radiosity, Indirect Illum. Only**—Includes only indirect illumination from the radiosity solution in the vertex color assignments.

  This choice disables the Shadows toggle, because shadows don't need to be recomputed.

Reminder field—Displays a message that says whether regathering is enabled or disabled. Regathering provides the most accurate radiosity results, but it can add considerable time to radiosity calculations.

Radiosity Setup—Click to display the *Advanced Lighting panel (page 3–44)* of the Render Scene dialog, where you can set up and generate a radiosity solution.

This button is not available if the mental ray renderer is the active renderer.

Assign to Selected—Assigns vertex colors to the selected objects based on the assigned material, and the choices specified in the preceding group boxes. Assign to Selected creates a VertexPaint modifier and adds it to the stack of the selected objects.

Edit—Click to display the *VertexPaint Paintbox (page 1–941)*, the floating dialog that holds the vertex painting tools.

This button is unavailable if you haven't yet clicked Assign To Selected.

Channel Info Utility

Utilities panel > More button > Channel Info > Click Channel Info button.

Tools menu > Channel Info

The Channel Info utility gives game artists and others direct access to objects' channel information that might not otherwise be easily available. All objects in 3ds Max have mapping channels, which hold information pertinent to texture mapping as well as vertex color, illumination, and alpha. Mesh objects also have geometry and vertex-selection channels. The Channel Info utility lets you view an object's channels, give them meaningful names, delete unused channels, and copy information between channels.

The utility's Map Channel Info dialog shows all the channel data for selected objects. It displays the number of channels, the number of vertices per channel, and how much memory the channel uses. It also lets you name channels, as well as clear (or delete), copy, and paste channels. Each of these commands except renaming puts a modifier on the stack to achieve the results.

Note: Channel Info supports mesh, polygon, and patch objects. It does not support NURBS objects.

See also

This map and utility and these modifiers work in conjunction with Channel Info:

Vertex Color Map (page 2–1693)

Skin Utilities (page 2–700)

Select By Channel Modifier (page 1–785)

UVW Mapping Add Modifier (page 1–933)

UVW Mapping Clear Modifier (page 1–933)

UVW Mapping Paste Modifier (page 1–934)

Procedure

To use the Channel Info utility:

This procedure is a guide to basic usage of the Channel Info utility. For more detailed usage methods and examples, see the tutorial Using the Channel Info Utility.

1. Select an object or objects to use with the utility.
2. Open the utility.

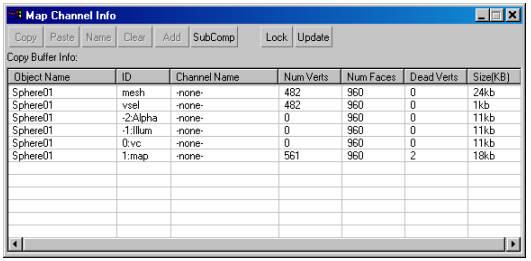
   The Map Channel Info dialog opens.
3. To create a map channel, click any channel and then click the Add button.

   The new, empty channel appears at the end of the list.
4. Most channels have three components. For example, a mesh or map channel has X, Y, and Z components, and an alpha channel has R, G, and B components. To expand all three-component channels, click the SubComp button. To collapse all expanded channels, click SubComp again.
5. To copy one channel to another, click the source channel, click Copy, and then click the destination channel and click Paste.

   In some cases, you might need to expand or collapse the component display (see previous step). For example, when copying a vertex selection (vsel) channel to a map channel, you must paste the vsel channel to a component channel.
6. To minimize a channel's memory footprint, click the channel and then click the Clear button.

   This removes most all or of the data from the channel, so first make sure the data is unnecessary or is available elsewhere. If the cleared channel is the last one in the list, it might be deleted from the list.

Interface



The primary user interface of the Channel Info utility is the Map Channel Info dialog, which you open by clicking the utility's Channel Info button on the command panel. This modeless dialog shows information about all map channels belonging to the current selection, at the object level. If you change the selection, the dialog automatically updates to reflect the selection.

The dialog consists of two parts: a button toolbar at the top, and a tabular display of map channels belonging to each object in the current selection.

Channel Info toolbar

Copy—Copies the channel data from the highlighted channel to the copy buffer, where it becomes available for pasting. After you copy a channel, its name appears on the line below the button toolbar.

Paste—Pastes the contents of the copy buffer to the highlighted channel.

You can copy and paste only between channels with the same topology, or you can copy from any channel to a channel with no vertices.

Source and destination channels need not be of the same type. For instance, you can copy from a mesh channel to a map channel, and vice-versa.

Name—Lets you rename the highlighted channel. Click this button to open a small dialog that displays the current channel name and lets you edit this name or enter a new one from the keyboard.

Clear—Use this function to remove channels or delete data from a map channel (including alpha, illumination, and vertex color channels). Clear has no effect on geometry or vertex selection channels.

The specific result depends on the type of object and which channel you clear. In terms of reducing the object's memory footprint, the utility is most effective with Editable Poly objects.

- **Geometric primitive or Editable/Edit Mesh object—**Deletes the highlighted texture map channel if it is the last map channel in the object, and it's not the default map channel (1:map). If the highlighted channel is not the last, Clear deletes all vertices in the channel. The faces remain, so the memory-footprint reduction is partial.

  Note: This also applies to objects that collapse to editable mesh, such as Loft objects.

- **Editable Poly object—**Deletes the highlighted texture map channel if it is the last map channel in the object, and it's not the default map channel. If the highlighted channel is not the last, Clear deletes all vertices and faces in the channel.

- **Patch object—**Deletes the highlighted texture map channel if it is the last map channel in the object, and it's not the default map channel. If the highlighted channel is not the last, Clear has no effect.

Note: When you use the Clear function, the software adds a UVW Mapping Clear modifier to the object's modifier stack. You can recover the deleted data by removing the modifier from the stack, or changing its Map Channel setting.

Add—Appends a new map channel to the object's channel list. If multiple objects are selected, Add becomes available only after you click a track, so the software knows which object to add the channel to.

Note: If you apply mapping with a channel number higher than any existing channels, the software automatically creates all intermediate channels. For example, if you apply a UVW Mapping modifier to a standard object and set Map Channel to 5 in the modifier, the software adds map channels 2, 3, 4, and 5.

SubComp—Toggles display of the channels' subcomponents. When displayed, you can rename, copy, and paste each subcomponent independently of its parent channel.

Each channel except vsel has three subcomponents. Mesh and map channels' subcomponents are labeled X, Y, and Z; those of alpha, illumination, and vertex color channels are R, G, and B (red, green, and blue).

Lock—Retains the current mapping data information in the table even if you change the selection.

For example, if you want to see mapping data for a specific object or objects constantly, first select the objects and then click Lock. Thereafter, if you select different objects in the viewport, the table continues to display the data for the selection when you clicked Lock. If you turn off Lock, the table updates to show data only for the current selection.

If you click Update when Lock is on, the software will refresh the table contents to reflect the current selection, and then retain that data.

Update—Refreshes the displayed data to reflect any changes in the objects or mapping, or, when Lock is on, the selection.

For example, if you apply mapping to an object, or change its mapping, click Update to display the changes in the Map Channel Info dialog.

Channel Info table

The table functions similarly to a spreadsheet. If not all rows or columns are visible, you can scroll the table using standard methods, including rolling the mouse wheel for vertical scrolling. To highlight a row, click anywhere in the row. You can highlight only one row at a time. To resize a column, drag the vertical divider at the right of the column heading. To automatically set a column's width to the size of the longest entry, double-click the vertical divider to the right of the column heading.

Following is a brief explanation of each of the columns in the table:

Object Name—The name of the object. If you change the name in the Modify panel, click the dialog's Update button to display the new name in the dialog.

ID—The type of channel. The available channel types are:

- **mesh/poly—**The object's mesh or poly data, depending on whether it's a mesh or poly object: vertices and faces. You can copy this channel and paste it to any other three-component channel.

  This channel is not available for patch objects.

- **vsel—**The vertex selection. You can copy this channel and paste it to other channels' subcomponents.

  This channel is not available for patch objects.

- **-2:Alpha—**The vertex alpha channel. You can transfer all vertex alpha values between objects with the same topology by copying and pasting this channel.

  You can apply vertex alpha information to objects with the *VertexPaint modifier (page 1–936)*, and to editable surfaces with the Vertex Properties settings (*editable poly (page 1–1034)*) and Surface Properties (*editable mesh (page 1–1005)* and *editable patch (page 1–978)*).

- **-1:Illum—**The vertex illumination channel. You can transfer all vertex illumination values between objects with the same topology by copying and pasting this channel.

  You can apply vertex illumination information to objects with the *VertexPaint modifier (page 1–936)*, and to editable surfaces with the Vertex Properties settings (*editable poly (page 1–1034)*) and Surface Properties (*editable mesh (page 1–1005)* and *editable patch (page 1–978)*).

- **0:vc—**The vertex color (vc) channel. You can transfer all vertex color values between objects with the same topology by copying and pasting this channel.

  You can apply vertex color information to objects with the *VertexPaint modifier (page 1–936)*, and to editable surfaces with the Vertex Properties settings (*editable poly (page 1–1034)*) and Surface Properties (*editable mesh (page 1–1005)* and *editable patch (page 1–978)*).

- **1:map—**The default mapping channel. You can transfer all UVW mapping information between objects with the same topology by copying and pasting this channel.

  You can create additional mapping channels by various means, including with the Channel Info utility.

Channel Name—The name of the channel. By default, a channel has no name, as indicated by the entry "-none-". To name or rename the channel, click the channel to highlight it and then click the Name button at the top of the dialog, or right-click

the channel and choose Name from the right-click menu.

Note: Most channels can be split into *subcomponents*. You can name the subcomponents separately from the channel itself.

Num Verts—The number of vertices in the channel. To paste one channel to another, they must have the same number of vertices.

Some channels have faces but no vertices. This is typically the case with Alpha, Illumination, and vertex color channels in newly created non-poly objects. In such cases, these channels function as placeholders for the corresponding data should you add it later. They do consume a small amount of memory, so if you have no intention of using a channel, you can save some memory by converting the object to Editable Poly.

Num Faces—The number of faces in the channel.

If a channel has faces but not vertices, that means it's a placeholder. See Num Verts, above, for more information.

Dead Verts—The number of unused map vertices in the channel. Such vertices can be left over from sub-object editing.

Size(KB)—The approximate amount of memory consumed by the channel. Use this figure to check for unused channels that are using up memory.

Clean MultiMaterial Utility

Material Editor > Utilities menu > Clean MultiMaterial

Select an object. > Utilities panel > More button > Clean MultiMaterial > Click Find All button.

The Clean MultiMaterial utility parses Multi/Sub-Object materials and displays any that contain sub-materials are not assigned to any material IDs in the scene. You can then choose to remove any unused sub-materials, thus consolidating your Multi/Sub-Object materials.

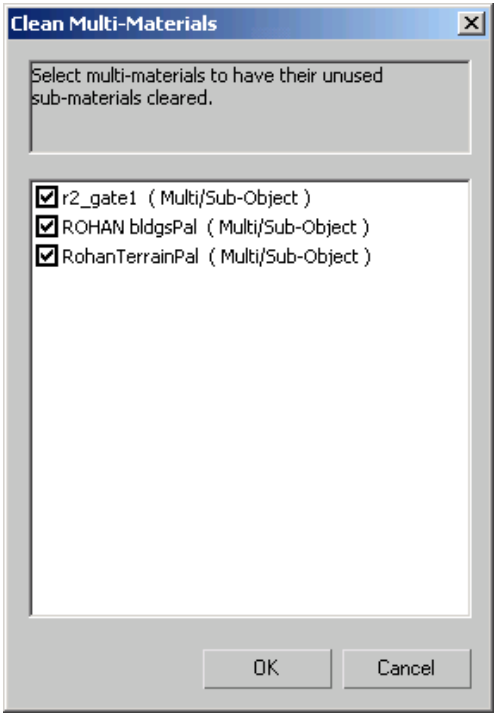
This utility searches an entire scene. You do not need to select objects or materials.

Procedures

To clean all materials:

1. Open a scene.
2. On the Material Editor, open the Utilities menu and choose Clean MultiMaterial.

   The Clean Multi-Materials dialog opens, displaying the following:



   The dialog displays a list of all Multi/Sub-Object materials that contain unassigned sub-materials. All the Multi/Sub-Object materials are automatically turned on, and thus subject to cleaning.

3. Click the OK button.

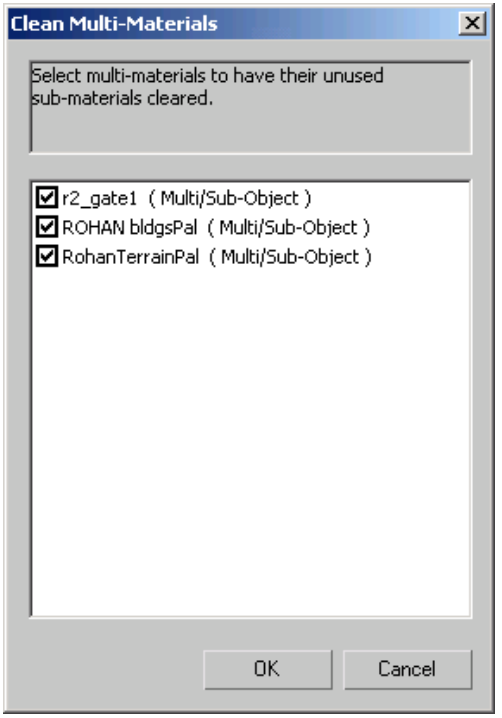
All unused sub-materials are deleted and the dialog closes.

To clean specific Multi/Sub-Object materials:

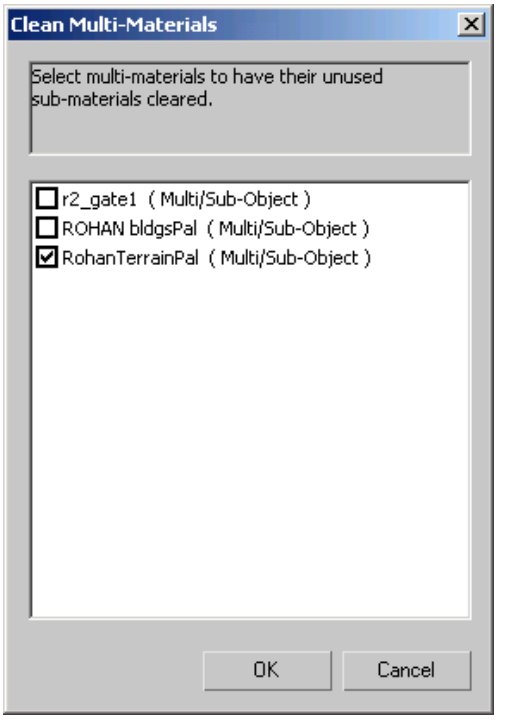
When you create a Multi/Sub-Object material, you might create extra sub-materials in anticipation of objects that haven't been added to the scene yet. In such cases, you can prevent the utility from removing the extra sub-materials.

1. Open a scene.
2. On the Material Editor, open the Utilities menu and choose Clean MultiMaterial.

   The Clean Multi-Materials dialog opens.



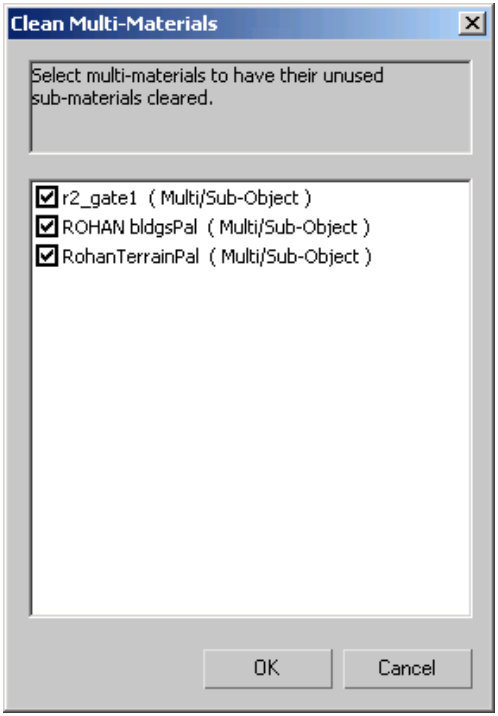
3. In the list, turn off any Multi/Sub-Object materials you do not want to clean.



4. Click OK.

   All unused sub-materials are deleted from the indicated materials and the dialog closes.

Interface



The Clean dialog presents you with a list of all Multi/Sub-Object materials in the scene that contain unassigned sub-materials. You can then select materials from which to remove unassigned sub-materials.

The dialog consists of two parts: a status field at the top and a list of Multi/Sub-Object materials.

Status Field

The Status Field is not interactive. It displays prompts about the unused materials.

Materials list

This list displays Multi/Sub-Object materials that have sub-materials that are not currently being used in the scene. When the Clean MultiMaterial utility opens, it lists all Multi/Sub-Object materials with unused sub-materials and turns them on to be cleaned. Click the check box next to a material name to turn it off and prevent the utility from cleaning it.

Clean—Deletes unused sub-materials from Multi/Sub-Object materials that are turned on.

Cancel—Cancels the operation.

Instance Duplicate Maps Utility

Material Editor > Utilities menu > Instance Duplicate Map

Utilities panel > More button > Instance Duplicate Maps > Click Find All button.

The Instance Duplicate Maps utility searches an entire scene for materials that have duplicate Bitmap maps and give you the option to instance them. If your scene has different materials that use the same texture maps, creating instances will reduce the load on your video card, which can improve viewport performance.

For example, if you render a scene containing three materials that reference the texture map *MyMap.bmp*, the software searches for that texture map three times: once for each material it is used in. However, if you use the Instance Duplicate Maps utility, you can create instances of the duplicate maps so the renderer will find the first reference to *MyMap.bmp* and use it for any subsequent material that uses the same maps.

This utility searches an entire scene. You do not need to select objects or materials.

Note: In order to be eligible for instancing, the Bitmap maps must be identical in all aspects with regard to their initial settings. For example, if two materials use the same bitmap image applied as Diffuse maps, but have different initial Tiling settings, their maps aren't eligible for instancing.

Animation is not supported as a criterion for determining duplication, so any differences in the animation of Bitmap map parameters will be

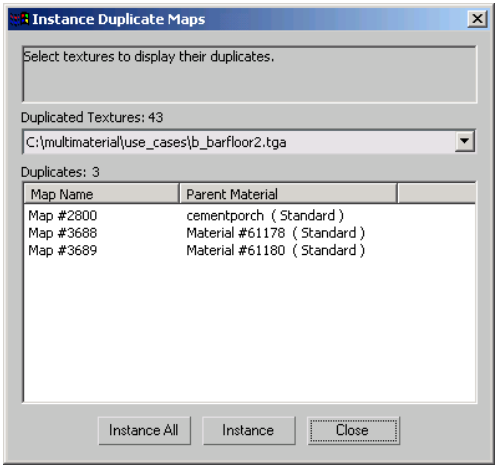
lost from the use of this utility. For example, if two materials use the same bitmap image applied as Diffuse maps, and have the same initial Tiling settings, but their Tiling settings are animated to different subsequent values, after using the utility both maps will have the same animation as the first map.

Procedures

To instance all duplicate maps:

1. Open a scene.
2. On the Material Editor, open the Utilities menu and choose Instance Duplicate Map.

   The Instance Duplicate Maps dialog opens, displaying the following:



   The dialog displays a list of all texture maps found to have exact duplicates in the scene. The search will include all copies found to have identical paths and settings.

3. Click the Instance All button.

   All identical maps are instanced and the dialog closes.

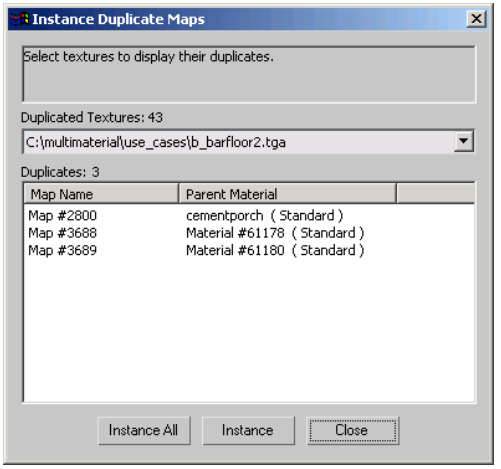
To instance only specific maps:

Perhaps you only want to instance a few of the texture maps found in specific materials.

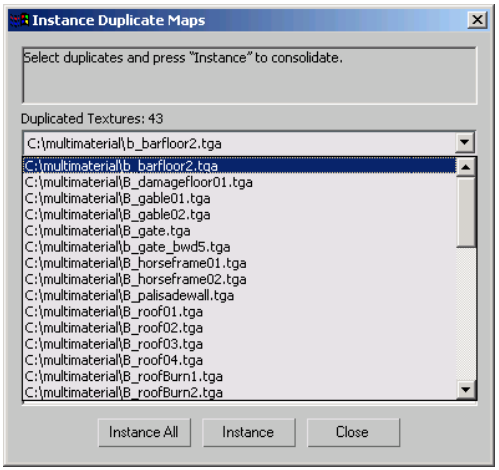
Note: You have to select at least two maps from the Duplicates list

1. Open a scene.
2. On the Material Editor, open the Utilities menu and choose Instance Duplicate Map.

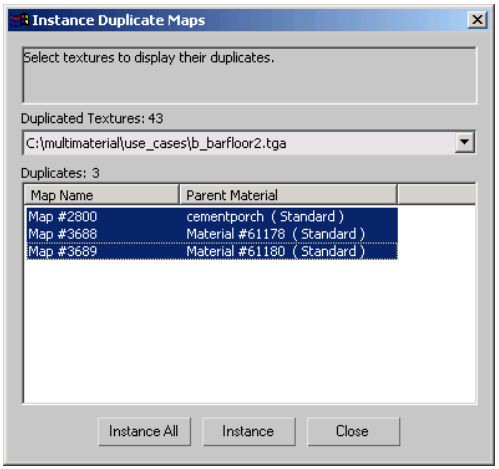
   The Instance Duplicate Maps dialog opens.



3. From the Duplicated Textures list, click the arrow button to open the list of duplicated textures and choose a texture you want to instance.

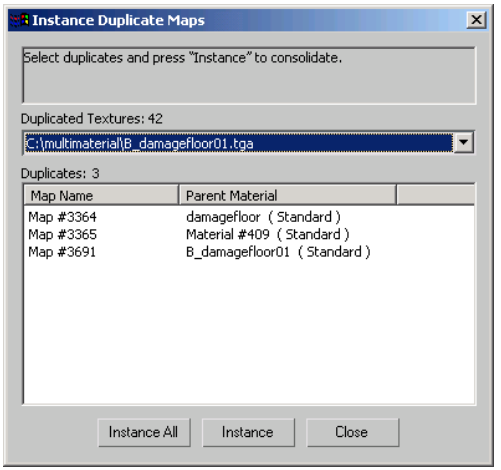


4. From the Duplicate list, select at least two map entries.



If you don't select at least two entries, you will see a warning in the status field at the top of the dialog.

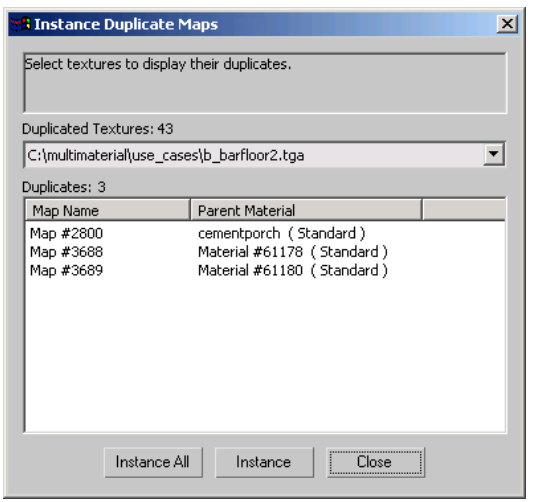
5. Click Instance.



Since all three of the maps were instanced, the dialog now shows there are 42 duplicated textures.

6. Continue instancing texture or click Close when you are done.

Interface



The primary user interface of the Instance Duplicate Maps utility is the Instance Duplicate Maps dialog, which you open from the Utilities menu in the Material Editor. This utility works

globally, so you do not need to select objects or materials in order to use it.

The dialog consists of three parts: a status field at the top, a drop menu containing duplicated texture maps and a list of map names and the materials that belong to.

Status Field

The Status Field is not interactive. It shows you prompts and warnings about the duplicate textures and maps. Some of the messages displayed in this section include:

- "No duplicate textures were found in the scene." – when the utility is run in a scene containing no duplicates.
- "Select duplicates and press "Instance" to consolidate." – when the utility is run in a scene containing duplicates.
- "The selected duplicates contain parameters that are animated. Animation is not supported as criteria for determining duplication so differences in the animation will be lost." – when animation is present in a set of duplicates.
- "At least two maps must be selected in order to proceed." – when the Instance button is pressed with one or no duplicates selected.

Duplicated Textures list

This drop-down list contains all the maps in the scene found to be identical in every way, including texture map path and name, and initial parameter settings. The number after "Duplicated Textures" indicates how many sets were found in the scene. This list appears with the first entry visible and once active can be scrolled using the up/down arrow keys.

Duplicates list

When you choose a texture in the Duplicated Textures list, the software displays its duplicates in the Duplicates list, showing the map name and the name of the parent material. The number after "Duplicates" indicates how many copies were found in the scene. Textures in this list can be chosen individually. Only duplicates chosen in this list will be consolidated into the final instance if you click the Instance button. The name of the resulting instance is that of the first chosen duplicate in the list.

Instance All—Performs the consolidation on all duplicates in the scene regardless of selections made from the Duplicated Textures or Duplicates lists.

Instance—Performs the consolidation on only the duplicates chosen on the Duplicates list. Selected duplicates will disappear from this list after consolidation. If all are selected, the corresponding texture will disappear from the first list as well.

Close—Closes the operation at its current point.

Index

Symbols & Numerics

1-rail sweep surface 1–1204
2 3 4 links 2–1111, 2–1150
2 feet down 2–988, 2–997
2.5D snap 2–35
2-point perspective 2–1392
2-rail sweep surface 1–1209
2-sided 3–855, 3–901
2D
  2D coordinates rollout 2–1625
  2D images 3–608
  2D maps 2–1624
  2D snap 2–35
2D map
  glossary 3–901
32–bit floating-point output 3–613
3D
  3D coordinates rollout 2–1663
  3D maps 2–1662
  3D snap 2–35
  viewing and navigating 3D space 1–21
3D displacement shader (mental ray) 2–1714
3D DWF
  exporting 3–555
  glossary 3–901
3D map
  glossary 3–902
3DS files
  3DS import dialog 3–530
  exporting 3–532
  importing 3–530
3ds Max
  bones 2–834, 2–1080
  knowledge of 2–832
  main window 1–9
  materials 3–83
3dsmax.ini file 1–17 to 1–18

A

about
  about MAXScript 1–xvii
  buttons 2–1106
  footstep animation 2–856
  freeform animations 2–886
absolute 2–1150
absolute snap 2–35, 2–41
absolute/offset coordinate display 3–709
abut selected (video post) 3–329
acceleration (raytrace) 2–1531
acceleration techniques (raytracer) 2–1528
acceleration test (particle flow) 2–233
accuracy 2–120, 3–815
AccuRender materials
  in 3ds Max 3–455
actions (particle flow) 2–141, 3–903
  action recovery 2–121
  adding 2–131
  and order/priority 2–124
  and time frames 2–141
  editing parameters 2–131
  order of 2–123, 2–130
activate all maps 1–50
activate options 2–974
activating
  grid object 2–34
  home grid 2–34
  joint axes 2–485
  maps 1–50
activating footsteps 2–865
active 2–1136
active link (glossary) 3–903
active time segment 2–286
active time segment (glossary) 3–904
active viewport 1–22
active/inactive footsteps 3–904

ActiveShade 3–17
commands (quad menu) 3–22
floater 3–21
glossary 3–904
quick render 3–17
viewport 3–21
actual stride height 2–992, 2–995
actual stride length 2–992, 2–995, 2–997
actual stride width 2–992, 2–995, 2–997
adapt locks 2–871, 2–980, 3–905
adaptation 3–905
adapting
keyframes to edits 2–871
keys to footstep edits 2–871
adaptive antialiaser dialogs 2–1533 to 2–1534
adaptive control 1–167, 1–828, 2–1534, 2–1698
adaptive degradation
glossary 3–905
override 1–34
viewport configuration dialog 3–859
adaptive subdivision dialog 1–706
add
alpha compositor (video post) 3–381
atmosphere 2–1351, 3–304
contrast filter (video post) 3–343
cross fade compositor (video post) 3–381
default lights to scene 1–49
effect 2–1351
external event (video post) 3–340
fade filter (video post) 3–344
image alpha filter (video post) 3–344
image filter event (video post) 3–335
image input event (video post) 3–332
image layer event (video post) 3–337
image output event (video post) 3–339
lens effects filter (video post) 3–345
loop event (video post) 3–342
negative filter (video post) 3–345
note track (Track View) 2–552
pseudo alpha compositor (video post) 3–382
pseudo alpha filter (video post) 3–346
scene event (video post) 3–329
simple additive compositor (video post) 3–383
simple wipe compositor (video post) 3–383
simple wipe filter (video post) 3–347
starfield filter (video post) 3–347
time tag dialog 3–710
visibility track 2–556
add change option 2–1108
add keys
function curves 2–581
Track View toolbar 2–560
add layer 2–326, 2–334
add selection to current layer 3–667
add texture elements dialog 3–164
add to track set 2–591
add twist pose 2–950
add/edit DBR host dialog 3–128
adding
comments (particle system) 2–134
editable spline vertices 1–297
footsteps 2–863
poses 2–1096
splines 1–303, 1–308
adding controllers to bipeds 2–896
additive opacity (glossary) 3–906
adjust animation range (track bar) 3–703
adjust color dialog
vertexpaint modifier 1–949
adjust pivot rollout 2–488
adjust talent pose option 2–1065
adjust transform rollout 2–489
adjusting
default envelope shape 2–1086
keys in Track View 2–875
link parameters 2–1091
normals and smoothing 1–166
object transforms 2–432
pivots 2–423
talent pose 2–1065, 3–906
Adobe Illustrator files
exporting 3–534
importing AI 88 format 3–533
advanced effects rollout 2–1341
advanced file link settings 3–431
advanced key info 2–306
advanced lighting
object properties 1–123
select advanced lighting rollout 3–44
advanced lighting override material 2–1601
advanced quad menu options 3–801
advanced ray-traced parameters rollout 2–1356
advanced rollout 1–791
advanced shaders rollout
mental ray material 2–1548
advanced surface approximation dialog 1–1245
AEC design elements 1–210
AEC extended 1–210
editing wall objects 1–228
foliage 1–214
railing 1–217
wall 1–223
affect region 1–557
soft selection rollout (EMesh) 1–963
soft selection rollout (NURBS) 1–1147
affine transformation (glossary) 3–906

after trajectory 2–944
age test 2–211
AI import dialog 3–524
airborne option 2–988, 2–995, 2–997
airborne periods 2–878, 3–906
aliasing/antialiasing
  alias against background 3–826
  and supersampling 2–1459
  fast adaptive 2–1533
  filters 1–567, 3–38
  glossary 3–907
  multiresolution adaptive 2–1534
align 1–462
  align geometry dialog (edit poly) 1–679
  and pivot point 2–488
  camera 1–468
  dialog 1–462
  editable mesh objects 1–1011
  editable mesh vertices 1–1011
  flyout 1–462
  grid to view 2–35
  keys (Track View) 2–556
  normals 1–465, 2–10
  objects 2–8
  selected left (video post) 3–328
  selected right (video post) 3–328
  to view (dialog) 1–468
  to view (toolbar) 1–468
all bipeds 2–944
all links 2–1150
allow non-vertical jambs 1–210
alpha channel 3–272, 3–907
alpha compositor (video post) 3–381
alpha map (baking) 3–149
alternate 2–992, 2–995, 2–997
altitude (sunlight and daylight systems) 1–421
ambient
  and diffuse map lock 2–1474
  and raytrace materials 2–1514
  color (glossary) 3–908
  light 2–1276, 2–1279 to 2–1280, 3–908
  lighting (rendered environment) 3–272
  mapping 2–1497
ambient occlusion map (baking) 3–149
analysis of lighting 3–76
analyze errors 2–1017
analyze panel 2–1017
analyzing
  Shockwave 3D files 3–585
  W3D files 3–585
anatomy of biped 2–846
anchor (VRML97 helpers) 3–597
anchor patches 1–968
anchors 2–962
angle 1–286, 2–1070
  transition editor 2–1051
angle of incidence 2–1276, 2–1279
angle separation 3–826
angle snap 2–37
angular dashpot 2–732
animatable IK attachments to 3ds Max objects 2–908
animated
  bitmap 2–1450
  material previews 2–1450
  reference objects (particle flow) 2–178
  texture 3–908
animating
  a biped with footsteps 2–856
  a tablecloth (FFD(cylinder) space warp) 2–95
  attachment 2–433
  cameras 2–1381
  lens effects properties (video post) 3–349
  lights 2–1282
  links 2–430
  materials 2–1449
  NURBS models 1–1099
  parameters 2–121
  scene 1–8
  shift+rotate 1–482
  shift+scale 1–482
  sub-object geometry 1–998
  transforms 1–432
  vertices 2–346
  with applied IK 2–481
  with IK solvers 2–446, 2–461, 2–472
  with interactive IK 2–480
animating a quadruped 2–907
animating hair 1–520, 1–540, 1–545
animation
  and particle system time frames 2–123
  auto key mode 2–278
  combining animations 2–1026
  concepts 2–275 to 2–276
  constraints 2–392
  controllers 2–307, 3–909
  dynamics & adaptation rollout 2–980
  expanding tracks 2–886
  glossary 3–909
  hair and fur modifier 1–520, 1–540, 1–545
  importing and exporting 2–921
  importing from another scene 3–466
  insert animation 1–114
  layers 3–910
  loading 3–474
  mapping 3–478 to 3–479, 3–481
  methods 2–275

NURBS 1–1091
preferences settings 3–828
previewing animations after attaching Physique 2–1084
sample animations in this release 2–920
saving 3–476
selecting and moving tracks 2–886
show ghosting 1–46
tips (NURBS) 1–1099
toggle animation mode 2–278
utilities 2–653
with radiosity 3–60
animation controls 3–716
animation layer weight 2–325 to 2–326
animation layers 2–325 to 2–326, 2–333 to 2–334, 3–690
animation layers toolbar 3–690
animation menu 3–681
bone tools 1–411
constraints 2–371, 2–392 to 2–393, 2–396, 2–398, 2–401, 2–406, 2–409
create bones 1–404
dummy 2–16
IK solvers 2–440, 2–473
make preview 3–168
parameter collector 1–138
parameter editor 1–129
previews 3–168
rename preview 3–170
view preview 3–170
wire parameters 2–411 to 2–412
animation mode
set key 2–280
animation modifiers 1–557
animation quad menu 3–697
animation range
ignore 2–521
respect 2–521
animation workbench 2–1012
anisotropic highlights 2–1492
anisotropic shader
anisotropy mapping 2–1504
basic parameters rollout 2–1480
ankle attach option 2–846, 2–984
ankle tension 2–959
antialiasing 3–98, 3–939, 3–1005
append
.mfe file 2–1032
motion flow editor 2–1032
append footsteps 2–863, 2–936
applied ik
glossary 3–910
applied IK 2–439, 2–481, 2–491, 2–497
apply ease curve 2–584
apply ease or multiplier curve (Track View) 2–584
apply increment 2–962, 2–965
apply multiplier curve 2–584
applying materials 2–1405, 2–1409
arc 1–274
arc rotate 3–744
arch & design material
main reference 2–1549
overview 2–1562
tips & tricks 2–1569
architectural material 2–1535
advanced lighting override rollout 2–1540
and radiosity solution 2–1540
physical qualities rollout 2–1536
special effects rollout 2–1539
templates rollout 2–1536
archive
file menu 3–499
program (files preferences) 3–819
archiving scenes 1–19
area light rollout 2–1354
area light sampling rollout 2–1354
area lights 3–910
omni 2–1298
spotlight 2–1299
area omni light 2–1298
area shadows 3–911
area shadows rollout 2–1357
area spot light 2–1299
arms
arm link 2–984
resizing 2–852
turning on 2–846
array 1–450
arraying objects 1–484
button 1–450
creating 1–471
creating arrays 1–487, 1–489
dialog 1–450
flyout 1–448
using the array dialog 1–485
array dialog 1–450
array flyout 1–448
artificial light 2–1280
ASCII files: exporting 3–534
ASE files 3–534
aspect ratio 3–30, 3–911
assemblies
and groups 1–98
and selection sets 1–98
attach to 1–111
close 1–109
create 1–107
disassemble 1–110

explode 1–110
open 1–109
using 1–98
assembly commands 1–107
assembly heads helper objects 1–111
assembly menu
assemble 1–107
attach 1–111
detach 1–110
disassemble 1–110
explode 1–110
asset browser 3–504
internet download dialog 3–515
preferences 3–514
using 1–17
asset tracking
dialog 3–487
icons 3–498
open from vault 3–389
asset tracking dialog 2–920, 3–487
asset tracking dialog icons 3–498
asset tracking prompts 3–498
assign
controller (Track View) 2–546
controller rollout (motion panel) 3–774
material to selection 2–1441
object effects dialog 2–696
random colors 1–161
vertex colors utility 2–1734
assign controller rollout (character studio) 2–934
assign random colors 1–161
assign renderer rollout 3–35
assign to link 2–1089, 2–1150
assigning
colors to objects 1–159
controllers 2–292
materials 2–1405, 2–1409
associate bipeds with delegates 2–1199
associate objects with delegates 2–1196
assume skin pose 1–116
asterisk (in modifier stack) 1–508
atmospheres and effects
adding atmosphere 3–276
and raytracing 2–1528
atmospheric effect 3–272, 3–282, 3–284, 3–288
for atmospheric apparatus 3–304
for lights 2–1349
atmospheric apparatus 3–304
BoxGizmo 3–304
CylGizmo 3–306
SphereGizmo 3–307
attach
attach options 1–1018
attaching and importing 3ds Max objects 1–1120
controls dialog (block controller) 2–388
editable mesh edges 1–1011
editable mesh vertices 1–1011
editable patch object 1–986
object 1–968, 1–988
splines 1–295, 1–297, 1–308
to assembly 1–111
to groups 1–106
attach points/tendon 2–1147
attach to deforming mesh constraint 2–799
attach to node 2–1106
attach to rigid body constraint 2–798
attached links 2–1147
attaching
mesh to a biped using Physique 2–1106
tendon to another link 2–1096
attachment
animating 2–433
constraint 2–393
attachments (IK) 3–912
attenuation 2–1276, 2–1279, 3–912
and lights 3–821
parameters 2–1345
raytrace attenuation rollout 2–1706
attribute holder modifier 1–559
attributes (custom) 1–129
audio controller 2–309, 2–386
AudioClip (VRML97 helpers) 3–606
auto 2–1070
auto clip names 2–1027
auto expand
animated 2–527
base objects 2–527
keyable 2–527
materials 2–527
transforms 2–527
xyz components 2–527
auto key 2–278, 3–717
and set key 2–280
auto termination rollout 2–499
auto timing 2–992, 2–995, 2–997
AutoCAD
blocks 3–456 to 3–457
importing DWG and DXF files 3–536
importing DXF files 3–551
instanced objects 3–456
AutoCAD Architecture
files 3–444
materials 3–445
objects 3–444, 3–461
styles 3–461
AutoCAD blocks in 3ds Max 3–441

AutoCAD DWG/DXF import options dialog 3–536
AutoCAD, AutoCAD Architecture, and Revit
working with 3–440
autodesk inventor files
importing 3–552
Autodesk Vault 3–487
Autodesk VIZ files 3–525
autogrid 3–913
AutoGrid 2–7
automatic
auto archive 3–819
auto backup 1–19, 3–819
auto secondary (lens effects) 3–238
auto termination (IK) 2–499
automatic exposure control 3–295
unit conversion 3–815
automatic mapping rollout
rendering to texture 3–163
autoplay preview file 3–815
AVI files 3–168, 3–609
avoidance behavior 2–1164, 2–1211, 3–913
preventing collisions 2–1240
awning window 1–256
axis constraints 1–437, 3–687
and hierarchies 2–500
and rollouts 2–500
and snaps 2–41
axis constraints toolbar 3–687
axis order 2–916, 2–948
axis ordering 2–1012
axis tripod 1–45
and transform managers 1–433
and World Axis 1–424
axonometric views 1–24, 3–913
azimuth (sunlight and daylight system 1–421

B

B-spline (glossary) 3–913
backburner
network rendering 3–201, 3–208
Backburner 3–173
backburner command line control 3–215
backface cull on object creation 3–821
backfacing - ignore 1–996, 1–1011, 1–1019
background 1–44
and antialiasing 1–567, 3–38
color 3–272, 3–276
image 3–272
reset background transform 1–45
sample slot 2–1433
select background image 1–42
update viewport image 1–44
viewport 3–731
viewport background 1–38
VRML97 helpers 3–605
backlight (sample slot) 2–1432
backup
auto 3–819
backing up and archiving scenes 1–19
backup on save 3–819
files 3–819
recovered files 1–20
backward knees (creating characters with) 2–891
baked material rollout
rendering to texture 3–162
baking animation 2–120
baking textures 3–144
target map slot 3–150
texture elements 3–146
balance
animating 2–876 to 2–877
shifting 2–876
shifting for entire footstep animation 2–876
balance factor 2–876 to 2–877, 2–954, 3–914
balance parameters dialog 2–629
balance track 2–622
ballistic gait 2–878, 3–914
ballistic tension 2–846, 2–878, 2–945, 2–954, 3–914
barycentric
coordinates (glossary) 3–914
morph controller 2–300, 2–309
morph controller key info dialog 2–346
base layer 2–974
basic file link settings 3–429
basic key info 2–304
basic parameters rollout
materials 2–1470
PArray 2–258
basics
basic building blocks 1–155
creating and modifying objects 1–153
file linking 3–416
selecting objects 1–64
batch
rendering 3–201, 3–203, 3–208
batch file conversion (motion capture) 2–1065, 2–1075
batch rendering 3–201, 3–203, 3–208
backburner 3–201 to 3–202, 3–208
batch render dialog 3–203
batch render tool 3–203
cameras 3–201, 3–203, 3–208
error dialog 3–203
errors 3–203
network rendering 3–201, 3–203, 3–208
presets 3–203
quick start 3–201
scene states 3–201, 3–203, 3–208

using 3–202
batch rendering completed 3–203
before trajectory 2–944
behavior assignments and teams dialog 2–1200
behavior rollout 2–1211
behaviors 2–1159, 3–915
avoid 2–1211, 3–913
fabric 1–579
obstacle-avoidance 2–1164, 3–983
orientation 2–1214, 3–985
patch-based 3–991
path follow 2–1216, 3–992
repel 2–1218, 3–1003
scripted 2–1220, 3–1008
seek 2–1220, 3–1008
space warp 2–1221, 3–1013
speed vary 2–1222, 3–1015
surface arrive 2–1223, 3–1018
surface follow 2–1226, 3–1019
wall repel 2–1227, 3–1033
wall seek 2–1229, 3–1033
wander 2–1231, 3–1033
bend 2–990, 2–1136
bend links mode 2–895, 2–936, 2–952, 3–915
bend modifier 1–560
bend parameters (links) 2–1091
bending
center of mass track 2–914
footstep path 2–869
bevel
bevel modifier 1–562
bevel profile modifier 1–565
deformation 1–366
faces and polygons 1–1011
patches 1–986
types of beveling 1–366
bevel polygons dialog 1–1066
bezier
controllers 2–310
handle control 2–582
bezier curves 3–915
bias 2–950, 2–956, 2–1091, 2–1136
mental ray shadow maps 2–1360
bifold door 1–252
billboard (VRML97 helpers) 3–607
binding
objects 2–437, 2–461, 2–491
to space warp 2–58
vertices 1–297
BioVision motion capture data files 2–1065
BIP files 2–919 to 2–920, 2–1065, 2–1263, 3–916
adding to motion mixer 2–609
adjust time in motion mixer 2–615
adjusting in motion mixer 2–611
combining motions 2–924
combining with mixer 2–604
export with motion mixer 2–624
filtering in motion mixer 2–612
loading 2–942
path 2–1041
saving 2–882, 2–941
transitions in motion mixer 2–616
biped 2–833, 3–916
add to motion mixer 2–607
and physique 2–834
assign controller rollout 2–934
body parameters 2–844, 2–846
center of mass 2–833
colored keys 2–1005
creating 2–844
display options 2–944
dummies 2–922
dynamics 2–833, 2–980
dynamics parameters 2–954
edit 2–1038
editing keys in Track View 2–875
figure files (.fig) 2–855
figure mode 2–847
IK key colors 2–1005
keyboard shortcuts 2–1006
layer 2–974
load motion file 2–936
moving keys 2–1004
naming 2–847
playback 2–936, 3–916
previewing motion 2–929
root object 2–846
select keys based on foot states 2–965
setting keys 2–904
shifting balance 2–876
structure 2–847
tracks in Track View 2–888
user interface 2–932
working with 2–843
biped apps rollout 2–935
biped balance, motion mixer 2–622
biped clips 2–649
biped colored keys 2–1005
Track View 2–1005
biped crowds 2–1172
biped dynamics 3–916
biped IK key colors 2–1005
biped links
selecting and rotating 2–895
biped object 2–636
biped playback 2–936

biped rollout 2–936
bipeds
and crowd simulation 2–1187
correcting posture 2–925
deleting 2–854
linking objects to 2–854
moving objects 2–890
posing 2–847
posing a biped 2–925
rotating objects 2–891
scaling after physique is applied 2–1099
visible in playback 2–944
bipeds dialog 2–643
birth event 3–916
birth operator 2–143
birth script operator 2–145
bitmap map 2–1631
bitmap pager 3–828
bitmap pager statistics dialog 3–514
bitmap proxies 3–32
bitmap proxies dialog 3–496
bitmap/photometric path editor 3–510, 3–516
bitmaps
choosing 2–1635
display 3–840 to 3–841, 3–844
glossary 3–917
Material Editor 2–1631, 2–1635
path configuration 3–503, 3–811
path editor 3–510
blend
blend curve (NURBS) 1–1158
blend element parameters 3–140
blend material 2–1588
blend object (glossary) 3–917
blend surface (NURBS) 1–1183
blend from/to 2–1140
blend map (baking) 3–149
blend materials
limitations when baking textures 3–147
blend weight 1–807
blending
between links 2–1085, 2–1111, 2–1150
envelopes 2–1086
blending envelope display options dialog 2–1125, 2–1128
Blinn highlights 2–1493
Blinn shader basic parameters 2–1480
blizzard 2–251
BlobMesh object 1–331
block
controller 2–389 to 2–390, 2–392
block controller 2–313
block controllers (Track View) 2–1179
block reference 3–917
block/style parent 3–918
blocks 3–457
and linking to 3ds Max 3–918, 3–1031
AutoCAD 3–441, 3–456
AutoCAD and 3ds Max 3–438
material assignment 3–458
materials 3–458
multi-view 3–459
propagation 2–1432
blowup (render) 3–13
blue vertices 2–1150
blur
and blur offset (glossary) 3–918
particle motion blur 2–240
rendering effect 3–260
BMP files 3–610
body 2–960
horizontal tracks 2–846, 2–945
parameters (biped) 2–846
space 3–918
tracks 2–846
turning track 2–945
vertical tracks 2–846, 2–945
bomb space warp 2–105
bone tools 1–411
bone editing 1–411
fin adjustment 1–413
object properties 1–414
bones 1–404, 2–1080
and IK solvers 2–440, 2–472
bone base 2–944
bone tip 2–944
display 2–853
exporting 3–580
floating bones rollout (Physique) 2–1110
linking to follow objects 2–461
spline IK solver 2–477
stretch factor 1–415
used with physique 2–1079, 2–1082
using objects as 1–410, 2–440
using unlinked bones with Physique 2–1082
Boolean controller 2–316
Booleans
alignment 1–338
and editable splines 1–308
Boolean objects 1–338, 1–378, 1–388
Boolean operation (glossary) 3–919
colinear edges 1–338
combining objects that have materials 1–345
complexity between operands 1–338
coplanar faces 1–338
face normals 1–338
inverted meshes 1–338

material IDs 1–338
overlapping elements 1–338
surface topology 1–338
troubleshooting problems with 3–885
with maps and materials 1–338
bound vertex 1–297, 3–919
boundary conditions (and tendons) 2–1147
bounding box (and envelope creation) 2–1111
bounding box (glossary) 3–919
bounds (inner/outer) 2–1085
box
BoxGizmo 3–304
standard primitive 1–171
box caustics filter 3–106
box method 2–1242
box selected
render bounding box/selected dialog 3–16
branching events (particle flow) 2–123
break
spline at selected vertex 1–297
vertices 1–1011
breathe option (links) 2–1091, 2–1136
bricks 2–1658
bridge dialog 1–1067
bridge edges dialog 1–1068
brightness and contrast effect 3–265
browse 2–1070
browser
material/map 2–1412
browsing from 3ds Max 3–504
brush
styling hair with a 1–529
brush options 1–960
brush preset manager 3–692
brush presets 3–690
BSP method 3–129
BSP method, raytrace acceleration 3–1000
bubble motion rollout (PArray) 2–270
bubble motion with Particle Flow 2–123
bubble notification
communication center 3–716
buckets, distributed rendering 3–124
buffer mode 2–936
build face 1–1011
bulge angle display properties dialog 2–1127
bulge angles 2–1114, 2–1141, 3–920
adding 2–1095
changing 2–1095
choosing for editing 2–1095
color 2–1141
deleting 2–1095
parameters 2–1114
setting 2–1095
bulge editor 2–1096, 2–1106, 2–1114, 2–1135, 2–1141
bulge sub-object 2–1127, 2–1141
bulges 2–1111, 2–1113, 3–920
creating 2–1094
fine-tuning 2–1096
overview 2–1093
shaping 2–1096
workflow 2–1094
bump mapping 2–1506, 2–1539
bump shader (mental ray) 2–1716
button
2.5D snap 2–35
2D snap 2–35
3D snap 2–35
ActiveShade floater 3–21
align 1–462
align camera 1–468
align to view 1–468
angle snap 2–37
animate 3–717
arc rotate 3–744
array 1–450
auto key 3–717
bind to space warp 2–58
button sets (utilities) 3–779
clone and align tool 1–459
crossing 1–93
current frame 3–724
dolly camera 3–746
dolly light 3–751
dolly target 3–746, 3–751
edit current event 3–324
edit scene event (video post) 3–329
full screen 3–738
get material 2–1439
go forward to sibling 2–1447
go to end 3–724
go to frame 3–724
go to parent 2–1446
go to start 3–722
light falloff 3–753
light hotspot 3–752
lock selection 2–555
make unique (Material Editor) 2–1442
material and map type 2–1449
Material Editor 2–1427
Material Editor options 2–1436
maximize viewport 3–738
mirror 1–448
new sequence 3–323
next frame 3–724
normal align 1–465
open sequence 3–323

orbit/pan camera 3–749
orbit/pan light 3–755
pan (Track View) 2–595
pan (user interface) 3–743
percent snap 2–38
perspective 3–747
pick material from object 2–1448
place highlight 1–467
play/stop 3–723
previous frame 3–723
quick align 1–465
quick render 3–17
quick render (Production) 3–17
render scene 3–12
roll camera 3–747
roll light 3–753
save sequence 3–324
scale keys 2–559, 2–580
scale values 2–581
select and link 2–422
select and manipulate 2–15
select and move 1–439
select and rotate 1–439
select and uniform scale 1–441
select by material 2–1439
select by name 1–77
select object 1–77
selection center 1–447
selection lock 3–707
sets of modifiers 3–772
show curves 3–705
show end result 2–1446
snapshot 1–453
spacing tool 1–455
spinner snap 2–38
squash 1–442
transform coordinate center 1–447
truck camera 3–748
truck light 3–755
unlink selection 2–422
use pivot point center 1–446
zoom 3–739
zoom (Track View) 2–596
zoom all 3–740
zoom extents 3–740
zoom extents all 3–737
zoom extents all selected 3–737
zoom extents selected 3–740
zoom horizontal extents 2–595
zoom region 3–742
zoom region (Track View) 2–597
zoom selected object 2–588
zoom value extents 2–596
button appearance 3–803
BVH files 2–919, 2–1061, 2–1065, 2–1263, 3–920
by layer 3–920
bylayer 3–655

C

C-Ext 1–200
cache operator (particle flow) 2–197
CAL files 2–1070
calculation order (joint precedence) 2–467
calculator 1–12
calibrating marker files 2–1065
camera correction modifier 2–1392
camera effects rollout 3–101
camera map modifier
  object space 1–567
  world space 1–513
camera map per pixel map 2–1732
camera match
  camera match helper 2–1391
  camera match point 2–1391
  camera match utility 2–1387
  camera point 2–1391
  CamPoint 2–1391
camera tracker
  batch track rollout 2–678
  camera tracker utility 2–667
  error thresholds rollout 2–677
  match move rollout 2–680
  motion trackers rollout 2–673
  move smoothing rollout 2–682
  movie rollout 2–670
  movie stepper rollout 2–676
  movie window 2–671
  object pinning rollout 2–683
  position data rollout 2–679
  requirements for camera tracking 2–669
  troubleshooting 2–685
camera view
  right-click menu 3–731
camera viewports 1–33, 3–745
cameras 2–1365, 2–1372
  align camera button 1–468
  animating 2–1381
  camera object icons 2–1365
  camera view 1–24
  camera viewport controls 3–745
  choosing for vertical views 2–1365
  common parameters 2–1373
  create camera from view 1–48
  depth of field parameters (mental ray renderer) 2–1383
  dolly or target 3–746
  free 2–1370
  match camera to view 1–468

multi-pass parameters 2–1383, 2–1386
orbit/pan 2–1381, 3–749
placing 1–7
roll 3–747
setting lens size 2–1373
target 2–1371
truck 3–748
using clipping planes to exclude geometry 2–1379
using horizon to match perspective 2–1380
using move and rotate to aim 2–1379
using transforms to aim 2–1379
with target 2–1371
zoom 2–1381
candela 3–965
canopy mode 1–214
cap holes modifier 1–569
cap surface 1–1195
capsule 1–195
capture viewport 1–35
car paint material and shader 2–1576
car-wheel constraint 2–757
cartoon shading 2–1605
casement window 1–257
category, hiding and unhiding objects by 1–72
caustics 3–92, 3–106
caustics (mental ray) 3–80
caustics and global illumination rollout 3–106
CCB files 1–950
cellular map 2–1664
center 1–435
center of mass 2–833, 2–933, 3–920
object 2–846
selecting tracks 2–888
shadow 2–846
shifting balance with 2–876
tracks in Track View 2–945
chains (kinematic) 2–471, 3–960
chamfer
and editable splines 1–297
chamfer curve (NURBS) 1–1161
ChamferBox 1–191
ChamferCyl 1–192
editable mesh edges 1–1011
glossary 3–920
chamfer dialog 1–1070
change of value over time 2–578
changed feature 1–334, 2–66, 2–543, 2–1546, 3–149, 3–398, 3–539
changing
biped body parameters 2–844
biped name 2–847
controller properties 2–291
initial biped anatomy 2–846
light objects 2–1282
link inheritance 2–434
smoothing 1–167
video system 2–1434
changing biped to bones 2–921
channel 1–285
channel (map) 3–966
channel info
skin utilities 2–700
channel info utility 2–1738
character 1–112
character assembly 1–102
and parameter wiring 1–104
create character 1–112
destroy character 1–115
insert animation 3–466
insert character 1–115
lock/unlock 1–115
save character 1–115
skin pose 1–116
character modeling 1–842
character studio
assign controller rollout 2–934
definition 2–831
file formats 2–1263
space warp behavior 2–1221
checker map 2–1638
child overlap 2–1130
children
don't affect 2–489
choose directory dialog 3–808, 3–810
choose renderer dialog 3–36
choosing
child-to-parent precedence 2–469
colors for realism 2–1400
parent-to-child precedence 2–470
playback speed and frame rate 2–288
transform center 1–435
CIBSE files 3–921
cineon image file format dialog 3–610
circle 1–273
circular
falloff graph 3–254
circular arrays 1–489
circulating
materials 2–1432
clean multimaterial utility 2–1742
clean remove 1–1039
clear UVW mapping 1–933
clip
ratio 2–615
replace 2–634
timing 2–615

transition 2–1048
clip controllers 2–1179, 3–921
clip frame numbers
motion mixer 2–615
clip mode 2–1027
clip properties dialog 2–1027, 2–1045, 2–1059
clipping planes 2–1373, 2–1379, 3–921
clips 2–1045
combining 2–1026
create 2–1027
looping with motion-capture filtering 2–1061
menu 2–632
move 2–1027
path 2–1041
ClipState dialog 2–1253
clone 1–476
clone and align tool 1–459
clone options dialog (particle flow) 2–132, 2–136
cloning 1–476
clone 1–476
materials 2–1432
objects 1–453, 1–474, 1–483
shape sub-object selections 1–289
shift+move 1–479
shift+rotate 1–480
shift+scale 1–481
sub-object geometry 1–998
cloning characters 2–922
close
assembly 1–109
group (group menu) 1–105
close curve dialog 1–1228, 1–1235
cloth 2–778
cloth modifier 1–578
cloth simulation 1–571
collection 2–781
collision detection 1–572
dT messages during simulation 1–583
effect of geometry on 1–576
fabric behavior 1–579
garment maker modifier 1–607
how it works 1–576
mesh density 1–577
modifier 2–778
object properties 1–602
overview 1–571
troubleshooting and error codes in garment maker 1–622
units of measure 1–579
cloth modifier 1–578
user interface 1–582
clothing design 1–575
CLR files 3–799
cmdjob.exe 3–215
codec (glossary) 3–921
cognitive controller 3–921
state dialog 2–1207
state transition dialog 2–1208
cognitive controller editor 2–1206
cognitive controllers 2–1170
coincident - making splines 1–842
collapse
collapse utility 1–966
stack 1–504, 1–966
vertices 1–1011
collapse controller tool (Track View) 2–522
collapse layer 2–326, 2–333
collapsing
animation tracks 2–886
layers 2–974
collections
cloth 2–781
creating and deleting 2–966
deforming mesh 2–794
loading and saving 2–966
rigid body 2–723
rope 2–792
soft body 2–788
collision 2–243
collision detection 2–891
cloth 1–572
collision tests (particle flow)
collision 2–212
collision spawn 2–215
collisions
storing and accessing 2–774
collisions rollout 2–810, 2–1240
color
and light 2–1276, 2–1279, 2–1331
and particle view display operator 2–131
and program state 1–12
and realism 2–1400
assigning to objects 1–159
balance (render effect) 3–265
bleeding 3–93
changing vertex color 1–1009
color bleeding 3–45
color selector 1–161
copying 1–165, 2–1452
display 1–52
illegal video colors 2–1434
name and color rollout 3–757
object color dialog 1–159
parameters 2–1345
selecting vertices by 1–652, 1–1029
temperature (light color) 2–1276

color clipboard files 1–950
color clipboard utility 1–165
color coding 2–945
color controls 2–1485
color modifier maps 2–1692
color palette
  vertexpaint modifier 1–950
color RGB controller 2–317
color selector 1–161, 3–815
color space 3–1
colors
  assign random 1–161
  biped IK/FK keys 2–1005
  biped keys in Track View 2–947
  biped trajectories 2–1005
  footsteps 2–869
  in Track View 2–944
  vertex type 2–1089
colors panel (customize UI) 3–799
COM 2–933
COM/DCOM server control utility 3–792
combining animations 2–1026
combining motions
  motion mixer 2–604
combining objects 1–338, 1–378, 1–388
Combustion
  adding workspace 3–135
combustion map 2–1639
combustion workspace file 3–611
command line
  rendering 3–209, 3–211, 3–215
  startup options 3–671
command panel
  troubleshooting when missing 3–893
command panels
  create 3–757
  display 3–775
  hierarchy 3–773
  modify 3–758
  motion 3–774
  overview 3–756
  utilities 3–778
command-line options (MAXScript) 3–783
command-line options (starting 3ds Max) 3–671
command-line rendering 3–209, 3–215
  backburner command line 3–215
  batch render 3–209
  DOS 3–211
  pre-render scripts 3–209, 3–215
commands provided only from the keyboard 3–669
comments
  particle system 2–134, 2–206
comments on the documentation 3–874
common panel
  render scene dialog 3–27
common parameters rollout 3–27
common procedures
  video post 3–315
commonality 1–509
communication center 3–712 to 3–713
  bubble notification 3–716
  configure 3–713
  new information 3–716
  notifications 3–716
  refresh content 3–715
  settings 3–713
  welcome wizard 3–713
compare dialog (loft objects) 1–374
compass helper object 2–27
complete map (baking) 3–147
component color - specular 3–1014
components
  hair and fur feature 1–517
composite
  glossary 3–922
  map 2–1688
  material 2–1589
composite materials
  limitations when baking textures 3–147
compositor
  compositor maps 2–1687
compound materials
  glossary 3–922
  kinds of 2–1587
compound objects 1–313
  BlobMesh 1–331
  Boolean 1–338, 1–378, 1–388
  conform 1–324
  connect 1–328
  mesher 1–374
  morph 1–314
  ProBoolean 1–378
  ProCutter 1–388
  scatter 1–318
  ShapeMerge 1–336
  terrain 1–347
compound rigid bodies 2–722
concepts 3–88
cone 1–172
cone (spotlight) 2–1338
cone angle manipulator 2–27
cone caustics filter 3–106
configuration
  network rendering 3–175
configuration settings 3–119

configure
communication center 3–713
Direct3D 3–844
driver 3–821, 3–840
key mode 3–725
modifier sets 3–772
OpenGL 3–841
presets (video post) 3–327
software display driver 3–840
system paths 3–810
time 3–725
track bar 3–703
user paths 3–808
utilities button sets 3–779
viewports 3–853
configure paths 3–808
configure preset dialog 3–33
configure system paths 3–810
plug-ins path configuration 3–814
configure user paths 3–808
bitmaps 3–189
external files 3–811
file i/o path configuration 3–810
FX files 3–811
using 3–189
conform
compound object 1–324
space warp 2–103
connect
to child link 2–1147
to parent link 2–1147
connect compound object 1–328
connect edges dialog 1–1070
connect parameter to shader dialog (mental ray) 2–1713
constant
coordinate system 3–815
facet shading 3–937
constrained point
glossary 3–922
constraint solver 2–736
constraint spaces 2–725
constraints 2–289, 2–392, 2–724
and bones 1–404
and cloth 1–580
angular dashpot 2–732
attachment 2–393
axis constraints 1–437, 3–687
breakable 2–735
car-wheel 2–757
concepts 2–725
constraint solver 2–736
constraint spaces 2–725
cooperative constraints 2–735
deformable constraints 2–795
hinge 2–747
linear dashpot 2–730
link 2–403
look-at 2–406
orientation 2–409
path 2–398
point-path 2–762
point-point 2–750
position 2–401
prismatic 2–754
rag doll 2–737
simple constraints 2–727
spring 2–727
surface 2–396
contact object (particle flow) 3–922
contacting us 1–xiv
containers (glossary) 3–922
continuity 2–956, 2–1091, 2–1136, 3–923
continuity level (glossary) 3–923
NURBS concepts 1–1091
continuous time frame 2–141
contour shading
mental ray 3–96
contrast filter (video post) 3–343
contrast sampling thresholds 3–98
control lattice (glossary) 3–923
control objects (IK) 2–435, 2–446
control points 2–1114, 2–1130, 2–1141, 2–1147, 3–923
and bulges 2–1096
and envelopes 2–1088
rotating 2–1088
control vertex (CV) 3–926
control vertex (glossary) 3–924
controller 3–924
controller menu, Track View 2–521
controller toolbar, Track View 2–540
controller type 2–333 to 2–334
controller window, Track View 2–512
controllers 2–295 to 2–297, 2–307, 2–375, 2–896
assigning 2–292
audio 2–309
barycentric morphing 2–309
bezier 2–310
block 2–313
Boolean 2–316
categories of 2–289
changing length 2–502
changing properties 2–291
changing range 2–502
collapsing procedural controllers 2–522
color RGB 2–317
copy 2–544

default settings 3–828
Euler XYZ rotation 2–318
expression 2–320
frame duration 2–502
general-purpose controllers 2–295
limit 2–335
linear 2–341
list 2–342
local euler XYZ rotation 2–344
look at 2–344
make unique 2–550
master point 2–346
morph 2–300
motion capture 2–347
noise 2–353
on/off 2–355
paste 2–545
point3 XYZ 2–317
position XYZ 2–356
properties 2–560
PRS 2–357
reaction 2–358
scale XYZ 2–371
script 2–372
slave 2–313
smooth rotation 2–374
specifying default 2–294
TCB 2–377
time duration 2–502
transform script 2–379
types of 2–289, 2–546
understanding 2–289
viewing types 2–289
waveform 2–381
working with 2–289
xref 2–383
controlling
colors 2–299
display performance 1–28
flipping on path 2–398
IK precision 2–463
object display 1–51
position 2–298
rotation 2–299
time 2–285
transforms 2–298
viewport rendering 1–27
controls
camera viewport 3–745
light viewport 3–750
perspective and orthographic 3–738
special 1–12
viewport 3–735
conversion modifier
turn to mesh 1–871
turn to patch 1–873
turn to poly 1–874
convert 2–936
between footstep and freeform animations 2–885
data in motion capture buffer 2–1065
from buffer 2–1065
to freeform 2–886
to freeform/footsteps dialogs 2–999
convert curve dialog 1–1225
convert surface dialog 1–1227
convert to mesh 2–206
converting
event-level selection to particle level 2–138
TRC into CSM 2–665
convex hull property (glossary) 3–924
cool (glossary) 3–924
cooperative constraints 2–735
coordinate display (mouse position) 3–708
coordinate space 3–924
coordinate system 1–443
coordinates
absolute/offset display 3–709
barycentric (glossary) 3–914
coordinate display 3–708
coordinate system 1–443
coordinates rollouts 2–1625, 2–1663
mapping 2–1405
copies 1–472
creating 1–471
overview 1–472
copy 2–1141, 2–1147
a material, map, bitmap, or color 2–1418
footsteps 2–990
pose 2–962
poses and postures 2–910
posture 2–962
selected cross section 2–1114
copy biped animation to clip dialog 2–635
copy controller 2–544
copy keys 2–579
copy layer 2–325 to 2–326
copy map dialog 2–1451
copy tangent handles 1–297
copy time (Track View) 2–568
copy tracks 2–926
copy/paste rollout 2–966
copying
actions, events (particle flow) 2–127, 2–132, 2–134
and pasting items/objects 2–575 to 2–576
colors 1–165, 2–1452
copy controller (Track View) 2–544

copy time (Track View) 2–568
copy track (Track View) 2–568
joint parameters 2–495
materials 2–1409, 2–1432
modifiers 2–544
patch surface 1–968
presets 3–437
splines 1–308
transform keys 2–283
CPY files 2–1263
crash recovery 1–20
crease
at link's joint 2–1140
at parent's joint 2–1140
create
clip 2–1027, 2–1045
envelopes 2–1111
footsteps 2–988
keys for inactive footsteps 2–865, 2–990
layer 2–974
random motion 2–1035, 2–1055
script 2–1030, 2–1048
separate tracks for biped arms 2–980
shared motion flow 2–1039, 2–1056
transition 2–1045, 2–1051
unified motion 2–1038
create assembly (assembly menu) 1–107
create assembly dialog 1–107
create biped rollout 2–844
create camera from view 1–48
create character 1–112
create key dialog 2–284
create layer 2–326, 2–334
create material preview dialog 2–1452
create menu 1–347, 3–675
AEC objects 1–214, 1–217, 1–223, 1–232, 1–235, 1–239, 1–243, 1–251 to 1–252, 1–256 to 1–261
cameras 2–1365, 2–1370 to 2–1371
compound objects 1–318, 1–328, 1–331, 1–336, 1–338, 1–352
extended primitives 1–186
lights 2–1272
NURBS 1–1102 to 1–1103, 1–1106, 1–1110
particles 2–237
patch grids 1–993 to 1–995
photometric lights 2–1302 to 2–1305, 2–1307, 2–1309
shapes 1–262, 1–270, 1–272 to 1–274, 1–276 to 1–278, 1–281 to 1–282, 1–284 to 1–288
standard lights 2–1288 to 2–1290, 2–1292 to 2–1293, 2–1295
standard primitives 1–170
create method rollout 2–1242
create multiple footsteps 2–988, 2–992
jump 2–988, 2–997
run 2–988, 2–995
walk 2–988, 2–992
create new layer 3–667
create new map files 2–332
create out of range keys 2–533
create out-of-range keys utility (Track View) 2–562
create panel 1–154, 3–757
cameras 2–1365
helpers 2–2
lights 2–1272, 2–1301
space warps 2–55
systems 1–404
create position lock key 2–310
create rotation lock key 2–310
create shape (editable patch) 1–988
creating
1-rail sweep surface 1–1204
2-rail sweep surface 1–1209
a script 2–1045
an object 1–157
animated material previews 2–1450
arrays 1–471
biped character 2–844
biped skin 2–1076
blend surface 1–1183
bulges 2–1094
cap surface 1–1195
chamfer curve 1–1161
circular and spiral arrays 1–489
copies 1–471
crowd system 2–1155
curve sub-objects 1–1151
custom sample object 2–1425
CV curve on surface 1–1172
CV curve sub-object 1–1153
CV surface 1–1103
CV surface sub-object 1–1179
dependent curve point 1–1220
dependent curve-curve point 1–1223
dependent curves 1–1151
dependent offset point 1–1219
dependent surface point 1–1222
dependent surface-curve point 1–1224
dependent surfaces 1–1177
editable mesh edges from shapes 1–1006
editable mesh vertices 1–1011
extrude surface 1–1188
faces 1–1011
fillet curve 1–1164
fillet surface 1–1216
footsteps 2–863

footsteps automatically 2–862
freeform animations 2–886
independent surfaces from NURBS curve objects 1–1114
individual footsteps 2–863
iso curve 1–1168
lathe surface 1–1190
linear arrays 1–487
mirror curve 1–1160
mirror surface 1–1187
models with NURBS 1–1094
multicurve trim surface 1–1214
multiple footsteps 2–862
multiple slices 1–1019
multisided blend surface 1–1213
normal projected curve 1–1169
NURBS curves from splines 1–1115
NURBS CV curve 1–1110
NURBS models 1–1079
NURBS point curve 1–1106
NURBS sub-objects 1–1081
NURBS surfaces 1–1101, 1–1116
objects (basics) 1–153
offset curve 1–1159
offset surface 1–1186
particle emitter 2–239
particle system 2–238
particles 2–143, 2–145
physique links and envelopes 2–1111
physique modifier 2–1083
point curve on surface 1–1175
point curve sub-object 1–1155
point curve with curve fit 1–1157
point sub-objects 1–1078, 1–1219
point surface 1–1102
point surface sub-object 1–1181
primitives from the keyboard 1–169
ruled surface 1–1193
shapes 1–262
skin 2–1076
sub-objects 1–1177
surface edge curve 1–1177
surface offset curve 1–1167
surface sub-objects 1–1177
surface-surface intersection curve 1–1166
tendons 2–1096
transform curve 1–1157
transform surface 1–1182
U loft surface 1–1196
UV loft surface 1–1200
vector projected curve 1–1171
creation method rollout 1–354
creation parameters 2–844
glossary 3–925
cross fade compositor (video post) 3–381
cross section editor viewport options 2–1141
cross sections 1–282, 1–284 to 1–288, 1–623, 1–842, 2–1108, 2–1114, 2–1130, 2–1141, 2–1147, 3–925
and bulges 2–1096
and envelopes 2–1088
and tendons 2–1147
initialization 2–1125
parameters 2–1114
view (bulge editor) 2–1096
cross-hairs cursor 3–815
crossing selection 1–93
CrossSection modifier 1–623
crowd
behaviors 2–1159, 2–1164
creating crowd systems 2–1155
crowd object 2–1157
definition 3–925
delegate object 2–1159
keyboard shortcuts 2–1182
with bipeds 2–1172
crowd animation
and bipeds 2–1187
ClipState dialog 2–1253
motionclip parameters dialog 2–1252
orientation behavior 2–1214
path follow behavior 2–1216
repel behavior 2–1218
scripted behavior 2–1220
seek behavior 2–1220
speed vary behavior 2–1222
surface arrive behavior 2–1223
surface follow behavior 2–1226
synthesis dialog 2–1246
user interface 2–1182
vector field space warp 2–1241
wall repel behavior 2–1227
wall seek behavior 2–1229
wander behavior 2–1231
working with 2–1154
crowd behaviors
path follow 2–1162
seek 2–1162
wall seek 2–1162
crowd helper object 2–1187, 3–925
crowd object
associate bipeds with delegates dialog 2–1199
behavior assignments and teams dialog 2–1200
behavior rollout 2–1211
collisions rollout 2–1240
edit multiple delegates dialog 2–1197

geometry rollout 2–1240
global clip controller 2–1241
icon size 2–1240
object/delegate associations dialog 2–1196
priority rollout 2–1235
scatter objects dialog 2–1189
select behavior type dialog 2–1205
select delegates dialog 2–1205
set start frames dialog 2–1237
setup rollout 2–1188
smoothing rollout 2–1238
solve rollout 2–1232
state dialog 2–1207
state transition dialog 2–1208
crowd simulation
avoid behavior 2–1211
cognitive controller 2–1206
solving 2–1168
crowd system 3–925
CS amplitude option (links) 2–1091, 2–1136
CSM files 2–919, 2–1061, 2–1065, 2–1263, 3–925
cube method 2–1242
cubic morph controller 2–300
CUI files 3–804 to 3–806
current bulge angle 2–1114, 2–1141
current event (video post) 3–324
current frame 3–701, 3–724
current value editor 2–533, 2–565
currentdefaults.ini 3–790
currently assigned links only 2–1150
currently installed driver 3–821
curve editor 2–501, 2–507, 2–535, 2–1002
display menu 2–530
modes menu 2–521
curve view 3–925
curves
curve approximation 1–1238
curve fit 1–1157
curve point 1–1220
curve-curve intersection point 1–1223
curve-curve point 1–1223
freeze non-selected 2–587
function 2–837, 2–1008
curves menu
Track View 2–525
curves toolbar 2–535
custom attributes 1–129
custom grid 2–20
custom splash screen 1–17
custom UI and defaults switcher 3–789
custom UI scheme 3–804
customize
keyboard shortcut 2–1432
toolbar 2–1432
transitions 2–1034
customize menu 3–683
configure system paths 3–810
configure user paths 3–808
customize user interface 3–792
grid and snap settings 2–41
load custom UI scheme 3–805
lock UI layout 3–788
plug-in manager 3–788
preferences 3–815
revert to startup UI layout 3–807
save custom UI scheme 3–806
show UI 3–788
viewport configuration 3–853
customize user interface 3–792
colors 3–799
keyboard shortcuts 3–793
load UI scheme 3–805
lock UI layout 3–788
menus 3–798
overview 3–785
quad menus 3–795
revert to startup layout 3–807
save UI scheme 3–806
shortcuts 3–793
toolbars 3–794
customizing biped characters in figure mode 2–847
customizing user interface
Track View 2–599
cut
and slice 1–1011, 1–1019
and snaps 1–1019
time (Track View) 2–567
cut time (Track View) 2–567
cutout mapping 2–1540, 2–1542
cutting hair 1–529
CV 3–926
CV curve
CV curve 1–1110
CV curve (glossary) 3–926
CV curve on surface 1–1172
CV curve sub-object 1–1153
CV sub-objects 1–1085
CV surface
CV surface 1–1103
CV surface (glossary) 3–926
CV surface sub-object 1–1179
CWS file 3–135
CWS file (combustion workspace) 3–611
cycling
animation 2–551

cylinder
chamfer 1–192
CylGizmo 3–306
standard primitive 1–177
cylindrical area omni light 2–1298

D

damper 1–396, 2–66
damping joint action 2–466
dashpots
angular 2–732
linear 2–730
data files
path for 3–813
data management
asset tracking dialog 3–487
open from vault 3–389
daylight
IES sky 2–1312
IES sun 2–1309
DDS files (glossary) 3–611
deactivate all maps 1–50
deactivate footsteps 2–865, 2–990
decay 2–1345
default
controller settings 2–294, 3–828
controllers 2–294
heights 2–1
keyboard shortcuts 1–900, 2–140, 2–510, 2–1006, 2–1111, 2–1182, 3–871
lighting 2–1272, 2–1274
material settings 2–1442
tangent types 3–721
default color 2–1141
default scanline rollout 3–38
default tangent types 3–721
defaults
setting 3–790
setting and changing 1–17
switching 3–789
define
append 2–1048
inser below 2–1048
insert above 2–1048
script 2–1048
defining
search terms (HTML help viewer) 3–876
time tags 3–710
deflector space warp 2–90
deflector space warps
deflector 2–90
PDynaFlect 2–81
POmniFlect 2–78
SDeflector 2–87
SDynaFlect 2–85
SOmniFlect 2–84
UDeflector 2–89
UDynaFlect 2–86
UOmniFlect 2–85
deform
deform bevel 1–366
deform fit 1–367
deform scale 1–364
deform teeter 1–365
deform twist 1–364
deformable bodies 2–777
cloth 2–778
constraints 2–795
deforming meshes 2–793
ropes 2–789
soft bodies 2–783
soft selection 2–800
deformable constraints 2–795
attach to deforming mesh 2–799
attach to rigid body 2–798
fixing vertices 2–796
keyframe points 2–797
deformable envelopes 2–834, 2–1085, 2–1090, 2–1108, 2–1111, 2–1130, 3–927
deformation 2–1083, 3–927
deformation spline 2–1083, 2–1085, 2–1091, 2–1135, 3–927
deformations (and loft objects) 1–363, 1–368
deforming meshes 2–793
collection 2–794
degradation override 1–34
degree 1–1091, 3–927
degree of freedom and rotating links 2–891
delegate
adjusting parameters 2–1159
geometry parameters rollout 2–1183
motion parameters rollout 2–1183
delegates 2–1157
definition 3–928
helper objects 2–1183, 3–928
using bipeds with 2–1172
delete 1–95, 2–1141, 2–1147
biped 2–854
bulge angle 2–1114, 2–1141
bulge angles 2–1095
bulge cross sections 2–1096
clip/transition 2–1027, 2–1045
control points 2–1114
controller 2–549
cross section slice 2–1114
current event (video post) 3–324
footsteps 2–869, 2–990

key 2–904, 2–956, 2–962
keys 2–554, 3–703
layers 2–974
maps 2–1413
material 2–1413
mesh modifier 1–626
operator (particle flow) 2–146
patch modifier 1–627
Schematic View 3–652
script 2–1030, 2–1048
selected keys (track bar) 3–703
spline modifier 1–627
tendons 2–1096
time (Track View) 2–567
Track View 2–598
transition 2–1051
transitions 2–1034
delete keys 2–502, 2–554
delete selected animation 3–698
delete time (Track View) 2–567
delete Track View 2–598
delete twist pose 2–950
delete UVW mapping 1–933
deleting
blocks of time 2–567
editable mesh edges 1–1011
editable mesh vertices 1–1011
isolated editable mesh vertices 1–1011
particles 2–146
patch surfaces 1–968
splines 1–308
vertices 1–297
dent map 2–1667
dependencies (views menu) 1–47
dependent 3–928
dependent sub-objects 1–1087, 3–928
depot (particle view)
display of 2–129
glossary 3–928
depth of field 2–1383, 3–90, 3–101, 3–269
description panel (particle view)
display of 2–129
glossary 3–928
deselect all 1–88
Design Web Format 3–555
designing clothing (garment maker) 1–575
designing materials 2–1395
destination clip
transition editor 2–1051
destination directory 2–1075
destroy character 1–115
detach 2–1147
detach (assembly) 1–110
detach (group menu) 1–106
detach dialog (edit poly) 1–679
detach dialog (NURBS curve/surface) 1–1228
editable mesh vertices 1–1011
editable patches 1–986
patch surface 1–968
DGS material (mental ray) 2–1580
dgs material shader (mental ray) 2–1717
diagnostics
mental ray renderer 3–123
dialog
asset tracking 3–487
bevel polygons 1–1066
bitmap pager statistics 3–514
chamfer 1–1070
color selector 1–161
connect edges 1–1070
extrude polygons 1–1072
filter selected euler tracks (Track View) 2–564
flatten mapping 1–907
MAXScript debugger 3–783
mixer transition editor 2–638
normal mapping 1–908
pelt map parameters 1–909
pick nodes 2–641
pivot selection 2–959
prompts 3–498
relax tool 1–912
render UVs 1–914
shared motion flow 2–1039, 2–1056
track sets editor 2–591
unfold mapping 1–919
unwrap options 1–920
weight tool 1–807
XRef objects 3–397
dialogs
bitmap proxies 3–496
global settings and defaults for bitmap proxies 3–496
toggling 3–670
dielectric material shader (mental ray) 2–1719
different ambient and different diffuse materials
dialog 3–512
diffuse
diffuse color (glossary) 3–929
diffuse distribution 2–1323
diffuse level 2–1489
diffuse level mapping 2–1499
diffuse mapping 2–1498
roughness mapping 2–1500
diffuse map (baking) 3–147
diffuse parameters rollout 3–143
diffuse texture element rollout 3–143
direct manipulation mode 1–1022

Direct3D driver 3–838, 3–843 to 3–844
Direct3D driver setup dialog 3–843
directional parameters 2–1348
directories
  for network rendering 3–187
  mounting 3–188
  sharing 3–188
DirectX 10 3–847
DirectX 9 shader material 2–1613
DirectX 9 shaders, FX file 3–946
DirectX shaders 2–1464, 2–1613 to 2–1614
disable layer 2–325 to 2–326
disable particle system 2–121
disabling playback 2–1084
disassemble 1–110
disc (circular) area light 2–1299
displace
  disp approx modifier (OSM) 1–628
  displace mesh (world space) 1–514
  displace modifier (OSM) 1–629
  displace NURBS (world space) 1–515
  displace space warp 2–76
displacement mapping 1–628, 2–1511, 2–1539
displacement shading
  mental ray renderer 3–96
display
  backface cull 3–775
  coordinate display 3–708
  cross-hairs cursor 3–815
  display controls for NURBS models 1–1117
  display driver (specifying at startup) 3–671
  display floater 3–775
  display floater (Schematic View) 3–651
  display image 3–502
  display operator (particle flow) 2–202
  display performance 1–28
  display plane 2–6
  display properties rollout 1–55
  grid settings 3–709
  hide/unhide (glossary) 3–951
  hide/unhide objects 3–775
  key bracket display 3–828
  layer properties 3–656
  line parameters for NURBS surfaces 1–1119
  marker 2–1065
  nth frame 3–821
  NU scale warning 3–815
  options/preferences 2–847, 2–853, 2–931, 2–944, 2–1090, 2–1130, 2–1141
  properties 3–775
  reflectance 2–1430
  selection floaters 1–79
  stack collapse warning 3–815
  topology-dependence warning 3–815
  track bar 3–707
  trajectories 2–931
  transmittance 2–1430
  world axis 3–821
display color rollout 1–52
display driver setup dialog 3–838
display menu
  curve editor 2–530
  particle view 2–129
display menu (Schematic View) 3–644
display panel 3–775
  display color rollout 1–52
  display properties rollout 1–55
  freeze rollout 1–54
  hide by category rollout 1–52
  hide rollout 1–53
  link display rollout 1–58
  object display 1–51
display rollout 1–791, 2–812
  hair and fur modifier 1–549
display subtree 2–1130
display trajectories;trajectories,displaying 2–957
displaying
  links 2–421
  selected key statistics (Track View) 2–595
  selected keys 2–594
distance
  distance from origin (accuracy setting) 3–815
  measuring 2–13, 2–15
distributed bucket rendering rollout 3–124
distributed maps 3–124
distributed rendering 3–124, 3–1001
distribution
  materials 2–1432
dithering (glossary) 3–930
divide
  edges 1–1019
  editable mesh edges 1–1011
  faces 1–1011
  segments 1–303
divisions 2–1114, 2–1141
docking 3–930
documentation for 3ds Max 1–xiv
DOF 2–1383, 3–269
dolly
  camera 3–746
  light 3–751
  target 3–746, 3–751
don't affect children 2–489
donut 1–276
doors 1–210, 1–246
  bifold 1–252

pivot 1–251
sliding 1–251
dope sheet 2–501, 2–507, 2–1002
modes menu 2–521
toolbars 2–538
Dope Sheet
editing footstep timing 2–869
DOS
command-line rendering 3–211
double support 2–988, 2–992, 3–930
double-sided 3–901
double-sided material 2–1591
download options (asset browser) 3–515
drag and drop
and copied/instanced maps 2–1451
and instanced objects 3–456
content from web pages 3–523
maps and materials 2–1423
modifier 1–499
sub-object material assignment 2–1424
with i-drop indicator 3–523
drag space warp 2–66
draw control points 2–1114
draw in profile view 2–1114
draw links as lines 3–821
drawing aids 2–1, 2–52
DRF files 3–527, 3–529
driver setup/configuration 3–838, 3–840 to 3–841, 3–844
dummies 2–922
dummy helper object 2–16
dummy object (glossary) 3–930
dummy objects
using 2–429
duplicate name dialog (material library) 2–1453
DWF
exporting 3–555
DWG files 3–931
exporting 3–550
importing 3–536
DWG/DXF import options dialog 3–536
geometry panel 3–539
layers panel 3–544
spline rendering panel 3–545
DXF files 3–931
exporting 3–552
importing 3–536, 3–551
dynaflectors 2–81, 2–85 to 2–86, 3–931
dynamic names (particle flow) 2–131
dynamics 2–122, 3–916, 3–932
and footsteps 2–1002
dynamics blend 2–945, 2–954, 3–932
hair and fur modifier 1–520, 1–540, 1–545
options 2–846
dynamics & adaptation rollout 2–980
dynamics deflectors 2–81, 2–85 to 2–86
dynamics objects 1–395
dynamics rollout
hair and fur modifier 1–545
dynamics utility 2–686
dynamics properties rollout (Material Editor) 2–686, 2–1479
edit object dialog 2–696
edit object list dialog 2–700

E

ease curve
applying 2–584
deleting 2–585
enable toggle 2–585
glossary 3–932
ease options
key info rollout 2–958
transition editor 2–1051
ease out-of-range types (Track View) 2–585
edge - definition 3–932
edge count 1–1253, 3–861
edge visibility threshold 1–1006
edges
aligning 1–1011
and rendering 3–826
attaching 1–1011
chamfer 1–1011
creating shapes from 1–1006
cut and slice 1–1011
deleting 1–1011
divide 1–1019
dividing 1–1011
extruding 1–1011
make planar 1–1011
rotating 1–1011
welding 1–1011
edit
alpha compositor (video post) 3–381
biped 2–1038
button appearance 3–803
clip 2–1048
contrast filter (video post) 3–343
cross fade compositor (video post) 3–381
current event (video post) 3–324
external event (video post) 3–340
fade filter (video post) 3–344
filter event (video post) 3–335
footsteps 2–936, 2–1000
freeform 2–1000
ghosts 2–1034
image alpha filter (video post) 3–344
image input event (video post) 3–332

layer event (video post) 3–337
lens effects filter (video post) 3–345
loop event (video post) 3–342
negative filter (video post) 3–345
normals 1–634
output image event (video post) 3–339
preset settings 3–438
pseudo alpha compositor (video post) 3–382
pseudo alpha filter (video post) 3–346
range bar (video post) 3–327
ranges (Track View) 2–573
scene event (video post) 3–329
simple additive compositor (video post) 3–383
simple wipe compositor (video post) 3–383
simple wipe filter (video post) 3–347
starfield filter (video post) 3–347
tag 3–711
time 2–566, 2–1002
time tag dialog 3–711
transition 2–1048
transitions 2–1034
edit commands 1–94, 2–1130
and envelopes 2–1088
edit curve on surface dialog 1–1229
edit geometry rollout
edit poly modifier 1–673
editable mesh 1–1011
editable poly 1–1055
edit keys 2–1002
edit keys (Track View) 2–528, 2–554, 2–559
edit keys mode 2–508
edit macro button dialog 3–803
edit menu 3–673
clone 1–476
delete 1–95
edit named selections 1–84
fetch 1–95
hold 1–95
move 1–439
object properties 1–117
particle view 2–127
region 1–92
rotate 1–439
scale 1–440
select all 1–87
select by 1–88
select by color 1–88
select by name 1–88
select invert 1–88
select none 1–88
select region crossing 1–93
select similar 1–88
selection method 1–92
transform type-in 1–431
undo/redo 1–94
edit menu (Schematic View) 3–642
edit modifiers
and editable objects 1–506
edit mesh modifier 1–634
edit patch modifier 1–638
edit poly modifier 1–640
edit spline modifier 1–680
edit multiple delegates dialog 2–1197
edit named selections 1–84
edit normals
edit normals modifier 1–634
edit object dialog 2–696
edit object list dialog 2–700
edit poly
align geometry dialog 1–679
bevel polygons dialog 1–1066
border 1–663
bridge borders/polygons dialog 1–1067
bridge edges dialog 1–1068
chamfer dialog 1–1070
connect edges dialog 1–1070
detach dialog 1–679
edge 1–656
extrude edges dialog 1–1073
extrude polygons along spline dialog 1–1071
extrude polygons dialog 1–1072
extrude vertices dialog 1–1073
hinge from edge dialog 1–1073
inset polygons dialog 1–1074
meshsmooth selection dialog 1–1074
object 1–651
paint deformation rollout 1–1064
polygon/element 1–666
preserve map channels dialog 1–1075
relax dialog 1–1076
tessellate selection dialog 1–1077
vertex 1–652
weld dialog 1–1077
weld edges dialog 1–1077
weld vertices dialog 1–1077
edit poly modifier 1–640
edit geometry rollout 1–673
selection rollout 1–647
edit ranges 2–573
edit ranges mode (Track View) 2–509, 2–528
edit texture surface dialog 1–1230
edit time mode (Track View) 2–528, 2–566
edit time tag 3–711
edit track set 2–591
edit UVWs dialog 1–888
menu bar 1–895

edit wire 2–412
editable mesh 1–996, 3–932
aligning 1–1011
edge 1–1006
edit geometry rollout 1–1011
element 1–1009
exploding 1–1011
face 1–1009
object 1–1001
polygon 1–1009
selection rollout 1–999
vertex 1–1003
editable objects and edit modifiers 1–506
editable patch 1–968
attach 1–986
changing vertex type 1–986
deleting vertices 1–986
detach 1–986
edge 1–980
element 1–984
geometry rollout 1–986
handle 1–979
object 1–974
patch 1–981
vector 1–979
vertex 1–975
visibility of 1–989
editable poly 1–1022
bevel polygons dialog 1–1066
border 1–1044
bridge borders/polygons dialog 1–1067
bridge edges dialog 1–1068
chamfer dialog 1–1070
connect edges dialog 1–1070
edge 1–1035
edit geometry rollout 1–1055
extrude edges dialog 1–1073
extrude polygons along spline dialog 1–1071
extrude polygons dialog 1–1072
extrude vertices dialog 1–1073
glossary 3–933
hinge from edge dialog 1–1073
inset polygons dialog 1–1074
meshsmooth selection dialog 1–1074
object 1–1028
paint deformation rollout 1–1064
polygon/element 1–1048
preserve map channels dialog 1–1075
relax dialog 1–1076
selection rollout 1–1024
subdivision displacement rollout 1–1063
subdivision surface rollout 1–1060
tessellate selection dialog 1–1077
vertex 1–1029
weld dialog 1–1077
weld edges dialog 1–1077
weld vertices dialog 1–1077
editable spline 1–289, 1–842
and overlapping vertices 1–289
attaching to 1–303
general rollout (for object and sub-objects) 1–289
identification numbers and 1–289
object 1–295
rendering options 1–289
segment 1–303
setting vertex type 1–297
spline 1–308
vertex 1–297
vertex area selection 1–289
editing
action parameters (particle view) 2–131
active footsteps in time 2–871
animation 2–304
bones 1–411, 1–413
curve cv sub-objects 1–1127
curve sub-objects 1–1135
fins (bones) 1–413
footstep placement 2–869
footsteps 2–988
footsteps in time 2–869
footsteps in Track View 2–869
modifier stack 1–504
named selection sets 1–67
point sub-objects 1–1123, 1–1219
strokes 3–865
surface cv sub-objects 1–1130
surface sub-objects 1–1141
time (Track View) 2–566
time tags 3–711
wall objects 1–228
editing track sets 2–590
editor
cognitive controller 2–1206
edtiable patch
selection rollout 1–971
effects 3–217
auto secondary lens effects 3–238
blur lens effects 3–260
brightness and contrast lens effects 3–265
color balance lens effects 3–265
depth of field lens effects 3–269
effects (rendering menu) 3–218
effects panel 3–219
environment and effects dialog 3–217
file output lens effects 3–266
film grain lens effects 3–268

glow lens effects 3–226
lens effects 3–223
manual secondary lens effects 3–242
merging from other files 3–220
ray lens effects 3–234
rings lens effects 3–230
star lens effects 3–246
streak lens effects 3–250
elements 3–933
of rendered textures 3–146
rendering 3–130
ellipse 1–274
email notification
network rendering 3–196
rendering 3–33
emission rollout (particle view) 2–136
emit start/stop values, and frame rate 2–144
emitter (particle flow) 2–135, 3–933
empty flow operator 2–209
enable ease or multiplier curve toggle 2–585
enable layer 2–326
encapsulated PostScript files 3–612
end effector 2–1080
end effectors 2–437, 2–440, 2–463, 2–471, 3–933
animating 2–461
linking to parent 2–461
entering frames 2–594
entire link 2–1114, 2–1141
entities
AutoCAD 3–441
envelope parameters 2–1130
envelope sub-object 2–1125, 2–1128
exclude envelopes dialog 2–1126
envelopes 2–1130, 3–934
adjusting shapes 2–1086
and control points 2–1088
and cross sections 2–1088
and edit commands 2–1088
and weighted vertices 2–834
blending types 2–1086
choosing default fit 2–1083
choosing default types 2–1083
copying 2–1086
copying to mirrored link 2–1088
display options 2–1090
exclude for selected links 2–1126
excluding influence 2–1086
overview 2–1085
scaling size 2–1086
selecting 2–1086
types of 2–1085
updating display manually 2–1090
using transforms with 2–1086
working with 2–1090
working with both envelope types 2–1091
working with rigid 2–1091
environment 3–271
and raytrace materials 2–1514
environment map (glossary) 3–934
environment panel 3–272
exposure controls 3–293
environment and effects dialog 3–217
effects panel 3–219
environment panel 3–272
environment effect
fire 3–276
fog 3–282
volume fog 3–284
volume light 3–288
environment shader (mental ray) 2–1721
environments 3–217
environment and effects dialog 3–217
EPS files 3–612
error threshold (camera tracker) 2–677
errors 1–20
garment maker 1–622
euler filter 2–564
euler rotation 2–916, 2–948
controller 2–891
Euler XYZ rotation controller 2–318
euler/tangent 2–916
events (particle flow)
and action sequence 2–123
branching 2–123
event display 2–131, 2–133, 3–936
event level 3–936
glossary 3–935
inputs 3–957
local 3–963
notating 2–206
properties 2–134
events (video post) 3–325
every step update script (particle flow) 2–139
exclude
exclude left end point (Track View) 2–570
exclude right end point (Track View) 2–571
exclude/include lights 2–1283, 2–1335
exclude envelopes dialog 2–1126
exclude left end point (Track View) 2–570
exclude option 2–1086, 2–1126
exclude right end point (Track View) 2–571
excluding layers 3–438
excluding particles from lighting 2–121
execute network rendering 3–182
execute sequence (video post) 3–325
exit command (file menu) 3–503

expanding animation tracks 2–886, 2–888
expert mode 1–51
explicit axis keys 2–297, 2–357
explode
  assemblies 1–110
  editable mesh objects/sub-objects 1–1011
  explode angle threshold 1–1011
  groups 1–106
  objects into faces 2–105
  particle system 2–68
  splines 1–308
explode angle threshold 1–1011
exploding objects 2–68
explosion 2–68
explosions 2–120
export animation
  motion mixer 2–624
export selected (file menu) 3–486
exporting
  3D DWF 3–555
  3DS files 3–532
  Adobe Illustrator files 3–534
  animation 2–921
  ASCII files 3–534
  bones 3–580
  DWG files 3–550
  DXF files 3–552
  export (file menu) 3–486
  FBX 3–558
  HTR/HTR2 3–578
  IGES files 3–562
  m3g files 3–563, 3–565 to 3–566
  MTL 3–590
  OBJ 3–589
  objects 3–486
  selected objects 3–486
  Shockwave 3D files 3–580 to 3–581
  stereolithography 3–588
  STL files 3–588
  texture coordinates 1–914
  to IGES 3–563
  UVs 1–914
  VRML97 files 3–591
  W3D files 3–580 to 3–581
exporting materials 2–1407
expose transform
  helper object 2–17
exposetm 2–17
  helper object 2–17
exposure control 3–293
  automatic 3–295
  linear 3–296
  logarithmic 3–297
  pseudo color 3–300
expression controller 2–320, 2–324
expression evaluator 1–12
expression techniques 1–146
extended parameters rollout (materials) 2–1471
extended primitives 1–186
  C-Ext 1–200
  capsule 1–195
  chamfer box 1–191
  chamfer cylinder 1–192
  gengon 1–199
  hedra 1–187
  hose 1–206
  L-Ext 1–198
  L-Extrusion 1–198
  oil tank 1–194
  prism 1–205
  ringwave 1–202
  spindle 1–196
  torus knot 1–189
extended shapes
  angle 1–286
  channel 1–285
  tee 1–287
  wide flange 1–288
  wrectangle 1–284
extended splines 1–266
extents
  scene extents 3–1007
  zoom extents (particle view) 2–129
extents (glossary) 3–936
external event (video post) 3–340
external reference 3–1037
external reference, AutoCAD (glossary) 3–1036
extraction tolerance 2–1070
extras dope sheet toolbar
  Track View 2–541
extras toolbar 3–688
extrude
  editable mesh edges 1–1011
  extrude modifier 1–680
  extrude surface (NURBS) 1–1188
  faces 1–682
  faces and polygons 1–1011
  patches 1–986
extrude edges dialog 1–1073
extrude polygons along spline dialog 1–1071
extrude polygons dialog 1–1072
extrude vertices dialog 1–1073
extruded shapes 1–198, 1–200
eyedropper tool 2–1448

F

f-curves 2–837, 2–1008

f-stop 3–90, 3–101
face - definition 3–936
face extrude modifier 1–682
face/edge thresholds (optimize modifier) 1–748
faces
  assigning to smoothing groups 1–1009
  beveling and extruding 1–1011
  creating 1–1011
  dividing 1–1011
  tessellating 1–1011
faceted (glossary) 3–937
facial animation 2–1100
facial expression 1–729
fade filter (video post) 3–344
fade in/out (lights) 2–1345
falloff 2–1111, 2–1130, 2–1136, 2–1338
  falloff map 2–1670
  glossary 3–954
  light falloff 3–753
family elements
  from Revit 3–457
  rendering properties 3–457
fast adaptive antialiaser 2–1533
fast view display mode 3–853
favorite location dialog 3–516
favorites
  asset browser 3–516
  HTML help viewer 3–878
FBX
  exporting 3–558
  importing 3–558
FBX files 2–921
features
  crowd 2–839
feedback about the documentation 1–xiv, 3–874
fence selection region 1–90
fetch (edit menu) 1–95
FFD soft bodies 2–786
FFD(box) space warp 2–91
FFD(cyl) space warp 2–95
FFDs 2–834, 3–937
  and physique 2–1104
  FFD 2x2x2 1–683
  FFD 3x3x3 1–683
  FFD 4x4x4 1–683
  FFD modifier 1–683
  FFD select modifier 1–689
  FFD(box) modifier 1–685
  FFD(box) space warp 2–91
  FFD(cyl) modifier 1–685
  FFD(cylinder) space warp 2–95
FGM files 3–937
field of view
  flyout 3–741
  glossary 3–937
field-of-view
  field-of-view button 3–741
fields (glossary) 3–938
FIG (figure) files 2–936, 2–1070
  saving and loading 2–855
FIG files 2–1263
  loading 2–942
  saving 2–941
figure mode 2–834, 2–936, 2–982, 2–984, 3–939
figure structure 2–1070
file corruption 3–883
file formats 2–919, 2–1263
file i/o path configuration 3–810
file link 3–903
  advanced settings 3–431
  basic settings 3–429
  basics 3–416
  excluding layers 3–438
  file link settings dialog 3–428, 3–435
  including layers 3–438
  manager utility 3–422
  presets 3–429, 3–431
  tips for using 3–419
  working with drawing files 3–417
  xref resolution 3–439
file menu 3–386, 3–673
  archive 3–499
  exit 3–503
  export 3–486
  export selected 3–486
  file link manager 3–422
  file properties 3–500
  import 3–485
  load animation 3–474
  merge 3–463
  merge animation 3–466
  new 3–386
  open 3–387
  open recent 3–390
  replace 3–470
  reset 3–387
  save 3–390
  save animation 3–476
  save as 3–391
  save copy as 3–392
  save selected 3–392
  summary info 3–499
  view image file 3–502
  XRef objects 3–394
  XRef scene 3–407

file output 3–266
file properties 3–500
file types
BIP 2–920, 3–916
BVH 3–920, 3–969
CAL 2–1070
CSM 2–1065, 3–920, 3–925, 3–969
CWS 3–135
FIG 2–936, 2–1070
.mfe file 2–1045
MNM 2–1061, 2–1065
MOC 2–1065, 2–1070
PHY 2–1098, 2–1106, 3–994
STP 2–924
file-handling commands 3–386
files
backup and saving 3–819
compressed 3–819
finding 3–510
incremental saves 3–819
managing 1–15
mismatched units 3–852
motion flow editor 2–1041
preferences settings 3–819
recent in file menu 3–819
fillet
and editable splines 1–297
fillet curve 1–1164
fillet surface 1–1216
fillet/chamfer modifier 1–689
glossary 3–939
film grain effect 3–268
filter color mapping 2–1503
filter color/filter opacity (glossary) 3–939
filtering 2–837
euler tracks (Track View) 2–564
motion capture and marker data 2–1070
trackgroups 2–645
tracks in motion mixer 2–612
filtering bitmaps 3–939
filtering character animation 3–940
filters 3–98
add filter event (video post) 3–335
caustics 3–106
creating custom 1–68
edit filter event (video post) 3–335
environment backgrounds (viewports) 3–821
euler tracks (Track View) 2–564
filter (track bar) 3–703
filter combinations dialog 1–81
filtering selections 1–81
filters button (Track View) 2–541
filters dialog (Track View) 2–542
key 3–718
sampling 3–1005
filters panel 2–1023
final gather map (FGM file) 3–937
final gather rollout 3–111
final gathering 3–93, 3–111, 3–940
final step update script (particle flow) 2–139
find target test (particle flow) 2–218
finding errors 2–1012
fine-tuning envelopes 2–1088
fingers option 2–846
fins (bones) 1–404, 1–413
fire environment effect 3–276
first vertex 1–297, 3–941
fit 2–1130, 2–1147
fit (deformation) 1–367
fit to existing 2–1070
fix ambient utility 3–512
fix panel 2–1020
fix vertices constraint 2–796
fixed
transition editor 2–1051
fixed width text button 3–815
fixed window 1–258
fixing errors 2–837, 2–1012, 2–1020
fixing motion errors 2–1023
fixing problems 3–883
flag properties dialog (Material Editor) 2–1655
flag with black 3–826
flat mirror map 2–1695, 3–83, 3–942
flatten footsteps 2–1070
flatten mapping 1–898, 1–907
flatten sides 1–1011
flex modifier 1–691, 1–700
flexibility (neck and spine) 2–846
flip normals 1–166, 1–1009
float controllers 2–297
float limit controller 2–335
floaters 1–431, 3–775
bone tools 1–411
display floaters 3–775
Schematic View display floater 3–651
selection floaters 1–79
transform type-in 1–431
floating 3–930
floating bones 2–1082, 2–1110
flows (particle flow) 2–208
empty flow 2–209
glossary 3–942
standard flow 2–209
fluorescence (glossary) 3–942
flyouts 1–12
align 1–462

arc rotate 3–744
array 1–448
default tangent types 3–721
dolly camera/target 3–746
dolly light/target 3–751
field of view 3–741
flyout (glossary) 3–943
material ID channel 2–1444
orbit/pan light 3–755
quick render 3–17
select and scale 1–440
selection region 1–80
timing preferences 3–815
use center 1–445
zoom extents 3–740
zoom extents all 3–737
focus plane 3–90
fog
VRML97 helper 3–600
fog environment effect 3–282
folder
motion flow editor 2–1041
foliage 1–210, 1–214
follow object
binding to 2–461
glossary 3–943
follow/bank utility 2–653
foot states 2–833, 2–936, 2–965, 3–943
footcandle 3–955
footstep
adaptation 2–988
animation (glossary) 3–943
animation workflow 2–856
converting to freeform 2–885
creation 2–863
edge selection 2–1000
editing 2–988
keys 2–867
leg states 2–867
timing (gait parameters) 2–861
footstep creation 2–833, 2–936, 2–988
create multiple footsteps (jump) 2–997
create multiple footsteps (run) 2–995
footstep operations 2–990
footstep extraction 2–1070
using motion-capture filtering 2–1061
footstep keys
body horizontal 2–867
body turning/rotation 2–867
body vertical 2–867
footstep method 2–833
footstep mode 2–936, 2–988
footstep operations rollout 2–936, 2–990
footsteps 2–999
activating 2–865
airborne period 2–883
appending 2–863
bending path 2–869
convert to 2–999
creating 2–863
creating automatically 2–862
creating multiple 2–862
deleting 2–869
display 2–853
editing in time 2–869
editing placement 2–869
footstep mode dialog 2–1000
freeform period between 2–883
moving and rotating 2–869
saving 2–882
selecting in Dope Sheet mode 2–869
selecting in viewports 2–869
timing 2–869
timing gait parameters 2–861
footsteps method 3–943
force operator (particle flow) 2–204
force space warps
displace 2–76
drag 2–66
gravity 2–73
motor 2–61
path follow 2–71
PBomb 2–68
push 2–59
vortex 2–63
wind 2–75
forward kinematics 2–954, 3–944
and IK 2–435
manipulating hierarchies with 2–426
forward kinematics and bipeds 3–944
FOV
field-of-view button 3–741
fps 1–1253
fracture 2–770
tips 2–773
fragmentation (particle flow) 2–120
frame
transition editor 2–1051
frame rate 2–288, 3–725, 3–944
and emit start/stop values (particle flow) 2–144
frames (snapping) 2–554
frames per second 3–861
free area light 2–1309
free camera 2–1370
free key defaults 2–956

free lights
direct 2–1293
linear 2–1307
spot 2–1290
free-form deformation (FFD)
box 2–91
box/cyl modifier 1–685
cylinder 2–95
modifier 1–683
select modifier 1–689
freeform 2–999
animation 2–902
convert to 2–999
converting to footsteps 2–885
inserting period between footsteps 2–883
method 2–833
setting period in footstep animations 2–883
freeform animation 2–886, 2–1002, 3–945
freeform method 3–945
freeze
freeze rollout (display panel) 1–54
freeze/unfreeze (glossary) 3–945
freezing/unfreezing objects 1–70, 3–775
freeze non-selected curves (Track View) 2–587
frequently asked questions
reactor 2–821
frizz animation
hair and fur modifier 1–520, 1–540
frizz parameters rollout
hair and fur modifier 1–540
from z level 2–1070
full screen 3–738
function curve editor 2–507, 2–535
function curves 2–837, 2–1008, 2–1012
add keys mode 2–581
glossary 3–945
show tangents 2–582
Track View 2–578
funnel-like objects 2–63
fuse vertices 1–297, 1–303, 1–308
fusing (glossary) 3–946
FX files 3–946

G

G-buffer
glossary 3–946
layers (rendering preferences) 3–826
gait pattern 3–947
gait type 3–947
game engine - edit normals 1–634
gamma correction (glossary) 3–948
gamma preferences 3–824
garment maker
troubleshooting and error codes 1–622
garment maker modifier 1–607
user interface 1–613
using 1–575
general parameters (lights) 2–1331
general parameters rollout
hair and fur modifier 1–534
general preferences settings 3–815
general settings rollout
render to texture 3–157
generate colors 2–944
gengon 1–199
geographic location dialog 1–422
geometric primitives 3–948
geometric/deformable space warps
bomb 2–105
conform 2–103
FFD(box) 2–91
FFD(cyl) 2–95
ripple 2–102
wave 2–100
geometry
AutoCAD 3–441
AutoCADArchitectural Desktop 3–443
compound objects 1–313
doors 1–246
effect on cloth 1–576
extended primitives 1–186
file formats 3–523
geometric primitives 1–169
importing 3–524
loft object 1–352
standard primitives 1–170
types of 1–155
windows 1–253
geometry parameters rollout 2–1183
geometry rollout 2–1240
patch 1–986
GeoSphere 1–176
get material 2–1439
getting started 1–1
ghost
transition editor 2–1051
ghosts
ghost before/after current frame 3–821
ghost in wireframe 3–821
GI (global illumination) 3–51
GIF files 3–613
gizmo
box atmospheric apparatus 3–304
cylinder atmospheric apparatus 3–306
gizmo/center (glossary) 3–949
preferences 3–832
sphere atmospheric apparatus 3–307

tracker 2–671
types of 3–304
using transform gizmos 1–426
gizmos rollout 1–791
global and local exclude/include dialog (for raytraced maps and materials) 2–1531
global clip 2–1179, 2–1246
global clip controller 2–1241, 3–950
global event (particle flow) 3–949
global illumination 3–51, 3–61, 3–93, 3–106, 3–940
global illumination (mental ray) 3–80
global lighting (rendered environment) 3–272
global motion clip 2–1179
global raytracer settings dialog 2–1528
global settings and defaults for bitmap proxies dialog 3–496
global shadow parameters (Track View) 2–512
glossary 3–901
glossiness mapping 2–1502, 3–950
glow
render effect 3–226
go to rotation test 2–224
go to settings
Material Editor 2–1446 to 2–1447
time 3–707, 3–722, 3–724
goal (and IK chain) 2–440
goniometric diagrams 2–1326
grab viewport 1–35
gradients
gradient map 2–1650
gradient mapping 2–195
gradient ramp map 2–1652
lens effects gradient colors (video post) 3–381
lens effects gradient options (video post) 3–378
lens effects gradient types (video post) 3–379
lens effects gradients (Video Post) 3–377
graph editors menu 3–682
delete schematic view 3–652
delete Track View 2–598
new schematic view 3–638, 3–652
new Track View 2–597
saved schematic views 3–638, 3–653
saved Track View 2–599
Track View 2–598
Track View - curve editor 2–501
Track View - dope sheet 2–501
graphics driver setup dialog 3–838
GravAccel (gravitational acceleration) 2–846, 2–878, 2–980, 3–950
gravity 3–950
computation 2–878
gravity space warp 2–73
green
line 2–1114
rigid vertices 2–1150
grid
autogrid 2–7
grid and snap settings 2–41
grid nudge distance 3–821
grid setting display 3–709
home grid settings 2–49
options 2–46
snap override 2–45
snaps 2–41
user grids settings 2–51
grid method 3–129
grid method, raytrace acceleration 3–1000
grids 2–33, 3–951
activating 2–34
align to view 2–35
aligning to 1–1011
and resolution of patch model surface 1–991
grid and snap settings 2–41
grid helper object 2–20
show home grid 2–34
using 2–4 to 2–5
viewing 2–6
ground plane (and collision detection) 2–891
group
script 2–1048
group menu 1–104, 1–109, 3–674
attach 1–106
close 1–105
detach 1–106
explode 1–106
group 1–104
open 1–105
ungroup 1–106
groups 1–96
and assemblies 1–98
and attaching physique 2–1083
and selection sets 1–96, 3–674
closing nested groups 1–105
detach from 1–110
explode 1–106
smoothing 1–167
using 1–96, 3–674
grow 1–809
growth objects
hair and fur modifier 1–517
guide hairs
hair and fur modifier 1–518, 1–526

H

hair
brush for styling 1–529
compositing method 3–222
cutting 1–529

light attributes 2–1351
shadows 3–223
styling 1–526
styling rollout 1–526
hair and fur
and lighting 1–519, 3–220
rendering options 3–221
hair and fur feature
components 1–517
hair and fur modifier 1–516
animation 1–520, 1–540, 1–545
display rollout 1–549
dynamics 1–520, 1–540, 1–545
dynamics rollout 1–545
frizz animation 1–520, 1–540
frizz parameters rollout 1–540
general parameters rollout 1–534
growth objects 1–517
guide hairs 1–518, 1–526
instanced hair 1–523
kink parameters rollout 1–542
material parameters rollout 1–537
mr parameters rollout 1–540
multi strand parameters rollout 1–544
quad menu 1–532
selection rollout 1–521
splines 1–517
styling hair 1–518, 1–526
styling rollout 1–526
surfaces 1–517
tools rollout 1–523
user interface 1–521
hair and fur render effect 3–220
hair and fur render element 3–140
hair light attributes 2–1351
handle display size 3–822
HD IK solver 2–440, 2–461, 2–467, 2–491 to 2–492, 2–494 to 2–496
HD solver 2–496
HD Solver 2–461, 2–463, 2–468, 2–471, 2–491 to 2–492, 2–494 to 2–495
HDR file format 3–613
HDRI files 3–613
head object 1–111
head object (glossary) 3–951
hedra 1–187
height map (baking) 3–148
height map displacement shader (mental ray) 2–1722
height option 2–846, 2–984
helix 1–281
help 3–873
about HTML help 3–874
contents 3–874
favorites tab 3–878
index 3–874
search 3–874
searching for help topics 3–876
help menu 3–684
helper object 3–951
helpers 2–2, 2–13
atmospheric apparatus 3–304
camera match 2–1391
compass 2–27
cone angle 2–27
dummy 2–16
expose transform 2–17
exposetm 2–17
grid 2–20
luminaire 1–111
manipulators 2–27
plane angle 2–29
point 2–23
protractor 2–26
reactor 2–715
slider 2–31
standard 2–16
tape 2–24
VRML97 3–597
HI IK solver 2–440, 2–446, 2–449
IK display options rollout 2–458
IK solver properties rollout 2–456
IK solver rollout 2–453
sliding and rotational joints 2–459
HI Solver 2–456
hide 1–53, 2–1150, 3–951
hide attached nodes 2–1108
hide by category 1–52
hide reference geometry 1–772
hide/show all 2–944
hiding and unhiding 1–53
by category 1–52, 1–72
by selection 1–70
edges 1–986
editable spline vertices 1–297
hide rollout 1–53
hierarchical linkage 3–951
hierarchical subdivision surfaces 1–701
hierarchies
hierarchical linkage (glossary) 3–951
joint limits 2–421
navigating 2–425
terminology 2–416
using multiple 2–418
viewing 2–424
hierarchy of biped objects (Track View) 2–886

hierarchy panel 3–773
commands 2–487
IK 2–491
link info rollouts 2–499
pivot 2–487
hierarchy right-click menu (Track View) 2–516
hierarchy window (Track View)
placing selected objects 2–588
selecting by name 2–589
high dynamic range images 3–613, 3–621
high-resolution rendering 3–197
highlights
anisotropic 2–1492
Blinn 2–1493
metal 2–1494
multi-layer 2–1495
Oren-Nayar-Blinn 2–1493
Phong 2–1493
specular color 3–1014
hinge constraint 2–747
hinge polygons from edge dialog 1–1073
history list 3–390, 3–502, 3–641
history-dependent IK solver 2–440
history-independent IK solver 2–440, 2–446
hold (edit menu) 1–95
home grid 1–23
glossary 3–952
settings 2–49
using 2–4
views based on the world coordinate axes 1–23
hopping (dynamics of) 2–878
horizon (glossary) 3–953
horizontal (move key) 2–579
horizontal bezier handle control 2–582
horizontal text in vertical toolbar 3–815
hose 1–206
hosts file 3–124
hot (glossary) 3–953
hot keys 1–900, 2–140, 2–510, 2–1006, 2–1111, 2–1182, 3–871
hotspot 2–1338, 3–752, 3–954
how many 2–992, 2–995, 2–997
how to (NURBS)
fix objects 1–1098
improve performance 1–1099
make things 1–1094
how tos 2–1264
HSDS modifier 1–576 to 1–577, 1–701, 1–706
HSV (glossary) 3–1001
HTML help viewer
favorites tab 3–878
keyboard shortcuts 3–879
right-click menus 3–879
searching in 3–876
toolbar 3–878
using 3–874
HTR/HTR2
exporting 3–578
importing 3–576
hue/saturation/value (glossary) 3–1001

I

i-drop Indicator 3–523
IAM files
importing 3–552
ICB targa files (glossary) 3–633
icons
color scheme 3–806
path for additional 3–813
reactor 2–707
ID
material ID channel 2–1443
IES 2–1328
IES sky 2–1312
IES sun 2–1309
IFL files 3–616
and view file command 3–5, 3–502
IFL manager utility 3–619
image file list control dialog 3–618
IGES
and NURBS surfaces 3–558
export/import log file 3–560, 3–562
exporting to 3–563
file translation 3–558
glossary 3–954
history 3–558
IGES import dialog 3–560
import table to 3ds Max 3–561
log files 3–560
overview 3–558
temporary files 3–560
ignore animation range 2–549
ignore backfacing 1–996, 1–1011, 1–1019
IK
and control objects 2–435
and set key 2–281
animating with interactive IK 2–480
IK joints 2–437
IK solution (glossary) 3–955
overlapping chains 2–446
preferences 3–830
IK blend 2–959, 3–954
IK constraints 2–900, 2–902
IK limb solver 2–440, 2–472
IK object 2–954
IK only option 2–980

IK rollouts 2–491
auto termination 2–499
display options 2–458
IK solver rollout 2–453
inverse kinematics 2–497
object parameters 2–491
spline IK solver rollouts 2–478
IK solvers 2–440, 2–446, 2–453, 2–461, 2–472
illegal video colors 2–1434
illuminance 3–955
image alpha filter (video post) 3–344
image file formats 3–608
image file list
IFL control dialog 3–618
IFL manager utility 3–619
image filter event (video post) 3–335
image input event (video post) 3–332
image input options (video post) 3–334
image layer event (video post) 3–337
image motion blur (glossary) 3–955
image output event (video post) 3–339
image sequence 3–5
images (2D) 3–608
import
animations 3–466
file menu 3–485
IGES files 3–560 to 3–561
import options 3–586
merge animation 3–466
importing
3DS files 3–530
Adobe Illustrator 88 files 3–533
and attaching 3ds Max objects 1–1120
animation 2–921
DDF 3–571
DEM Models 3–571
DWG and DXF files 3–536
DXF files 3–551
FBX 3–558
HTR/HTR2 3–576
IAM files 3–552
IGES files 3–560
IPT files 3–552
landXML 3–571
marker file 2–1061
motion-capture file 2–1061
PRJ files 3–531
scenes 1–16
SHP files 3–533
STL files 3–586
TRC 3–577
VRML files 3–591
importing geometry 3–524
merge or replace scene 3–524
importing motion-capture data 2–925
IMSQ files 3–620
in
tangent 3–721
in place mode 2–930, 2–936, 3–956
using to adjust keyframes 2–930
include new bones 2–1111
include/exclude lights 2–1283
including layers 3–438
incremental saves 1–19, 3–819
independent 3–558, 3–956
index of refraction 2–1471, 2–1509, 2–1514, 2–1538, 2–1670, 2–1703
indirect illumination 3–106, 3–994
influence 1–78 to 1–79, 2–1114, 2–1141, 3–407, 3–465, 3–957
areas of and envelopes 2–1085
influenced vertices 2–1147
inherit rollout 2–500
inheritance 2–434
initial graphics exchange specification (IGES) 3–558, 3–954
initial pose 2–1111, 2–1113, 2–1130, 2–1141, 2–1147, 2–1150, 3–957
initializing
and ActiveShade 3–904
initializing physique 2–1083, 3–957
ink 'n paint material 2–1605
inline (VRML VRML97 helpers) 3–608
inner envelope 2–1130
inner/outer bounds 2–1085
input devices for motion capture 2–655
inputs (particle flow) 3–957
insert 2–1141, 2–1147
actions, events (particle view) 2–133
bulge angle 2–1114, 2–1141
control points 2–1114
cross section slice 2–1114
insert animation 1–114
insert character 1–115
insert tracks dialog 3–466
time (Track View) 2–570
tracks 3–466
vertices 1–295, 1–308
insert keys 2–558
inset polygons dialog 1–1074
inside 2–1136
installing
3ds Max (for network rendering) 3–186
instance duplicate maps utility 2–1744

instanced hair
hair and fur modifier 1–523
instanced modifiers 1–511
instanced objects
AutoCAD 3–456
rendering properties 3–457
instances 1–472
glossary 3–957
make unique 2–575, 2–577
of maps 2–1451
overview 1–472
propagating materials 2–1432
propagation 2–1432
shape instance operator (particle flow) 2–178
instances in motion mixer 3–958
integration steps (particle flow) 2–139
intensity (light) 2–1276, 2–1279
intensity mapping 2–1539
intensity/color/attenuation parameters 2–1345
intensity/color/distribution rollout 2–1352
interactive IK 2–480
interactive manipulation mode 1–1022
interactive redraw 2–1141, 2–1147
interactive rendering 3–17, 3–21, 3–1030
interactive reshade 3–17, 3–22
interactive update (Track View) 2–526
internet
access 3–522
connection 3–504
internet download dialog 3–515
interparticle collision 2–243
interpolation 2–992, 2–995, 2–997, 3–958
stride 2–862
intersection 3–821
introduction
dynamics 2–707
inverse kinematics 2–435
NURBS modeling 1–1078
object selection 1–61
particle flow 2–109
physics 2–707
reactor 2–703
rendering effects 3–218
sub-object selection 1–74
to this reference 3–873
introduction to character studio 2–831
inventor files
importing 3–552
inverse kinematics 2–954
controlling precision 2–463
glossary 3–958
introduction 2–435
methods 2–439
preferences settings 3–830
rollout 2–497
terminology 2–437
inverse kinematics with bipeds 3–959
invert selection 1–88
IOR 2–1471, 2–1509, 2–1514, 2–1538, 2–1670, 2–1703
IPT files
importing 3–552
iso curves 1–1168
iso line (glossary) 3–959
isolate selection tool 1–73
isometric views 1–24
isotropic light distribution 2–1323
iteration setting (HD IK solver) 2–463

J

jambs 1–210
jitter (antialiasing control) 3–98, 3–1005 to 3–1006
job archives
network job assignment 3–199
network rendering 3–173
job dialogs (network rendering) 3–190
join dialogs (NURBS) 1–1232 to 1–1233
join to previous ik key option 2–960
joint intersections 2–1108
parameters 2–1094, 2–1098
joint intersections rollout 2–1140
joint rotation data (in BVH files) 2–1061
joints
activating joint axes 2–485
joint limits (hierarchies) 2–421
joint parameters 2–483, 2–495
joint precedence 2–467 to 2–468, 2–494
joint resistance and spring back 2–466
limiting joint action 2–486
path 2–483
rotational 2–483
setting joint precedence 2–467
setting joint resistance 2–466
setting parameters 2–483
sliding 2–483
sliding and rotational 2–496
surface 2–483
using default joint precedence 2–468
JPEG files (glossary) 3–620
JSR-184
editing parameters 3–566
export/import files 3–563, 3–565 to 3–566
log files 3–569
m3g player 3–570
texture parameters 3–565
JSR-184 player 3–570
jump 2–936, 2–988

jumping
dynamics of 2–878
parameters 2–861

K

KBD files 3–793, 3–804
keep apart operator 2–172
key filters 2–590, 3–718
key info
Bezier controllers 2–310
key info rollouts 2–304, 2–306
master track key info dialog 2–391
key info rollout 2–954
key interpolation 2–833
key mode 3–724
key modes (links) 2–430
key reduction
settings 2–1070
using motion-capture filtering 2–1061
key tangents toolbar 2–535
key tools toolbar 2–535
keyable icons 2–531
keyboard
additional commands 3–669
creating primitives from 1–169
keyboard entry rollout 1–169
keyboard panel (customize UI) 3–793
keyboard shortcuts 1–900, 2–140, 2–510, 2–1006, 2–1111, 2–1182, 3–793, 3–871
HTML help viewer 3–879
override toggle 3–872
keyframe interpolation 2–305
keyframe mode 3–717
glossary 3–960
keyframe vertices constraint 2–797
keyframes
adapting to edits 2–871
adjusting with in place mode 2–930
keyframing
the biped 2–833
keyframing tools 2–962
keys
adding 2–560
aligning 2–556
colors in Biped 2–947, 2–1005
create out of range 2–533
create out-of-range 2–562
default tangent types 3–721
delete 2–502, 2–554
editing 2–554
glossary 3–960
interpolating 2–305
key mode 3–724
key properties (track bar) 3–703
key statistics (Track View) 2–595
key time display (Track View) 2–594
master track 2–346
moving 2–558, 2–579
moving a group of 2–558
moving horizontal and vertical (Track View) 2–579
randomize 2–533
randomize utility 2–562
reducing 2–572
select 2–502
select by time 2–533, 2–563
soft selection manager 2–533
keys (setting) 2–904
keys menu
Track View 2–524
keys windows (Track View) 2–504
kinematic chains 2–437, 3–960
kink parameters rollout
hair and fur modifier 1–542
knot (glossary) 3–961
Kodak Cineon 3–610

L

L-Ext 1–198
L-Extrusion 1–198
l-type stair 1–232
landXML importer 3–571
landXML/DEM model import dialog 3–571
Large BSP method 3–129
lasso selection region 1–90
lateral ratio 2–622
lathe
lathe modifier 1–707
lathe surface (NURBS) 1–1190
lattice modifier 1–709
lattice parameters rollout 2–1242
launch script (glossary) 3–961
layer controller dialog 2–325
layer defaults 3–815
layer list 3–666
layer manager 3–656
layer properties 2–333
layer properties dialog 3–662
layer track 2–607
layers 1–117, 2–974, 3–655, 3–961
AutoCAD and 3ds Max 3–438
excluding in file linking 3–438
from AutoCAD 3–421
from Revit 3–421
glossary 3–961
including in file linking 3–438
layer event (video post) 3–337
layer list button 3–666
layer manager 3–656

layer properties dialog 3–662
select dialog 3–438
layers toolbar 3–688
add selection to current layer 3–667
create new layer 3–667
select objects in current layer 3–667
set current layer to selection's layer 3–667
layout (viewports) 1–26, 3–856
layout menu (Schematic View) 3–643
layout mode
glossary 3–961
leg link 2–984
leg states 2–867
legacy DWG import 3–547
length 2–990
transition editor 2–1051
lens effects 3–223
auto secondary 3–238
blur 3–260
brightness and contrast 3–265
color balance 3–265
depth of field 3–269
file output 3–266
film grain effect 3–268
glow 3–226
gradients 3–377
manual secondary 3–242
ray 3–234
ring 3–230
star 3–246
streak 3–250
lens effects (video post)
animating properties 3–349
automatic secondary flare parameters 3–356
flare 3–350
flare glow parameters 3–355
flare inferno parameters 3–360
flare preferences 3–353
flare ray parameters 3–358
flare ring parameters 3–355
flare star parameters 3–359
flare streak parameters 3–360
focus 3–362
glow 3–364
glow inferno 3–368
glow preferences 3–367
glow properties 3–365
gradient colors 3–381
gradient options 3–378
gradient types 3–379
highlight 3–370
highlight geometry 3–374
highlight preferences 3–376
highlight properties 3–371
manual secondary flare parameters 3–357
lens effects filters (video post) 3–345
lens size (cameras) 2–1373
level of detail
utility 1–1253
VRML97 helpers (LOD) 3–602
lift 2–936, 2–965, 3–962
leg state 2–867
lift dynamics 2–878
light distribution
diffuse 2–1323
isotropic 2–1323
spotlight 2–1324
web 2–1325
light include/exclude tool 2–1283
light lister 2–1285
light map 3–962
light painting rollout (radiosity) 3–70
light parameters
mental ray indirect illumination rollout 2–1343
mental ray light shader rollout 2–1345
light shader rollout 2–1345
light shaders
mental ray 2–1345
light tracer 3–44 to 3–45
light viewports 1–24, 1–33, 3–750
lighting
exclude/include dialog 2–1335
general parameters 2–1331
guidelines 2–1280
hair and fur 1–519, 3–220
in 3ds Max 2–1279
lighting analysis 3–76, 3–300, 3–628
lighting analysis dialog 3–76
lighting data exporter utility 3–303
lighting exclusion 2–121
lighting map (baking) 3–148
lighting parameters rollout 3–141
lighting texture element rollout 3–141
lightmap shader 2–1614
lights 2–1272, 2–1301
add default lights to scene 1–49
advanced effects rollout 2–1341
and atmospheres 2–1351
and effects 2–1351
and materials 2–1399
and shading 2–1399
and shadows 2–1279
animating 2–1282
atmospheres and effects for 2–1349
dolly 3–751
free area 2–1309

free direct 2–1293
free linear 2–1307
free point 2–1304
free spotlight 2–1290
light falloff 3–753
light include/exclude tool 2–1283
light lister 2–1285
mental ray shadow maps 2–1360
mr sky 2–1318
mr sun 2–1319
name and color rollout 2–1273
omni 2–1295
orbit/pan 3–755
photometric lights 2–1301
placing 1–7
positioning 2–1282
properties of 2–1276
roll 3–753
standard 2–1288
target area 2–1307
target direct 2–1292
target linear 2–1305
target point 2–1303
target spotlight 2–1289
truck 3–755
types of 2–1272, 2–1301
using 2–1274
viewport controls 3–752 to 3–753
working with 2–1274
lights name and color rollout 2–1273
Lightscape
export 3–572
import 3–573
Lightscape import
Lightscape Materials utility 3–574
lightscape material 2–1604
Lightscape Materials utility 3–574
limit controller 2–335
limiting animation ranges 2–335
limiting joint action 2–486
line 1–270
linear arrays (creating) 1–487
linear controller 2–341
linear dashpot 2–730
linear exposure control 3–296
linear light rollout 2–1354
link 2–1130, 2–1141, 2–1147
biped hand to an object 2–960
blending 2–1108
envelopes list (left side) 2–1126
length 2–1111
length as basis for envelope creation 2–1083
linking drawing files 3–1004
linking objects to biped 2–854
name 2–1150
scale 2–1091, 2–1136
sub-object 2–1135
to root attach node 2–1106
link constraint 2–403
link rollouts 1–58, 2–499
link info inherit 2–500
link info locks 2–500
link settings rollout 2–1136
link sub-object level
joint intersections rollout 2–1140
link settings rollout 2–1136
linkage, hierarchical 3–951
linked file states 3–422
linked objects
assigning materials to 3–445, 3–454
conversion settings 3–428, 3–435
selecting when file linking 3–440
linked XForm modifier 1–712
linking
and unlinking objects 2–421
animatable parameters 2–411 to 2–412
bones to follow objects 2–461
end effectors to parent 2–461
strategy 2–418
linking files 3–422
links 3–963
adding and deleting 2–430
adjusting parameters 2–1091
and joint settings 2–1111, 2–1113
and pivots 2–426
animating links 2–430
blending between 2–1083, 2–1085
changing link inheritance 2–434
displaying 2–421
link inheritance (selected) utility 2–435
main toolbar 2–422
moving 2–890
parameters 2–1091
radial scale parameters 2–1091
rotating 2–891
scaling 2–851
setting parameters 2–1091
sliding parameters 2–1091
twist parameters 2–1091
lip sync 1–729
list
layers 3–666
list controller 2–342
named selection sets 1–83
selection filter 1–81
transformation axis coordinate system 1–443

list views (Schematic View) 3–645
listener
  listener window (glossary) 3–963
  MAXScript listener 3–781
load 2–1070
  buffer only 2–1065, 2–1070
  file option 2–936
  marker name file 2–1065
  .mfe file 2–1032
  motion capture file 2–1065
  motion flow editor 2–1032
  parameters 2–1070
  specification 2–1106, 2–1123
load animation 3–474
load custom UI scheme 3–805
load envelopes 1–805
load/save presets rollout (PArray) 2–274
loading
  BIP files 2–942
  biped figure files 2–855
  biped step files 2–924
  FIG files 2–942
  motion files 2–920
  STP files 2–942
loading animation 3–472
local
  working folder 3–487
local biped curve 2–1012
local center during animate 3–828
local coordinate system (glossary) 3–963
local euler XYZ rotation controller 2–344
local event (particle flow) 3–963
local illumination 3–51
locate vertical center of mass keys 2–945
lock
  character 1–115
  lock selection 2–555
  lock time tag 3–710
  lock UI layout 3–788
  locking object transforms 2–433
  selection lock 3–707
lock assignments 2–1089, 2–1150
lock com keying 2–945
lock selection
  status bar 3–707
  Track View 2–555
locks rollout 2–500
LOD
  level of detail utility 1–1253
  thresholds 1–1256
  VRML97 helpers 3–602
loft object 1–352
  creation method rollout 1–354
  deform bevel 1–366
  deform fit 1–367
  deform scale 1–364
  deform teeter 1–365
  deform twist 1–364
  deformation dialog 1–368
  deformations 1–363
  path commands 1–372
  path parameters rollout 1–356
  shape commands 1–373
  skin parameters rollout 1–358
  surface parameters rollout 1–354
lofting
  glossary 3–964
  shapes 1–262
log file 3–124, 3–964
log files
  IGES 3–560
logarithmic exposure control 3–297
LogLUV format (TIFF files) 3–303
look at controller 2–344
look at object (particle flow) 3–964
look-at constraint 2–406
loop 1–809
looping 2–1061, 2–1070
  animation 2–551
  animation (Track View) 2–570 to 2–571
  loop event (video post) 3–342
low res environment background 3–821
low-polygon modeling 1–1252
lower bound 2–1147
LS colors 1–550
LS colors modifier 1–550
LS mesh modifier 1–713
LTLI files 3–964
lume shaders 2–1713
lumen 3–965
LumeTools shaders 2–1713
luminaire helper object 1–111
luminance 3–964
luminous flux (glossary) 3–965
luminous intensity (glossary) 3–965
LUT preferences 3–824
lux 3–955
LZF files 3–965
LZG files 3–965
LZH files 3–965
LZO files 3–965
LZV files 3–965

M

m3g files 3–563
m3g player 3–570
macro recorder (MAXScript) 3–782

macros
path for additional 3–813
MACUtilities 2–665
main toolbar 3–686
main window 1–9
make absolute 3–809
make controller/object unique (Track View) 2–550
make curve on surface dialog 1–1226
make loft dialog 1–1234
make material copy 2–1442
make point curve dialog 1–1235
make point dialog 1–1235
make preview 2–1434, 3–168
make relative 3–809
make selected same size (video post) 3–328
make unique 1–504, 1–511, 2–577, 3–770
Material Editor 2–1442
particle view 2–127, 2–133
manage scene states 3–518
manage scene states dialog 3–520
manager
brush presets 3–692
manager (network rendering) 3–182
managers (transform) 1–433
managing
files 1–15
scenes and projects 3–385
manipulator 2–27
manipulator helper objects
cone angle 2–27
plane angle 2–29
slider 2–31
manipulators
built-in 2–15
select and manipulate 2–15
manual secondary flares 3–242
manual update (envelopes) 2–1130
map animation 3–478
map track to track rollout 3–481
motion mapping parameters rollout 3–479
retargeting rollout 3–481
map channel info dialog 2–1738
map track to track rollout 3–481
mapped material
glossary 3–967
mapping
ambient color 2–1497
anisotropy 2–1504
bump 2–1506
coordinates (glossary) 3–967
cutout 2–1542
diffuse color 2–1498
diffuse level 2–1499
diffuse roughness 2–1500
displacement 2–1511
filter color 2–1503
flatten 1–907
glossiness 2–1502
map network drive dialog 3–188
mapping coordinates 2–1405
mapping operator (particle flow) 2–195
metalness 2–1506
normal 1–908
opacity 2–1503
orientation 2–1505
reflection 2–1508
refraction 2–1509
self-illumination 2–1502
shininess 2–1502
shininess strength 2–1501
specular color 2–1500
specular level 2–1501
unfold 1–919
mapping biped motion 2–921
maps 2–1662, 3–124, 3–503
2D 2–1624
3D 2–1662
activate all 1–50
camera map per pixel 2–1732
cellular 2–1664
checker 2–1638
color modifier 2–1692
combustion 2–1639
composite 2–1688
compositor maps 2–1687
custom 3ds Max mental ray shaders 2–1711, 2–1714, 2–1716 to 2–1717, 2–1719, 2–1721 to 2–1724, 2–1728 to 2–1730
cutout mapping 2–1542
deactivate all 1–50
deleting 2–1413
dent 2–1667
dragging and dropping 2–1423
falloff 2–1670
flat mirror 2–1695
glossary 3–968
gradient 2–1650
gradient ramp 2–1652
hierarchy (glossary) 3–970
light map 3–962
lume shaders 2–1713
map bias (glossary) 3–966
map channel (glossary) 3–966
map types 2–1617
mapped materials 2–240, 2–1445
maps rollout 2–1474

marble 2–1673
mask 2–1689
mental ray shaders 2–1712
mix 2–1689
noise 2–1674
normal bump 2–1731, 3–150
"other" (in the material/map browser) 2–1695, 2–1698 to 2–1699, 2–1703, 2–1711 to 2–1714, 2–1716 to 2–1717, 2–1719, 2–1721 to 2–1724, 2–1728 to 2–1732
output 2–1692
particle age 2–1675
particle MBlur 2–1676
Perlin marble 2–1677
planet 2–1678
procedural 3–997
projected 2–1341
raytrace 2–1698
reflect/refract 2–1699
reflection and refraction 2–1695
RGB multiply 2–1691
RGB tint 2–1693
show in viewport 2–1445
smoke 2–1679
speckle 2–1680
splat 2–1681
stucco 2–1682
swirl 2–1656
thin wall refraction 2–1703
tiles 2–1658
to enhance material 2–1403
transparency 2–1542
type button (Material Editor) 2–1449
vertex color 2–1693
waves 2–1683
wood 2–1684
MapScaler object-space modifier 1–713
MapScaler world-space modifier 1–551
marble map 2–1673
marker data 2–1061, 3–969
marker display dialog 2–1075
marker files 2–1065, 3–920, 3–969
importing 2–1061
name files 2–1061
marker name file dialog 2–1061
market-specific defaults 3–790
marking a contact object (particle flow) 2–183
mask map 2–1689
mask viewport to safe region 3–821
master block parameters dialog (block controller) 2–390
master clip 2–1179
master motion clip 2–1179, 3–969
master object 1–494
master point controller 2–346, 2–391
material
xref material 2–1616
material assignment
blocks 3–458
material attach options dialog (Boolean objects) 1–345
Material Editor 2–1409
bitmap 2–1631
dynamics properties rollout 2–1479
maps rollout 2–1474
material ID channel 2–1443
menu bar 2–1428
morpher material 2–1592
options dialog 2–1436
tools 2–1427
type button 2–1449
material ID
and attaching objects 1–1018
and Booleans 1–338
and editable meshes 1–1009
and editable patches 1–981
and editable splines 1–308
and particles 2–190
changing (particle flow) 2–191
glossary 3–969
material ID channel
flyout 2–1444
Material Editor 2–1443
material operators
dynamic 2–191
frequency 2–189
static 2–187
material parameters rollout
hair and fur modifier 1–537
material propagation 2–1432
material shaders rollout
mental ray material 2–1544
material to shader 2–1723
material xml exporter utility 2–1407
material/map browser 2–1412
material/map navigator 2–1447
materialbyelement modifier 1–716
materials 2–1395, 2–1400, 3–83
adding to library 2–1406
advanced lighting override 2–1601
and attaching objects 1–1018
and blocks 3–457
and particle array 2–239
and particle system events 2–124
and particle systems 2–240, 2–242
and set key 2–281
and styles 3–461
animating 2–1449

applying in particle flow 2–186
applying to an object 2–1405
applying to objects 2–1409
arch & design (mental ray) 2–1549
architectural 2–1535
Architectural Desktop 3–446
assign to selection 2–1441
assigning 3–445, 3–454
AutoCAD Architecture 3–445
blend 2–1588
blocks 3–458
car paint (mental ray) 2–1576
changing 3–446, 3–454
combined when attaching objects/splines 1–295, 1–1011
components 2–1399
composite 2–1589
compound materials 2–1587
copying 2–1409
default material settings 2–1442
deleting 2–1413
designing 2–1395
DGS material (mental ray) 2–1580
DirectX 9 shader 2–1613
double-sided 2–1591
dragging and dropping 2–1423
editable 2–1535
exporting 2–1407
get 2–1439
getting from library 2–1409
glass (mental ray) 2–1582
glossary 3–971
hierarchy (glossary) 3–970
ID channel 2–1443
ink 'n paint 2–1605
lightscape 2–1604
loading from scene 2–1409
make copy 2–1442
material blending (morpher material) 2–1592
material modifier 1–714
material name field 2–1448
material properties rollout (NURBS) 1–1149
matte/shadow 2–1584
mental ray 2–1544
morpher 2–1592
multi/sub-object 2–1594
name 2–1396, 2–1409
pick from object 2–1448
propagation 2–1432
put to library 2–1443
put to scene 2–1440
raytrace 2–1512
Revit 3–453 to 3–454
saving 2–1406, 2–1409
scene 2–1440
select by 2–1439
shell 2–1600
shellac 2–1597
show end result 2–1446
SSS materials (mental ray) 2–1583
standard 2–1465
subsurface scattering (SSS) materials (mental ray) 2–1583
top-bottom 2–1599
type 2–1397, 2–1457
type button (Material Editor) 2–1449
types of 2–1457
updating 2–1440
using 1–6
using maps to enhance 2–1403
matte object (glossary) 3–971
matte parameters rollout 3–141
matte texture element rollout 3–141
matte/shadow material 2–1584
max clips 2–649
MAX file finder utility 3–510
MAX files and Autodesk VIZ 3–525
max object
add to motion mixer 2–607
max objects to mix dialog 2–642
maximum angular/positional deviation for a track 2–1070
MAXScript 2–120
about MAXScript 1–xvii
and particle flow 2–208, 2–229
command-line 3–783
glossary 3–972
listener 3–781
MAXScript listener 3–781
menu 3–684, 3–780
mini listener 3–699
open MAXScript 3–781
preferences settings 3–834
run script 3–781
visual MAXScript utility 3–783
MAXScript debugger dialog 3–783
MAXScript menu 3–684, 3–780
macro recorder 3–782
MAXScript listener 3–781
new script 3–781
open script 3–781
run script 3–781
maxstart.cui file 1–12, 1–17
maxstart.max file 1–17
measure distance 2–15
measuring 2–13, 2–24, 2–52
melt modifier 1–717

memory management 3–514
memory use 3–129
mental ray
  add/edit DBR host dialog 3–128
  arch & design material 2–1549, 2–1562, 2–1569
  car paint material and shader 2–1576
  DGS material 2–1580
  distributed bucket rendering 3–128
  distributed bucket rendering rollout 3–124
  glass material 2–1582
  material 2–1544
  object properties 1–126
  satellite processors 3–128
  satellites 3–124
  subsurface scattering materials 2–1583
mental ray Connection rollout 2–1461
mental ray indirect illumination rollout 2–1343
mental ray light shader rollout 2–1345
mental ray material
  advanced shaders rollout 2–1548
  material shaders rollout 2–1544
mental ray materials 2–1543
mental ray messages 3–124
mental ray renderer 3–78, 3–940
  contour shading 3–96
  diagnostic tools 3–123
  displacement shading 3–96
  feature enhancements 3–84
  FGM file 3–937
  final gather map 3–937
  materials 2–1543
  messages window 3–87
  MI files 3–972
  object properties 1–126
  PASS file 3–990
  photon map 3–994
  preferences 3–837
  processing panel 3–86
  shadow map rollout 2–1360
  volume shading 3–95
mental ray shaders 2–1461, 2–1710, 2–1712
  3D displacement 2–1714
  bump shader 2–1716
  car paint shader 2–1576
  connect parameter to shader dialog 2–1713
  custom 3ds Max shaders 2–1711
  dgs material shader 2–1717
  dielectric material shader 2–1719
  environment shader 2–1721
  height map displacement 2–1722
  material to shader 2–1723
  mr physical sky 2–1321
  shader list 2–1723
  third-party shaders 2–1711
  uv coordinate 2–1728
  uv generator 2–1724
  uv generator parameters rollout 2–1725
  uv generator shaders rollout 2–1727
  XYZ coordinate 2–1730
  XYZ generator 2–1729
  XYZ generator parameters rollout 2–1729
  XYZ generator shaders rollout 2–1730
menu bar
  curve editor 2–521
  dope sheet 2–521
  Material Editor 2–1428
  particle view 2–126
  Track View 2–521
menus
  animation 3–681
  controller 2–521
  create 3–675
  customize 3–683
  edit 3–673
  file 3–673
  graph editors 3–682
  group 3–674
  help 3–684
  material editor copy and paste 2–1418
  MAXScript 3–684, 3–780
  menu bar 3–672
  menus panel (customize UI) 3–798
  modifiers 3–678
  particle view 2–126
  reactor 2–706, 3–681
  rendering 3–683
  Schematic View 3–642
  tools 3–674
  views 3–675
merge 3–463
  animation (file menu) 3–466
  custom sections 1–859
  effects 3–220
  insert tracks command 3–466
  merge dialogs 2–1453 to 2–1454, 3–406, 3–463, 3–465
  scenes 1–16
  shapes 1–859
merge from file
  sweep modifier 1–859
merge xref controller 2–383
merging characters 2–922
mesh
  editable mesh 1–996
  mesh select modifier 1–719
  meshsmooth modifier 1–722
  skin morph modifier 1–812

skin wrap modifier 1–818
skin wrap patch modifier 1–824
turbosmooth modifier 1–868
working with mesh sub-objects 1–998
mesh - definition 3–972
mesh conversion 2–206
mesh density, and cloth 1–577
mesh editing 1–935
mesh object (as Physique skin) 2–1076
mesh size (reducing) 2–1099
mesher object 1–374
meshing parameters rollout (radiosity) 3–67
meshsmooth modifier 1–722
MeshSmooth modifier and cloth 1–576
meshsmooth selection dialog 1–1074
messages 3–124
meta-operators
cache 2–197
metaballs 1–331, 3–972
metal bump shader 2–1614
metal highlights 2–1494
metal shader 2–1481
metalness mapping 2–1506
methods (IK) 2–439
.mfe file
append 2–1032
file types 2–1045
load 2–1032
save 2–1032
MFE files 2–652, 2–919, 2–1263
folder 2–1041
MI file 3–119, 3–124
MI files 3–972
middle button pan/zoom (viewports preferences) 3–821
MIDI time slider control (animation preferences) 3–828, 3–847
mini listener (MAXScript) 3–699
mini Track View (track bar) 3–703
minimum key spacing 2–1070
mirror 1–448, 2–913, 2–962, 2–1130, 2–1141, 2–1147
main toolbar 1–448
mirror curve (NURBS) 1–1160
mirror dialog 1–448
mirror modifier 1–728
mirror surface (NURBS) 1–1187
mirroring joint parameters 2–495
mirroring objects 1–491
splines 1–308
mirror dialog 1–448
mirror parameters rollout 1–791
mirror selected cross section 2–1114
mirrored link
copying envelope settings to 2–1088
mirrored UVs 3–155
mirroring 3–972
mirroring motion 2–913
missing external files dialog 3–503
missing map coordinates dialog 2–1623
missing XRef paths dialog 3–415
MIX files 2–935, 2–1263
mix map 2–1689
mix menu 2–629
mixdown 2–624
mixer
and BIP files 2–924
load mix 2–652
mode 2–652
rollout 2–652
save mix 2–652
mixer clip source options dialog 2–634
mixer transition editor dialog 2–636, 2–638
MNM file 3–973
MNM files 2–919, 2–1061, 2–1065, 2–1263
MNU files 3–795, 3–798, 3–804
mobile gaming
editing JSR-184 parameters 3–566
exporting JSR-184 files 3–563
MOC files 2–1065, 2–1070
modal (glossary) 3–973
mode
motion flow 2–1026, 2–1043
modeless (glossary) 3–973
modeling
objects 1–5, 1–842
modes 2–936, 2–982, 2–988
in place 2–930
mixer 2–652
rubber band 2–876
talent figure 2–1061
Track View 2–521
modes menu
curve editor and dope sheet 2–521
modifier list 3–758
modifier sets menu 3–771
modifier stack 3–760
collapsing 1–504
editing 1–504
glossary 3–973
modifier stack rollout 1–502
right-click menu 3–766
using 1–502
using at sub-object level 1–508
modifier-based space warps 2–107
modifiers 1–493, 1–497, 1–555, 1–557
affect region 1–557
and AutoCAD object transforms 3–442

and set key 2–281
and transforms 1–499
attribute holder 1–559
bend 1–560
bevel 1–562
bevel profile 1–565
camera correction 2–1392
camera map 1–513, 1–567
cap holes 1–569
cloth 1–578
conversion 1–871, 1–873 to 1–874
CrossSection 1–623
delete mesh 1–626
delete patch 1–627
delete spline 1–627
displace 1–629
displace mesh (world space) 1–514
displace NURBS (world space) 1–515
edit mesh 1–634
edit normals 1–634
edit patch 1–638
edit poly modifier 1–640
edit spline 1–680
extrude 1–680
face extrude 1–682
FFD 1–683, 1–685, 1–689
fillet/chamfer 1–689
flex 1–691
free-form deformation 1–683, 1–685, 1–689
garment maker 1–607
glossary 3–974
hair and fur 1–516
HSDS 1–701, 1–706
instanced 1–509, 1–511
lathe 1–707
lattice 1–709
linked XForm 1–712
list of 1–497
LS colors (world space) 1–550
LS mesh 1–713
make controller unique 2–550
MapScaler (object space) 1–713
MapScaler (world space) 1–551
material 1–714
materialbyelement 1–716
melt 1–717
mesh select 1–719
meshsmooth 1–722
mirror 1–728
morpher 1–729
multires 1–739
noise 1–743
normal 1–746
normalize spline 1–747
NSurf sel 1–747
object space 1–557
patch select 1–751
PatchDeform 1–552, 1–754
PathDeform 1–552, 1–755
point cache 1–555, 1–758
poly select 1–762
preserve 1–766
projection 1–769, 1–771 to 1–773, 1–775 to 1–777
projection holder 1–778
push 1–779
reactor cloth 2–778
reactor rope 2–789
reactor soft body 2–784
relax 1–779
renderable spline 1–781
ripple 1–783
select by channel 1–785
shell 1–785
skew 1–790
skin 1–791
skin morph 1–812
skin wrap 1–818
skin wrap patch 1–824
slice 1–825
smooth 1–828
spherify 1–829
spline IK control modifier 1–830
spline select 1–831
squeeze 1–833
STL check 1–834
stretch 1–836
substitute 1–840
surface 1–842
surface mapper (world space) 1–556
SurfDeform 1–557, 1–848
sweep 1–848, 1–857 to 1–858
symmetry 1–861
taper 1–863
tessellate 1–865
topology dependent 3–1023
trim/extend 1–866
turbosmooth 1–868
turn to mesh 1–871
turn to patch 1–873
turn to poly 1–874
turn-to modifiers 1–871, 1–873 to 1–874
twist 1–876
unwrap UVW 1–878
UVW map 1–922, 3–447, 3–455
UVW mapping add 1–933
UVW mapping clear 1–933

UVW mapping paste 1–934
UVW Xform 3–447, 3–455
UVW XForm 1–934
vertexpaint 1–936
volume select 1–952
wave 1–957
world space 1–512
WSM 1–512
XForm 1–959
modifiers menu 3–678
animation modifiers 1–557, 1–712, 1–754 to 1–755, 1–848
cache tools 1–758
free-form deformers 1–683, 1–685
mesh editing 1–569, 1–626, 1–634, 1–680, 1–682, 1–746, 1–748, 1–828, 1–834, 1–861, 1–865, 1–935 to 1–936
nurbs editing 1–557, 1–628, 1–848, 1–1101
parametric deformers 1–557, 1–560, 1–629, 1–709, 1–728, 1–743, 1–766, 1–779, 1–783, 1–790, 1–825, 1–829, 1–833, 1–863, 1–876, 1–957, 1–959
patch/spline editing 1–623, 1–627, 1–638, 1–680, 1–689, 1–707, 1–747, 1–842, 1–848, 1–857 to 1–858, 1–866
radiosity modifiers 1–555, 1–839
selection modifiers 1–719, 1–751, 1–831, 1–952
subdivision surfaces 1–722, 1–839
surface 1–628, 1–714, 1–716
UV coordinates 1–513, 1–551, 1–556, 1–567, 1–922, 1–934
modify child keys 2–529
modify child keys (Track View) 2–509
modify panel 1–499, 3–758
modify subtree (Track View) 2–509, 2–528
modifying
at sub-object level 1–506
multiple objects 1–509
NURBS models 1–1081
objects (basics) 1–153
morph 1–314
morph controllers 2–300
morpher material 2–1592
morpher modifier 1–729
morphing (glossary) 3–974
motion
combining BIP files 2–924
mapping 2–921
mirroring 2–913
Motion Analysis 2–665, 3–576 to 3–578
motion blending 3–975
motion blur 2–122, 2–240, 2–1386, 3–89, 3–101, 3–114, 3–269, 3–955, 3–981, 3–1007
and particle flow 2–191
motion capture 2–347, 2–655, 2–1059, 2–1064, 3–975
batch file conversion dialog 2–1075
buffer 2–1061, 2–1065, 2–1070
conversion parameters dialog 2–1070
converting data from buffer 2–1065
file 2–1070
importing files 2–1061
introduction to importing 2–925
rollout 2–1061, 2–1065
motion clip 2–1246, 3–975
motion clips
Track View pick dialog 2–1252
motion clips panel 2–1246
motion damper 2–66
motion editing 2–1012
motion files 2–919, 3–976
information saved in 2–920
loading motion files 2–920
samples 2–920
motion flow 2–837, 3–976
and BIP files 2–924
BIP file location 2–920
clip properties dialog 2–1045
compare with motion mixer 2–604
editor file 2–1045
graph 2–1027 to 2–1028, 2–1045
mode 2–936, 2–1026, 2–1030, 2–1043, 2–1045, 2–1048
optimize transition 2–1058
random motion 2–1035
random motion flow 2–1039, 2–1056
rollout 2–1045
script 2–1026 to 2–1027, 2–1048
shared 2–1056
shared motion flow 2–1039, 2–1056
transition 2–1028
unified motion 2–1038
workflow 2–1043
motion flow editor 3–976
append 2–1032
files 2–1041
load 2–1032
save 2–1032
motion flow graph
optimize 2–1058
motion flow mode 2–936
motion fow scripts 3–976
motion mapping parameters rollout 3–479
motion mixer 2–604
adding bipeds 2–607
adding max objects 2–607
adding tracks 2–607
adjusting balance 2–622
clip timing 2–615

cloning clips 2–611
compare with motion flow 2–604
editor 2–646
export animation 2–624
filtering biped parts 2–645
how to use 2–604
importing clips 2–609
menus 2–629
moving clips 2–611
optimize transition 2–641
preferences 2–651
replacing clips 2–611
reservoir 2–649
toolbar 2–642
trackgroups 2–612
transitions 2–616
user interface 2–628
weight curve 2–619
motion panel 2–301, 2–303 to 2–304, 2–306, 2–463, 2–933, 3–774
motion parameters rollout 2–1183
motion synthesis 2–1172, 3–977
global clip controller 2–1241
motionclip parameters dialog 2–1252
motor 2–765
motor space warp 2–61
mounting a directory (network rendering) 3–188
mouse sensitivity 3–821
MOV files 3–621
move 1–439, 2–936, 2–965
clip 2–1045
leg state 2–867
move keys (Track View)
curve editor 2–579
dope sheet 2–558
edit keys 2–558
function curves 2–579
moved pivot 2–959
movie window (camera tracker) 2–671
moving
actions, events (particle flow) 2–132
cameras 1–7
center of mass 2–876
keys 2–558
keys (biped) 2–1004
lights 1–7, 2–1282
links 2–890
through time 2–287
to first frame 3–722
to last frame 3–724
to next frame 3–724
to previous frame 3–723
to transform keyframes 3–724
moving biped keys 2–1004
MPEG files 3–621
mr parameters rollout
hair and fur modifier 1–540
mr physical sky shader 2–1321
MSP files 3–977
MTL
exporting 3–590
mtl files (wavefront) 3–588
multi strand parameters rollout
hair and fur modifier 1–544
multi-layer basic parameters 2–1481
multi-layer highlights 2–1495
multi-level shader 2–1504
multi-pass parameters (cameras)
depth of field 2–1383
motion blur 2–1386
multi-pass rendering effects 3–77
multi-pass rendering effects (cameras) 2–1382
multi/sub-object material 1–834, 2–242, 2–1594
multi-threading 3–826
multi-view
blocks 3–459
multicurve trim surface 1–1214
multiple biped links 2–895
selecting and rotating 2–895
multiple instanced objects 2–121
multiplicity (glossary) 3–977
multiplier (glossary) 3–977
multiplier curve
applying 2–584
deleting 2–585
enable toggle 2–585
glossary 3–978
multiplier out-of-range types (Track View) 2–586
MultiRes modifier 1–739
multiresolution adaptive antialiaser 2–1534
multisided blend surface 1–1213
multithreading and rendering 3–828
MVBlocks 3–459

N

N blend surface 1–1213
n links 2–1111, 2–1150
N links 3–978
name
object name 3–757
name and color rollout 3–757
for lights 2–1273
named selection sets 1–67, 1–83 to 1–84, 1–508
names
material 2–1396
selecting by (Track View) 2–589
naming layers 3–655

naming materials 2–1409
natural light 2–1280
navigating
3D space 1–21
blocks 3–460
camera and light views 1–33
hierarchies 2–416, 2–425
rendered panorama 3–173
viewports 3–735
navigating the workbench 2–1010
navigator (material/map) 2–1447
NavInfo (VRML97 helpers) 3–599
neck link 2–984
negative filter (video post) 3–345
nested expressions (HTML help viewer) 3–876
net render control (common parameters rollout) 3–79
network
working folder 3–487
network plug-in configuration 3–814
network rendering 3–173, 3–175, 3–1001
advanced settings 3–199
Backburner 3–173
configuration 3–175
email notification 3–196
error messages 3–183
glossary 3–979
how it works 3–180
installing 3ds Max for 3–186
job dependencies 3–196
job dialogs 3–190
job handling 3–199
manager 3–978
per-job timeouts 3–199
pre-render MAXScript 3–199
pre-render scripts 3–173
server (glossary) 3–979
set up 3–175
single computer 3–202
starting 3–182
TCP post number 3–199
troubleshooting 3–183
new
new command (file menu) 3–386
new Schematic View 3–652
new script 3–781
new sequence (video post) 3–323
new Track View 2–597
new feature 1–88, 1–163, 1–446, 1–500, 1–775, 1–1070, 2–289, 2–291, 2–293, 2–301, 2–332, 2–414, 2–425, 2–518, 2–533, 2–543 to 2–544, 2–561, 2–564, 2–596, 2–959, 2–1005, 2–1414, 2–1447, 3–6, 3–426, 3–462, 3–474, 3–537, 3–556, 3–697, 3–761, 3–847
new preset 3–437
new settings preset dialog 3–437
new Track View 2–597
newton (glossary) 3–979
next frame 3–724
next key 3–724
next key-previous key 2–955
next transition
transition 2–1051
next/previous 2–1130, 2–1141
next/previous key;finding, next/previous key 2–955
NGon 1–277
no blending 2–1111, 2–1150
no footsteps 2–1070
no key reduction 2–1070
node (glossary) 3–979
node track (glossary) 3–979
noise
and terrain effects 1–744
noise controller 2–353
noise map 2–1674
noise modifier 1–743
noise rollout (2D) 2–1630
noise threshold 2–1650, 2–1652, 2–1674, 3–282, 3–288
non-biped object 2–638, 2–641
non-vertical jambs 1–210
nonrelational NURBS surfaces 1–1116
nonscaling object size 3–821
normal bump map 2–1731
normal bump maps 3–150
troubleshooting 3–151
normal mapping 1–898, 1–908
normal projected curve 1–1169
normalize spline modifier 1–747
normalized 2–1150
normals 1–166
adjusting 1–166
aligning 1–465, 2–10, 2–488
editing 1–634
flipping 1–166
normal modifier 1–746
scaling vertex and face 1–996
unifying 1–166
viewing and changing 1–166
normals map (baking) 3–148
note keys 2–552
note track 2–552 to 2–553
notes
adding 2–197, 2–206
notes operator 2–206
notes dialog (parameter collector) 1–145
NSurf sel modifier 1–747
nth serial numbering 3–826
NTSC 3–826, 3–980

numbers
of links that can affect a vertex 2–834
show/hide all 2–944
numeric calculator 1–12
numerical expression evaluator 1–12
NURBS 2–834
and animation 1–1091
and modifiers 1–1089
animation tips 1–1099
blend curve 1–1158
blend surface 1–1183
cap surface 1–1195
chamfer curve 1–1161
concepts 1–1091
creating models 1–1094
curve approximation 1–1238
curve fit 1–1157
curve point 1–1220
curve sub-objects 1–1135
CV curve 1–1110
CV surface 1–1103
definition 1–1091
extrude surface 1–1188
fixing problems with models 1–1098
glossary 3–980
improving performance 1–1099
introduction 1–1078
lathe surface 1–1190
mirror curve 1–1160
mirror surface 1–1187
offset curve 1–1159
offset surface 1–1186
point 1–1219
point curve 1–1106
point point 1–1219
point surface 1–1102
ruled surface 1–1193
sub-object clone options dialog 1–1237
surf point 1–1222
surface approximation 1–1239
tips 1–1094, 1–1099
transform curve 1–1157
transform surface 1–1182
U and V iso curves 1–1168
U loft surface 1–1196
using toolbox to create sub-objects 1–1083
working with models 1–1080
NURBS curve/surface
detach dialog 1–1228
NURBS curves 1–1106
creating from splines 1–1115
fillet 1–1164
glossary 3–980
NURBS models 1–1078
creating 1–1079
creating sub-objects 1–1081
dependent sub-objects 1–1087
display controls for 1–1117
glossary 3–980
modifying 1–1081
objects and sub-objects 1–1078
overview 1–1080
sub-object selection 1–1084
working with 1–1080
NURBS surfaces 1–1101
and IGES 3–558
creating from geometric primitives 1–1116
display line parameters 1–1119
glossary 3–981
making rigid imported surfaces independent 3–558
surface approximation 1–1239
NURMS 1–722, 1–1003

O

OBJ
exporting 3–589
obj files (wavefront) 3–588
object bounding box 2–1111
object color dialog 1–159
object data flow 1–494
object display 1–51
object display culling 1–58
object fragmentation (particle flow) 2–120
object instance 3–981
object motion blur (glossary) 3–981
object motion inheritance rollout (PArray) 2–269
object parameters rollout 2–491
copying/pasting/mirroring joint parameters 2–495
position/orientation/bind to follow 2–492
precedence 2–494
sliding and rotational joints 2–496
object properties 1–117, 3–80
advanced lighting panel 1–123
cloth 1–602
edit menu 1–117
general panel 1–117
mental ray panel 1–126
user defined panel 1–127
object selection (introduction) 1–61
object space 3–982
object space (biped) 3–983
object space modifiers 3–983
object transforms 2–432 to 2–433
object/delegate associations dialog 2–1196
object-layer relationships 3–655
object-space modifier 1–557
MapScaler 1–713

objects 1–153, 2–960, 2–1282
aligning 2–8
arraying 1–484
binding 2–461
color 1–159
combining 1–338, 1–378, 1–388
copies/instances/references 1–472
creating 1–157
exporting 3–486
freezing and unfreezing 1–70
glossary 3–981
make controller unique (Track View) 2–550
modeling 1–5
modifying multiple objects 1–509
object properties 1–117
select and manipulate 2–15
select and move 1–439
select and rotate 1–439 to 1–440
selecting 1–61, 1–77
selecting by material 2–1439
techniques for cloning 1–474
using as bones 1–410
objects to bake rollout 3–158
obsolete file alert 3–390, 3–819
obstacle avoidance 2–1164
obstacle parameters rollout 2–1242
obstacle-avoidance behavior 3–983
odd/even 3–826
offset
offset curve 1–1159
offset point 1–1219
offset surface 1–1186
offset/absolute coordinate display 3–709
oiltank (extended primitive object) 1–194
omni light 2–1295, 3–983
omnidirectional light 3–983
omniflector 2–78, 2–84 to 2–85, 3–984
on selected objects/on all objects 2–1075
on/off controller 2–355
once time frame 2–141
online help
using HTML help viewer 3–874
online reference
introduction 3–873
searching in 3–876
using HTML help viewer 3–874
only extract footsteps within tolerance 2–1070
opacity 2–1488
falloff (glossary) 3–984
mapping 2–1503
open
assembly 1–109
file (file menu) 3–387
from vault 3–389
group (group menu) 1–105
new bitmap file 2–1635
particle view dialog 2–136
script (MAXScript menu) 3–781
video post sequence 3–323
open dialog 2–942
open from vault 3–389
open physique file 2–1106
open physique file button 2–1098
open recent 3–390
OpenEXR files
format 3–621
opening 3–626
saving 3–623
OpenGL driver 3–838, 3–841
opening screen 1–17
operands 1–338, 1–378, 1–388, 3–985
operator icon 3–985
operators 2–142
birth 2–143
birth script 2–145
delete 2–146
display 2–202
empty flow 2–209
force 2–204
glossary 3–985
keep apart 2–172
mapping 2–195
material dynamic 2–191
material frequency 2–189
material static 2–187
notes 2–206
operator time frames 2–141
position icon 2–147
position object 2–148
render 2–206
rotation 2–153
scale 2–156
script 2–208
shape 2–176
shape facing 2–176
shape instance 2–178
shape mark 2–183
speed 2–159
speed by icon 2–162
speed by surface 2–167
spin 2–154
standard flow 2–209
opposite tracks 2–945
optical markers 3–985
optimizations rollout 2–1361

optimize
motion flow 2–1058
motion mixer 2–641
transition 2–641, 2–1058
optimize modifier 1–748
optimize transition
transition editor 2–1051
optimizing performance (particle flow) 2–120
options 3–828
grid and snap 2–46
Material Editor 2–1436
rendering 3–826
viewports 3–821
options menu
Track View 2–526
options menu (particle view) 2–130
options menus (Schematic View) 3–644 to 3–645
orbit/pan
camera 2–1381, 3–749
light 3–755
Oren-Nayar-Blinn basic parameters rollout 2–1482
Oren-Nayar-Blinn highlights 2–1493
organic surfaces 1–842
orientation
changing 1–423
constraint 2–409
mapping 2–1505
orientation bar 2–1114, 2–1141
orientation behavior 2–1214, 3–985
origin (glossary) 3–986
origin point helper 3–431
origin slider 3–815
ortho snapping mode 2–38
orthographic view 3–986
orthographic views 1–24
out
tangent 3–721
out-of-range
keys (Track View) 2–562
types 2–551, 2–585 to 2–586, 3–987
outdoor lighting 3–45
outer envelope 2–1130
outline 1–308
output image event (video post) 3–339
output map 2–1621, 2–1692
output rollout 2–1621, 3–160
outputs (particle flow) 3–987
outside 2–1136
overlap 2–1111
overlapping IK chains 2–446
overlapping vertices and editable rollout (for object and sub-objects) 1–289
overlays
xref scenes 3–408, 3–412
overriding (degradation) 1–34
overshoot (glossary) 3–988
overview
cloth 1–571
clothing and pattern design 1–572
garment maker modifier 1–607
physique 2–834
workflow 2–839
overview of 3ds Max 1–1

P

pack UVs dialog 1–909
paint (vertexpaint modifier) 1–936
paint deformation rollout 1–1064
brush options 1–960
paint selection region 1–91
paint weights 1–960
paintbox
vertexpaint modfier 1–941
painter options 1–960
PAL 3–826, 3–988
palette
vertexpaint modifier 1–950
paletted 3–826
pan
panning views 1–29
particle view 2–129, 2–135
Track View 2–595
viewport controls 3–743
pan view 3–743
panels
create 3–757
customize UI 3–793 to 3–795, 3–798 to 3–799
display 3–775
hierarchy 3–773
modify (command panel) 1–499, 3–758
motion 3–774
render scene 3–61, 3–219
scripted utility 3–1008
utilities (command panel) 3–778
panorama exporter 3–170
render setup dialog 3–171
viewer 3–173
pants (cloth) 1–574
parallel projection 1–24
parameter (glossary) 3–989
parameter animation 2–121
parameter collector 1–138
menu bar 1–142
parameter collector menu bar 1–142
parameter collector notes dialog 1–145
parameter curve out-of-range types (Track View) 2–551

parameter editor 1–129
parameter space (glossary) 3–988
parameter wiring 1–104, 2–411 to 2–412
parameters
  custom attributes 1–129
  HD Solver 2–491
  notes 1–145
  parameter collector 1–138
  parameter collector menu bar 1–142
  parameter editor 1–129
  wiring 2–411
parameters panel (particle view)
  display of 2–129
  glossary 3–988
parameters rollout 1–791
parametric (glossary) 3–989
parametric stride length 2–992, 2–995, 2–997
parametric stride width 2–992, 2–995, 2–997
parent overlap 2–1130
PArray 2–256, 2–258, 2–260, 2–262, 2–268 to 2–271, 2–274
partial blending 2–1092, 2–1130
particle age map 2–122, 2–191, 2–1675
particle collision 2–243
particle deflector
  deflector 2–90
  SDeflector 2–87
  UDeflector 2–89
particle emission 2–145
particle flow
  frequently asked questions 2–118
  introduction 2–109
  keyboard shortcuts 2–140
  source 2–135
particle generation 2–145
particle level 3–990
particle MBlur map 2–122, 2–1676
particle motion blur 2–122
particle motion damper 2–66
particle rotation
  rotation operator 2–153
  spin operator 2–154
particle system 2–134, 2–209, 3–990
particle systems
  blizzard 2–251
  creating 2–238
  glossary 3–990
  overview 2–237
  PArray 2–256
  particle MBlur map 2–1676
  PCloud 2–253
  snow 2–246
  spray 2–244
  super spray 2–249
particle tests 2–210
  age 2–211
  collision 2–212
  collision spawn 2–215
  find target 2–218
  go to rotation 2–224
  scale 2–227
  script 2–229
  send out 2–230
  spawn 2–230
  speed 2–233
  split amount 2–234
  split selected 2–235
  split source 2–236
particle view
  depot 2–125
  description panel 2–125
  display tools 2–125
  event display 2–131
  introduction 2–109, 2–125
  menu bar 2–126
  open 2–136
  parameters panel 2–125
particles
  along a path 2–163
  and age 2–146
  and binding to space warps 2–121
  and deflector space warps 2–212, 2–215
  and materials 2–187, 2–189, 2–191
  and particle flow 2–110
  and stretch 2–122
  appearance when selected 2–138
  caching 2–197
  creating particle systems 2–238
  direction 2–123, 2–159, 2–162, 2–167, 2–172, 2–176, 2–178
  display in viewports 2–202
  emission 2–143
  emission rate 2–145
  generation 2–143
  geometry type 2–176, 2–178
  leaving a mark 2–183
  limiting life span 2–146
  mapping 2–195
  MBlur map 2–191
  motion blur 2–191
  number of 2–119, 2–145
  orientation 2–153
  parent particle 3–989
  particle age map 2–1675
  particle diagram 3–989
  particle emitter 2–239
  particle generation rollout (PArray) 2–260

particle motion blur 2–240
particle spawn rollout (PArray) 2–271
particle type rollout (PArray) 2–262
positioning 2–147 to 2–148
rendered as 2–206
rotation 2–153
scaling 2–156
size 2–176
spawn particles 3–1014
speed 2–123, 2–159, 2–162, 2–167, 2–172
spinning 2–154
synchronize animated bitmap textures 2–121
testing particle scale 2–227
unexpected spawning 2–122
viewport display 2–202
PASS file
mental ray renderer 3–990
paste 2–1141, 2–1147
a material, map, bitmap, or color 2–1418
paste controller (Track View) 2–545
paste footsteps 2–990
paste from buffer 2–1065
paste layer 2–325 to 2–326
paste posture 2–962
paste posture opposite 2–962
paste selected cs 2–1114
paste skin data dialog 2–700
paste tangent handles 1–297
paste time/track (Track View) 2–568
paste-pos tolerance 1–807
pasting 2–910
pasting joint parameters 2–495
patch (glossary) 3–991
patch grids 1–993
quad patch 1–994
tri patch 1–995
patch select modifier 1–751
patch surfaces 1–842, 1–993
copying 1–968
deleting 1–968
patch-based objects 3–991
PatchDeform
object-space modifier 1–754
world-space modifier 1–552
path
motion flow 2–1041
path constraint 2–398
and particles 2–163
path follow behavior 2–1162, 2–1216, 3–992
path follow space warp 2–71
path joints 2–485 to 2–486
PathDeform
object-space modifier 1–755
world-space modifier 1–552
paths
and AutoCAD xrefs 3–431
and particles 2–122, 2–163
and XRefs 3–411, 3–415
configuring paths 1–15
configuring system paths 3–810, 3–814
configuring user paths 3–808, 3–810 to 3–812
glossary 3–991
moving a camera along 2–1381
path commands (loft objects) 1–372
path follow space warp 2–71
path parameters rollout 1–356
pattern background 2–1433
pattern design (cloth) 1–572
PBomb space warp 2–68
PCloud particle system 2–253
PDynaFlect space warp 2–81
pelt map parameters dialog 1–909
per-bitmap resolution for bitmap proxies dialog 3–496
per-pixel camera map 2–1732
per-pixel camera projection 2–1732
percent snap 2–38
perform footstep extraction 2–1070
performance 3–106, 3–129, 3–1000
and biped's motion previewing 2–929
and weight painting 1–961
controlling display performance 1–28
improving in NURBS 1–1099
optimizing with physique 2–1099
while running 3ds Max 3–889
performance optimization (particle flow) 2–120
period 3–992
Perlin marble map 2–1677
perspective
and orthographic viewport controls 3–738
glossary 3–992
matching 2–1380
viewport control 3–747
perspective view 1–24
phases of leg motion 3–993
Phong highlights 2–1493
Phong shader 2–1480, 2–1482
photometric lights 1–7, 2–1272, 2–1301
area light sampling rollout 2–1354
common lamp values 2–1329
data file 2–1328 to 2–1329, 2–1355
example of photometric data file 2–1329
free area light 2–1309
free linear light 2–1307
free point light 2–1304
IES standard file format 2–1328
linear light rollout 2–1354

mr sky light 2–1318
mr sun light 2–1319
photometric webs 2–1326
preset lights 2–1302
target area light 2–1307
target linear light 2–1305
target point light 2–1303
web 2–1326
web parameters 2–1355
photometry 3–993
photon map 3–994
photon maps 3–93, 3–106
photorealistic renderer 3–38
PHY files 2–1098, 2–1106, 2–1263, 3–994
physical scale 3–295 to 3–297, 3–300
physique 2–834, 2–1080, 2–1083, 3–994
and changing geometry 2–1104
and FFDs 2–1104
and groups 2–1083
and other modifiers 2–1104
applying 2–1083
blending envelope display 2–1125, 2–1128
bulge angle display 2–1127
deformation spline 2–834, 2–1135
getting started with 2–1076
initializing 2–1083
joint intersections rollout 2–1140
keyboard shortcuts 2–1111
link settings rollout 2–1136
overview 2–834
reinitializing settings 2–1098
saving settings 2–1098
storing settings in PHY files 2–1098
tendon display 2–1128
user interface 2–1106
Physique
initialization 2–1123, 2–1125
physique initialization dialog 2–1083, 2–1111
physique level of detail rollout 2–1108
physique load specification dialog 2–1106, 2–1123
physique rollout 2–1106
Physique sub-objects 2–1129
PIC file format 3–613
PIC files 3–303, 3–628
pick material from object 2–1448
pick nodes dialog 2–641
pinch 2–1147
pinch bias 2–1147
ping-pong (playback direction setting) 3–723
pivot door 1–251
pivot points 2–487, 2–959
glossary 3–995
use pivot point center 1–446
using 1–509
pivot selection dialog 2–959
pivoted window 1–259
pivots
adjust pivot rollout 2–488
adjust transform rollout 2–489
adjusting 2–423
and links 2–426
resetting 2–423
pivots (IK extensions) 2–905
pixel 3–995
pixel data (rendered frame window) 3–7
place highlight 1–467, 2–1282
planar
constraints 1–437
make edges 1–1011
make vertices 1–1011
threshold 1–719, 1–996
plane 1–185, 2–764
plane angle manipulator 2–29
planet map 2–1678
plant 2–936, 2–965, 3–995
leg state 2–867
planted key defaults 2–956
plate match 3–38
plate match/MAX R2.5 antialiasing 1–567
play selected 3–723
playback 2–853, 2–936
real-time 2–197
playback speed 2–288
playing
animated material previews 2–1450
animation 3–723
preview 2–1434
plug-ins
color selector (general preferences) 3–815
glossary 3–995
plug-in manager 3–788
sharing over a network 3–814
system path configuration 3–814
plugin.ini file 1–17
PMAP file 3–995
PNG file (glossary) 3–628
point 1–1219, 2–1070
glossary 3–995
helper object 2–23
sub-object 1–1085, 1–1219
point cache modifiers 1–555, 1–758
point curve 1–1106
glossary 3–996
on surface 1–1175
sub-object 1–1155
point point 1–1219

point surface 1–1102
glossary 3–996
sub-object 1–1181
point-path constraint 2–762
point-point constraint 2–750
point3 XYZ controller 2–317
polar snapping mode 2–39
poly select modifier 1–762
polygon - definition 3–936
polygon count 1–1253, 3–861
polyhedra 1–187
POmniFlect space warp 2–78
ponytails 2–984
pose
copying and pasting 2–966
pose adjustment 2–1070
poses 3–996
adding 2–1096
copying between bipeds 2–910
reference 2–1076
posing bipeds 2–847
position
changing 1–423
ranges (Track View) 2–574
position constraint 2–401
position data (in CSM files) 2–1061
position operators
position icon 2–147
position object 2–148
position ranges (Track View) 2–574
position XYZ controller 2–356
position/orientation/bind to follow object 2–492
position/rotation threshold (IK) 2–463
position/rotation/scale (PRS)
controller (Track View) 2–357
parameters (motion panel) 2–303
positional markers 3–996
posture
copying and pasting 2–966
postures 2–910
power 2–1114, 2–1141
pre-calculating particle motion 2–120
pre-render scripts
advanced settings 3–173
command-line rendering 3–211
network rendering 3–190
precedence 2–494
child-to-parent 2–469
glossary 3–997
parent-to-child 2–470
setting manually 2–471
precedence, and keyboard shortcuts 3–872
precision and drawing aids 2–1
preferences 3–815
animation 3–828
asset browser 3–514
files 3–819
gamma 3–824
general settings 3–815
gizmos 3–832
inverse kinematics 3–830
MAXScript preferences 3–834
mental ray renderer 3–837
MIDI time slider control 3–847
motion mixer 2–651
preference settings dialog 3–815
radiosity settings 3–836
rendering 3–826
Schematic View 3–646
strokes 3–862, 3–867
viewports 3–821
preferences (display) 2–931
premultiplied alpha (glossary) 3–633, 3–997
preserve map channels dialog 1–1075
preserve modifier 1–766
preset lights 2–1302
preset rendering options 3–23
preset views 1–24
presets 3–23, 3–437
brush 3–690
configure (video post) 3–327
rendering 3–23
preview
animated material previews 2–1450
make 2–1434
play 2–1434
renderings 3–168
save 2–1434
preview and animation rollout 2–806
preview window 2–815
previewing
biped motion 2–929
motion 2–1084
Shockwave 3D files 3–585
W3D files 3–585
previous frame 3–723
previous key 3–723
previous link/next link 2–1114, 2–1147
previous transition
transition 2–1051
primitives
creating with keyboard 1–169
extended 1–186
standard 1–170
print size wizard 3–25
priority of actions (particle flow) 2–124

priority rollout 2–1235
prism 1–205
prismatic constraint 2–754
PRJ files 3–531, 3–902
problems 3–883
problems caused by unit settings 3–891
ProBoolean 1–378
procedural maps
  dent 2–1667
  glossary 3–997
  wood 2–1684
procedures 2–1264
process options rollout 3–124
processing panel
  mental ray 3–86
processing parameters rollout (radiosity) 3–64
ProCutter 1–388
production render 3–17
productivity 2–833
profile 2–1141
profile view 2–1114
program window 1–9
progressive morphing 1–737
project file format 3–531
project folder 3–393
project mapping projector 1–777
project workflow in 3ds Max 1–1
projected window 1–260
projection - preferences 3–821
projection holder modifier 1–778
projection modifier 1–769, 3–150
  cage rollout 1–773
  project mapping rollout 1–777
  projection rollout 1–776
  reference geometry rollout 1–772
  selection check rollout 1–775
  selection rollout 1–771
projection options dialog 3–165
projector light 2–1341, 3–998
projectors
  project mapping 1–777
projects - managing 3–385
prompt line 3–699
prop bone 2–1065, 3–998
propagate materials to instances 2–1432
propagation 2–622
  blocks 2–1432
  instances 2–1432
  materials 2–1432
  styles 2–1432
properties
  animation controllers 2–519
  changing layer properties 3–666
  controller (Track View) 2–560
  dialog (Track View) 2–519
  file menu 3–500
  of light 2–1276
  particle system 2–134
  rigid body 2–717
  viewports 3–731
  waveform controllers 2–519
properties (clips) 2–1027
  motion flow 2–1045
properties rollout 2–815
props 2–898
  using 2–898
protractor helper object 2–26
proximity test (particle flow) 2–218
ProxSensor (VRML97 helpers) 3–598
proxy object
  XRef object 3–414
proxy object rollout
  XRef object 3–414
PRS
  PRS controller (Track View) 2–357
  PRS parameters 2–303
PS files 3–612
PSD file (glossary) 3–629
pseudo alpha compositor (video post) 3–382
pseudo alpha filter (video post) 3–346
pseudo color exposure control 3–300
publish Shockwave 3–580
publishing to
  3D DWF 3–555
pull 2–1147
pull bias 2–1147
pull/pinch/stretch options (tendons) 2–1147
push
  modifier 1–779
  space warp 2–59
put material to scene 2–1440
put to library 2–1443, 2–1455
pyramid 1–182

Q

QOP files 3–801, 3–804
quad menu
  Schematic View 3–653
quad menus 1–1082, 2–534, 3–694, 3–696, 3–795, 3–801
  animation 3–697
  hair and fur modifier 1–532
  reactor 2–707
quad meshing 1–392
quad patch 1–994
quadruped
  animating a 2–907
quads panel (customize UI) 3–795

quadtree 3–999
quaternion rotation 2–916, 2–948
quaternion/tcb 2–916
quaternions 2–318
queue
  video post 3–312
queue manager 3–999
queue monitor
  client (glossary) 3–999
quick align 1–465
quick render
  ActiveShade 3–17
  flyout 3–17
  production 3–17
quick start (batch rendering) 3–201
quickslice 1–676, 1–1058
QuickTime movies 3–621

R

radial dialogs
  density 3–256
  falloff 3–257
  size 3–259
radial scale 2–1130, 2–1136
  parameters (links) 2–1091
radiance file format 3–613
radiance map 3–613
radiance picture files 3–628
radiosity 2–1279, 3–51
  and animation 3–60
  and architectural materials 2–1540
  choosing radiosity 3–44
  controls 3–61
  how it works 3–56
  light painting rollout 3–70
  lighting analysis 3–76
  lighting analysis dialog 3–76
  meshing parameters rollout 3–67
  preferences settings 3–836
  processing parameters rollout 3–64
  rendering parameters rollout 3–71
  skylight 2–1296
  statistics rollout 3–75
  workflows 3–57
radiosity solution 3–51, 3–56
radius 2–1147
rag doll constraint 2–737
railing 1–210, 1–217
RAM player 3–635, 3–637
random motion
  create 2–1035, 2–1055
  motion flow 2–1035
  script 2–1035, 2–1055
  transition 2–1035
random placement difficulty dialog 2–1195
randomize keys 2–533
randomize keys utility (Track View) 2–562
range bar (video post) 3–327
ranges
  editing 2–573
  positioning 2–574
  realigning with keys 2–574
  recoupling 2–574
ranges toolbar
  Track View - dope sheet 2–541
ray
  render effect 3–234
ray tracing 3–80, 3–116
ray-traced
  reflections and refractions 3–88
  shadows 3–89, 3–114, 3–1011
ray-traced shadows
  glossary 3–1000
  parameters 2–1363
  ray-trace bias (glossary) 3–1000
RAYHOSTS file 3–124, 3–1001
  specifying name and path 3–124
raytrace
  acceleration parameters 2–1531
  adaptive antialiaser dialogs 2–1533
  attenuation rollout 2–1706
  basic material extensions rollout 2–1707
  basic parameters rollout 2–1514
  dynamics properties rollout 2–1527
  extended parameters rollout 2–1519
  global settings 2–1528
  map 2–1698
  map and material 2–1531
  maps rollout 2–1523
  material 2–1512
  messages 2–1528
  raytracer controls rollout 2–1521
  refractive material extensions rollout 2–1708
raytrace acceleration 3–116, 3–1000
  parameters for BSP method 3–129
  parameters for Grid method 3–129
  parameters for Large BSP method 3–129
raytrace map 3–83
raytrace material 3–80, 3–83
raytracer parameters rollout 2–1704
raytracing acceleration parameters dialog 2–1531
re-attachment 2–1113
reaction controllers 2–358, 2–361
reaction manager
  dialog 2–361
reaction manager dialog 2–361

reactor
analyze world 2–813
angular dashpot 2–732
car-wheel constraint 2–757
cloth 2–778
cloth collection 2–781
cloth modifier 2–778
collisions rollout 2–810
compound rigid bodies 2–722
constraint solver 2–736
constraints 2–724
cooperative constraints 2–735
create animation 2–806
deformable bodies 2–777
deformable constraints 2–795
deforming mesh collection 2–794
deforming meshes 2–793
display rollout 2–812
FFD soft bodies 2–786
fracture 2–770
frequently asked questions 2–821
gravity 2–808
hinge constraint 2–747
icons 2–707
introduction 2–703
linear dashpot 2–730
menu 2–706
motor 2–765
plane 2–764
point-path constraint 2–762
point-point constraint 2–750
preview 2–815
preview and animation rollout 2–806
prismatic constraint 2–754
properties rollout 2–815
quad menu 2–707
rag doll constraint 2–737
rag doll script 2–817
real-time preview 2–815
reduce keys 2–813
rigid bodies 2–716
rigid body collection 2–723
rope collection 2–792
rope modifier 2–789
ropes 2–789
scale 2–711
scripts 2–817
setup scripts 2–817
simple constraints 2–727
soft bodies 2–783
soft body collection 2–788
soft body modifier 2–784
solver 2–806
special features 2–712
spring 2–727
storing and accessing collisions 2–774
toolbar 2–706
toy car 2–766
troubleshooting 2–823
utilities 2–813
utility 2–805
utils rollout 2–813
water 2–801
water space warp 2–801
wind 2–803
world rollout 2–808
reactor helpers 2–715
reactor menu 3–681
real time (glossary) 3–1001
real-time playback 2–197, 2–1084
real-time preview 2–815
real-world map size 2–1619
real-world mapping 2–1619
reassign globally 2–1108
reassigning vertices 2–1089
rebuild dialogs
CV curve 1–1236
CV surface 1–1236
texture surface 1–1236
recorder (MacroRecorder) 3–782
recouple ranges (Track View) 2–574
_recover.max files 1–20
recovered file 1–20
rectangle 1–272
rectangular area light 2–1299
rectangular region 3–742
rectangular selection region 1–89
recursion depth 2–1528
red
deformable vertices 2–1150
line 2–1114
red/green/blue (glossary) 3–1001
redefine
script 2–1048
redo 1–36, 1–94
redraw all views 1–50
reduce keys (Track View) 2–572
reducing mesh size 2–1099
reference biped 2–974
reference coordinate system 1–435, 1–443, 3–815
reference objects 2–178, 3–1002
reference pose 2–1076
references 1–472, 2–577
glossary 3–1002
overview 1–472
using XRefs 3–393

refine
editable spline segments 1–303
editable spline vertices 1–297
glossary 3–1002
refining curves and surfaces (concept) 1–1091
reflect/refract map 2–1699, 3–83
reflectance display 2–1430
reflection maps 2–1508, 2–1695
reflections 3–88
refraction maps 2–1509, 2–1695
refractions 3–88
refresh content
communication center 3–715
refresh viewport display 1–50
regathering 3–61
region 1–65
selection 1–65, 1–80, 1–89 to 1–91
selection method (edit menu) 1–92
selection preferences 3–860
zoom region (particle view) 2–129
zoom region (viewport control) 3–742
region net render 3–197
reinitialize 3–1002
reinitialize (physique) 2–1098, 2–1106, 2–1113
reinitialize selected links 2–1091, 2–1136
relative scale 2–1150
relative snap 2–35, 2–41
relax dialog 1–1076
relax mesh 1–986
relax modifier 1–779
relax tool dialog 1–912
relaxing texture coordinates 1–912
remap locally 2–1108
remove
ease or multiplier curve (Track View) 2–585
note track (Track View) 2–553
remove from link 2–1089, 2–1150
remove from track set 2–591
remove note track 2–553
rename objects tool 1–128
rename preview 3–170
rename settings preset dialog 3–437
render
ActiveShade 3–17, 3–21
blowup 3–13
common parameters rollout 3–27
default scanline 3–38
dialogs 2–1455, 3–2, 3–9, 3–633
presets 3–23
render operator (particle flow) 2–206
render scene 3–2, 3–12
render type list (main toolbar) 3–13
rendered output 3–272
rendering elements separately 3–130
scripts rollout 3–34
selected objects 3–13
to texture 3–144
vertex coloring 2–1693
VUE files 3–130
render bounding box/selected dialog 3–16
render effects 2–512, 3–218 to 3–219, 3–223, 3–226, 3–230, 3–234, 3–238, 3–242, 3–246, 3–250, 3–260, 3–265 to 3–266, 3–268 to 3–269
hair and fur 3–220
render elements 3–130, 3–137 to 3–138
hair and fur 3–140
velocity 3–142
render scene dialog
advanced lighting panel 3–44
common panel 3–27
render panel 3–2
renderer panel 3–36
render shortcuts toolbar 3–689
render to texture 3–144, 3–146, 3–150, 3–158, 3–160, 3–162 to 3–165
normal bump mapping 3–150
render to texture dialog 3–156
add texture elements dialog 3–164
automatic mapping rollout 3–163
baked material rollout 3–162
general settings rollout 3–157
objects to bake rollout 3–158
output rollout 3–160
projection options dialog 3–165
render UVs dialog 1–914
render UVW Template 1–914
renderable spline modifier 1–781
rendered frame window 3–5
rendered output 3–173
renderer
configure preset dialog 3–33
renderer panel
render scene dialog 3–36
renderers 3–35 to 3–36, 3–38, 3–130
interactive 3–1030
panorama exporter 3–170
photorealistic 3–1006
scanline 3–1006
viewport 3–1030
rendering 3–1, 3–79
and multithreading 3–828
batch 3–201 to 3–202, 3–208
bones 1–404
command line 3–209, 3–211, 3–215
commands 3–11
don't alias against background 3–826

elements separately 3–137
email notification 3–33
large images 3–828
on multiprocessor systems 3–828
portions of scene 2–1379
preferences settings 3–826
presets 3–23
reactor toolbar 3–688
render farms 3–180
render operator (particle flow) 2–206
rendering method 3–853
scene 1–9
shapes 1–262
to texture 3–144, 3–146, 3–150
with caustics (mental ray) 3–80
with global illumination (mental ray) 3–80
with motion blur 3–89
with shadow maps 3–89
rendering algorithms rollout 3–116
rendering commands
render last 3–25
rendering effects
multi-pass (cameras) 2–1382
rendering for print 3–197
rendering menu 3–683
ActiveShade floater 3–21
ActiveShade viewport 3–21
advanced lighting 3–44
effects 3–218
environment 3–271 to 3–272
lighting analysis 3–76
Material Editor 2–1409
material/map browser 2–1412
print size wizard 3–25
radiosity 3–61
RAM player 3–635
raytrace global exclude 2–1531
raytrace settings 2–1528
render 3–12
render scene 3–12
render to texture 3–156
show last rendering 3–25
video post 3–311
rendering parameters rollout (radiosity) 3–71
rendering properties
family elements 3–457
instanced objects 3–457
renderingmenu
panoramic exporter 3–170
reparameterize dialog 1–1237
repathing 3–487
repel behavior 2–1218, 3–1003
replace (file menu) 3–470
replace clip 2–634
replace dialogs 2–1456, 3–470
replace keys 2–579
replacing scenes 1–16
requirements
for camera tracking 2–669
system (see Installation Guide) 1–xiv
rescale world units utility 2–53
reservoir 2–626, 2–649
reservoir file groups dialog 2–635
reset 3–387
reset background transform (viewport image) 1–45
reset material settings to default 2–1442
reset XForm (transform) utility 1–438, 2–432
reset position 3–442
reset tangents 1–297
reshade 3–17, 3–22
resizing arms 2–852
resolution 2–1114, 2–1141, 2–1147
glossary 3–1003
resolve externally referenced file dialog 3–439
resource collector utility 3–512
resource information dialog 3–517
respect animation range 2–550
restore
active view (views menu) 1–37
to default settings (animation controllers) 3–828
restrict to axis buttons 1–437
restrict to x 1–437
restrict to xy plane 1–437
restrict to y 1–437
restrict to yz plane 1–437
restrict to z 1–437
restrict to zx plane 1–437
restructure biped to match file 2–936, 2–1065
retarget 2–974
retargeting rollout 3–481
reverse knees (creating characters with) 2–891
reverse time (Track View) 2–569
revert to startup UI layout 3–807
reviewing and editing strokes 3–865
Revit
files 3–447
materials 3–453
objects 3–447
settings 3–450
workflow 3–452
Revit material
3ds Maxobjects 3–455
RGB (glossary) 3–1001
RGB files 3–633
RGB maps
multiply 2–1691

tint 2–1693
right-click menus 3–694
additional quad menus 3–696
animation 3–697
customize display 3–787
display option 3–821
display preference setting 3–821
event display (particle view) 2–133
HTML help viewer 3–879
material editor copy and paste 2–1418
modifier stack 3–766
morpher modifier 1–729
named selection sets 1–85
NURBS 1–1082
sample slot 2–1422
Schematic View selection 3–653
snaps 2–41
spinner 2–282
Track View 2–516, 2–534
Track View toolbar 2–603
viewports 3–731
XRef entities list (XRef object dialog) 3–405
XRef files list (XRef object dialog) 3–403
rigid bodies 2–716
basics 2–717
collection 2–723
compound 2–722
constraints 2–724
fracture 2–770
properties 2–717
storing and accessing collisions 2–774
rigid envelopes 2–834, 2–1085, 2–1091, 2–1108, 2–1111, 2–1130
rigid surfaces (NURBS) 1–1089
ring 1–809
render effect 3–230
ring array 1–415
ringwave 1–202
ripple
modifier 1–783
space warp 2–102
RLA files 3–630
RMAT materials
in 3ds Max 3–446
roll angle manipulator 2–1334
roll viewport controls
camera 3–747
light 3–753
rollout
distributed bucket rendering 3–124
twist poses 2–950
rollouts 1–12, 2–488, 2–491
inverse kinematics 2–497
maps 2–1474
materials 2–1470
paint deformation 1–1064
PArray rollouts 2–258, 2–260, 2–262, 2–268 to 2–271, 2–274
root name 2–847, 2–984
ropes 2–789
collection 2–792
modifier 2–789
rotate 1–439
rotating
curling fingers 1–446
editable mesh edges 1–1011
elbows and knees 2–891
hierarchy 1–446
lights 2–1282
links 2–891
multiple biped links 2–895
multiple linked objects 1–446
particles 2–153
spine 2–891
views 1–29
rotation
and collision rollout (PArray) 2–268
go to rotation test 2–224
increment (viewport preference) 3–821
rotation operator 2–153
rotational joints 2–486, 2–496
rotoscoping (glossary) 3–1003
roughness 2–1490
roughness mapping 2–1500
RPF files 3–631
RPS files 3–23
rubber band mode 2–852, 2–876, 2–936
rubber-band mode 3–1003
rubber-banding
arms and legs 2–852
ruled surface sub-object 1–1193
run 2–936, 2–988
run footstep 2–988, 2–995
run script 3–781
running
dynamics of 2–878
in place 2–930
parameters 2–861
RVT link 3–1004

S

safe frames 3–733, 3–857, 3–1030
safe video threshold 2–1434
sample object 2–1432
preview 2–1434
UV tiling 2–1433
video color check 2–1434

sample range (glossary) 3–1004
sample rate 3–98
sample slot 2–1420, 2–1422
  active 2–1441
  adding bitmap 2–1631
  and material name 2–1396
  background 2–1433
  backlight 2–1432
  cool 2–1442
  default 2–1442
  display adjustment 2–1436
  hot 2–1442
  right-click menu 2–1422
sample type
  cube 2–1432
  cylinder 2–1432
  sphere 2–1432
sample UV tiling 2–1433
samples
  animations 2–920
  motion files 2–920
samples threshold (motion capture) 2–655
sampling 3–1005
  filters 3–98
  sampling quality rollout 3–98
  thresholds 3–98
satellites 3–124
save
  .mfe file 2–1032
  motion flow editor 2–1032
save animation 3–476
save as dialog 2–941
save commands
  hold 1–95
  save (file menu) 3–390
  save active view (views menu) 1–37
  save as (file menu) 3–391
  save character 1–115
  save copy as (file menu) 3–392
  save custom UI scheme 3–806
  save preview (Material Editor) 2–1434
  save selected (file menu) 3–392
  save sequence (video post) 3–324
  scene/settings in buffer 1–95
save copy as 3–392
save custom UI scheme 3–806
save file 2–936
save parameters 2–1070
save physique file 2–1106
save reservoir items dialog 2–650
save segment 2–936
save talent figure structure 2–1065
save talent pose adjustment 2–1065
saved schematic views 3–653
saving
  backup on save 3–819
  BIP files 2–920, 2–941
  biped figure files 2–855
  biped step files 2–924
  compressed file 3–819
  FIG files 2–941
  files from previous versions 3–390
  material 2–1406
  materials 2–1409
  physique data 2–1098
  STP files 2–941
  thumbnail image 3–819
  UI configuration on exit 3–815
saving animation 3–472
scale 1–440, 2–990, 2–1141
  a biped 2–1113
  factor 2–1070
  in reactor 2–711
  stride 2–936, 3–1006
  tail keys 2–965
scale operator (particle flow) 2–156
scale synchronization between AutoCAD and 3ds Max 3–422
scale test (particle flow) 2–227
scale values (Track View) 2–581
scale weight 1–810
scaling 1–441
  a biped 2–984
  a node 2–1091
  and system units 2–1099
  arm 2–852
  changing 1–423
  characters 2–1100
  deformation 1–364
  face and vertex normals 1–996
  function curves 2–580
  height 2–878
  keys (Track View) 2–559, 2–580
  links 2–851
  particles 2–156
  rendering preferences 3–826
  scale XYZ controller 2–371
  time (Track View) 2–570
  values (Track View Curve Editor) 2–581
scaling objects 1–440
scanline renderer 3–38, 3–1006
scanline rendering 3–116
scatter 1–318
scatter objects dialog 2–1189
scene event (video post) 3–329
scene extents 3–1007

scene motion blur (glossary) 3–1007
scene state
  camera properties 3–518
  camera transforms 3–518
  delete 3–520
  environment 3–518
  layer assignments 3–518
  layer properties 3–518
  light properties 3–518
  light transforms 3–518
  materials 3–518
  object properties 3–518
  rename 3–520
  restore 3–520
  save 3–520
  selected parts 3–518
scenegraph 3–581
scenes 1–4
  animating 1–8
  archiving 1–19
  backing up 1–19
  importing 1–16
  managing 3–385
  merging 1–16
  rendering 1–9
  replacing 1–16
Schematic View
  delete Schematic View 3–652
  display floater 3–651
  displaying in viewport 3–731
  glossary 3–1007
  list views 3–645
  menus 3–642
  new Schematic View 3–652
  preferences dialog 3–646
  right-click menu (selection) 3–653
  saved schematic views 3–653
  Schematic View window 3–638
  selecting with 1–69
  toolbars 3–649
  using 3–640
screen 3–738
script
  and scripting definitions 3–1008
  motion flow 2–1026
  random motion 2–1035, 2–1055
  script controller (Track View) 2–372
  shared motion flow 2–1039, 2–1056
script rollout (particle flow) 2–139
script wiring 2–134, 2–175
scripted behavior 2–1220, 3–1008
scripting
  birth operator 2–145
  script operator 2–208
  script test 2–229
  script wiring (particle flow) 2–175, 2–205
scripts 2–1030
  and controlling particles 2–120
  debugging 3–783
  define script 2–1030
  path for additional 3–813
  start frame 2–1030
  start position x 2–1030
  start position y 2–1030
  start position z 2–1030
  start rotation 2–1030
scripts group
  clip mode 2–1048
  copy 2–1048
  create 2–1048
  cut 2–1048
  define script 2–1048
  delete script 2–1048
  go to frame 2–1048
  group 2–1048
  paste 2–1048
  random motion 2–1048
  redefine 2–1048
  start frame 2–1048
  start position x 2–1048
  start position y 2–1048
  start position z 2–1048
  start rotation 2–1048
  unified motion 2–1048
scripts in motion flow 3–1008
scripts rollout 3–34
scrolling panels/toolbars 1–12
SDeflector space warp 2–87
SDynaFlect space warp 2–85
searching
  defining search terms 3–876
  for files 3–510
  for help topics 3–876
  using nested expressions 3–876
section 1–282
section view 2–1114
sections 2–1141, 2–1147
seed value (glossary) 3–1008
seek behavior 2–1162, 2–1220, 3–1008
segment (glossary) 3–1009
segments 2–1114
segments shadow mode 3–114
select 1–442, 2–1150
  and rotate control points 2–1114
  and scale control points 2–1114
  and translate cross section 2–1114

by link 2–1089, 2–1150
clip/transition 2–1027, 2–1045
multiple biped links 2–895
nearest bulge angle 2–1095, 2–1114, 2–1141
select all (edit menu) 1–87
select and link button 2–422
select and manipulate 2–15
select background image dialog 1–42
select bitmap image file dialog 2–1635
select invert (edit menu) 1–88
select keys by time (Track View) 2–563
select layers dialog 3–438
select linked objects dialog 3–440
select none (edit menu) 1–88
select object (main toolbar) 1–77
select objects dialog 1–77
select similar 1–88
select time (Track View) 2–566
selection floater 1–79
select and move 1–439
select and non-uniform scale 1–441
select and rotate 1–439
select and scale 1–440
select and squash 1–442
select and transform buttons
move 1–439
non-uniform scale 1–441
rotate 1–439
scale flyout 1–440
squash 1–442
uniform scale 1–441
select and uniform scale 1–441
select behavior type dialog 2–1205
select bitmap image file dialog 2–1635
select button 2–1089
select by 1–88
color 1–88
material 2–1439
material ID 1–303, 1–308, 1–981, 1–1009, 1–1238
name 1–77, 2–589
name (button) 1–77
name (edit menu) 1–88
select similar 1–88
smoothing group 1–981, 1–1009
time (Track View) 2–563
vertex color 1–1003
select by channel modifier 1–785
select delegates dialog 2–1205
select emitter objects dialog 2–151
select keys 2–502
select keys by time 2–533
select menu (particle view) 2–128
select menu (Schematic View) 3–643
select objects dialog 1–77
select objects in current layer 3–667
select pivot 2–959
select region
crossing 1–93
lasso 1–90
paint 1–91
window 1–92
window/crossing 1–93
select scale rotate control points 2–1114
select similar 1–88
select time (Track View) 2–566
select tool (particle view) 2–128
selected deformable/rigid envelope areas 2–1130
selecting
actions, events (particle flow) 2–132
and blocks 3–460
basics 1–64
by ID 1–303, 1–308, 1–981, 1–1009
by material 2–1439
by name 1–67
by particle event 2–138
by particle ID 2–138
by region 1–65
hierarchies 2–424
mesh sub-objects 1–998
named selection sets 1–67
objects 1–61, 1–64
open editable mesh edges 1–1011
particles 2–138
selection filters 1–68
shadow type 2–1331
shape sub-objects 1–289
successive vertices 1–297
time 2–566
with Schematic View 1–69
with Track View 1–69
selecting workbench tracks 2–1011, 2–1016
selection (particle view) 2–128
selection center (use center flyout) 1–447
selection commands 1–76
selection filter (main toolbar) 1–81
selection floater (tools menu) 1–79
selection list 3–718
selection lock toggle 3–707
selection region 1–80, 1–89 to 1–91
selection rollout
edit poly modifier 1–647
editable mesh 1–999
editable poly 1–1024
hair and fur modifier 1–521
selection sets 1–67, 1–81, 1–83, 1–511
selection statistics 1–1253

self-illumination 2–1480 to 2–1482, 2–1487
self-illumination (glossary) 3–1009
self-illumination mapping 2–1502
send out test 2–230
separate tracks options 2–888, 2–980, 2–1002
separating particles 2–172
server setup and managing (network rendering) 3–182
set all 2–1070
set as skin pose 1–116
set bulge angle 2–1114, 2–1141
set current layer to selection's layer 3–667
set free key 2–956
set key 2–280, 2–904, 2–956, 2–962, 3–718
set lowest starting foot height to Z=0 2–936
set multiple keys 2–965
set parents 2–962
set planted key 2–956
set project folder 3–393
set sliding key 2–956
set start frames dialog 2–1237
set twist pose 2–950
set weight 1–810
setting keys 2–904
setting up
  directories 3–187
  rendering software 3–186
  your scenes 1–4
settings
  communication center 3–713
  Revit 3–450
setup rollout 2–1188
setup rollout (particle view) 2–136
SGI image file formats (glossary) 3–633
shade selected (views menu) 1–47
shaded viewports 1–52
shader basic parameters rollout 2–1466
shader list 2–1723
shaders 2–1466, 2–1468, 2–1480 to 2–1484, 2–1504
  car paint (mental ray) 2–1576
  custom 3ds Max 2–1711
  DirectX 2–1464
  for standard materials 3–1010
  lume 2–1713
  LumeTools 2–1713
  mental ray 2–1710 to 2–1712, 2–1723, 3–1009
  mental ray (third party) 2–1711
  mr physical sky 2–1321
  viewport 2–1464
shading and lights 2–1399
shading type 2–1397 to 2–1398
shading, cartoon 2–1605
shadow (center of mass) 2–846, 2–933
shadow maps 2–1363, 3–1010
  mental ray 2–1360
shadow parameters (lights) 2–1337
shadow types 2–1331, 3–45
shadows 3–89
  shadow maps 3–89, 3–114, 3–1011, 3–1037
  shadow modes 3–114
  shadows rollout 3–114
shadows and rendering 2–1331
shadows from hair 3–223
shadows map (baking) 3–148
shape check utility 1–265
shape commands (loft objects) 1–373
shape operators (particle flow)
  shape 2–176
  shape facing 2–176
  shape instance 2–178
  shape mark 2–183
shape sub-objects
  cloning selections 1–289
  selecting 1–289
shape-file format 3–533
ShapeMerge object 1–336
shapes 1–262, 1–282, 3–1011
  creating from edges 1–656, 1–1035
  lofting 1–262
  rendering 1–262
shared
  motion flow 2–1056
shared motion flow
  create 2–1039, 2–1056
  dialog 2–1039, 2–1056
  script 2–1039, 2–1056
sharing
  materials 2–1432
sharing a directory (network rendering) 3–188
sharing plug-ins with a network 3–814
shell material 2–1600
shell modifier 1–785
shellac material 2–1597
shift+clone
  animating 1–482
  using 1–474, 1–478 to 1–481
shifting
  center of mass 2–876
shininess 2–1537
shininess and shininess strength 2–1501 to 2–1502, 3–950
shirt (cloth) 1–573
Shockwave 3D files
  analyzing 3–585
  exporting 3–580 to 3–581
  previewing 3–585

shortcuts 3–871
Biped 2–1006
Crowd 2–1182
particle flow 2–140
Physique 2–1111
shortcuts - default keyboard
Track View 2–510
unwrap UVW 1–900
show
show curves button 3–705
show dependencies (views menu) 1–47
show end result (Material Editor) 2–1446
show end result (modifier stack) 1–503
show frame numbers (viewports preference) 3–821
show ghosting (views menu) 1–46
show home grid (views menu) 2–34
show key times (views menu) 1–46
show last rendering (rendering menu) 3–25
show map in viewport (Material Editor) 2–1445
show safe frame 3–1030
show selected key statistics (Track View) 2–595
show tangents (Track View) 2–582
show UI 3–788
show vertices as dots (viewports preference) 3–821
transform gizmo 1–45
show buffer 2–1065
to show original motion 2–1061
trajectory 2–1065
show entire trajectory 2–944
show graph
motion flow 2–1045
show icon control 2–1298 to 2–1299
show markers 2–1065, 2–1075
show prop markers 2–1075
show recognized markers 2–1075
show safe frame 3–857
show selected key statistics (Track View) 2–595
show statistics 1–1253
show tangents (Track View) 2–582
show time 2–944
show unrecognized markers 2–1075
show/hide all 2–944
SHP files 3–533, 3–1012
shrink 1–809
shutter speed 3–101
sibling
go forward 2–1447
go to 2–1447
sides 2–1141
simple additive compositor (video post) 3–383
simple constraints 2–727
simple shadow mode 3–114
simple wipe compositor (video post) 3–383
simple wipe filter (video post) 3–347
simulation, cloth 1–571
single-axis constraints 1–437
size
of particle view icons 2–137
of particle view logo 2–137
or particles 2–156
size of grid square 3–709
size test (particle flow) 2–227
skeletal deformation tool 1–791
skeletons 2–833
and physique 2–1079
structure 2–834
used with physique 2–1082
sketch tool dialog 1–916
sketch UVWs 1–898
skew modifier 1–790
skin
attaching to skeleton 2–1076
creating 2–1076
defined 2–1076
deformable 2–1076
mesh 2–1099
optimizing 2–1099
rigid 2–1076
save/load envelopes 1–802
save/load vertex weights 1–802
sliding 2–1108
sliding parameters 2–1091
valid types 2–1076
skin modifier 1–791
paint weights 1–960
weight table 1–810
skin morph modifier 1–812
skin parameters rollout (loft objects) 1–358
skin pose commands 1–116
skin pose mode 1–116
skin utilities 2–700
skin wrap modifier 1–818
skin wrap patch modifier 1–824
skirt (cloth) 1–572
sky 2–1312
skylight 2–1296
radiosity 2–1296
skylight (glossary) 3–1012
slave controller 2–313
slave parameters dialog (block controller) 2–391
sleeves (cloth) 1–573
slice
editable mesh edges 1–1011
editable mesh vertices 1–1011
modifier 1–825
slide keys 2–558

slider manipulator 2–31
slider, time/frame 3–701
sliding 2–1136
  angle 2–1070
  distance 2–1070
  footsteps 2–1064, 3–1013
  parameters (links) 2–1091
sliding and rotational joints (HI IK solver) 2–459
sliding angle 2–1070
sliding distance 2–1070
sliding door 1–251
sliding joints 2–485 to 2–486, 2–496
sliding key defaults 2–956
sliding window 1–261
smart scale 1–440
smart select 1–77, 1–80
smoke map 2–1679
smooth 2–1111
smooth modifier 1–828
smooth rotation controller 2–374
smooth twist mode 2–952
smoothing groups 3–429, 3–1013
  assigning faces to 1–1009
  assigning patches to 1–981
  viewing and changing 1–166 to 1–167
smoothing rollout 2–1238
SMPTE (glossary) 3–1013
snap frames (Track View) 2–554
snap options 2–12
snap set key 2–974
snaps
  2D/2.5D/3D 2–35
  and cuts 1–1019
  grid and snap settings 2–41
  options/settings 2–12, 2–46
  setting standard 2–11
  snap commands 2–35
  snap override 2–45
snaps toggle
  angle 2–37
  percent 2–38
  spinner 2–38
snaps toolbar 3–690
snapshot 1–453
  cloning objects over time 1–483
  dialog 1–453
snapshot dialog 1–453
snow 2–246
soft bodies 2–783
  collection 2–788
  FFD soft bodies 2–786
  modifier 2–784
soft selection 2–525
soft selection manager 2–533
soft selection rollout
  brush options 1–960
soft selection rollouts 1–963, 1–1147
software display driver 3–838, 3–840
solve rollout 2–1232
solver plane 2–472
SOmniFlect space warp 2–84
sort order (select objects dialog) 1–77
sort shadow mode 3–114
sound
  options dialog (Track View) 2–520
  sound plug-in (animation preference) 3–828
  threshold 2–386
  VRML97 helpers 3–601
source clip
  transition editor 2–1051
source file selection 2–1075
source icon (particle flow) 2–135
source outputs (particle flow) 3–987
space warp (vector field) 2–1241
space warp behavior 2–1221, 3–1013
space warps 2–55
  binding to 2–58
  bomb 2–105
  conform 2–103
  deflector 2–90
  displace 2–76
  drag 2–66
  FFD(box) 2–91
  FFD(cyl) 2–95
  glossary 3–1014
  gravity 2–73
  modifier-based 2–107
  motor 2–61
  path follow 2–71
  PBomb 2–68
  PDynaFlect 2–81
  POmniFlect 2–78
  push 2–59
  reactor water 2–801
  ripple 2–102
  SDeflector 2–87
  SDynaFlect 2–85
  SOmniFlect 2–84
  UDeflector 2–89
  UDynaFlect 2–86
  UOmniFlect 2–85
  vortex 2–63
  wave 2–100
  wind 2–75
spacing tool 1–455, 1–491
spawned particles 2–242

spawning particles 2–215, 2–230, 3–1014
special controls 1–12
special-purpose controllers 2–296
specification 3–597, 3–954
specify conversion parameters once option 2–1075
specify parameters for each file option 2–1075
specifying
  default controller values 2–294
  default controllers 2–294
  reference coordinate system 1–435
speckle map 2–1680
specular
  color 3–1014
  color mapping 2–1500
  level mapping 2–1501
  specular highlight 2–1481
specular highlights
  anisotropic 2–1492
  Blinn 2–1493
  metal 2–1494
  multi-layer 2–1495
  Oren-Nayar-Blinn 2–1493
  Phong 2–1493
specular map (baking) 3–147
speed 2–992, 2–995, 2–997
speed operators (particle flow)
  keep apart 2–172
  speed 2–159
  speed by icon 2–162
  speed by surface 2–167
speed test (particle flow) 2–233
speed vary behavior 2–1222, 3–1015
sphere
  object 1–174
  SphereGizmo helper 3–307
spherical area omni light 2–1298
spherical deflector 2–87
spherify modifier 1–829
spin operator (particle flow) 2–154
spindle 1–196
spine
  flexibility 2–846
spine link 2–984
spinner right-click menu 2–282
spinners 1–12
  spinner precision 3–815
  spinner snap 2–38, 3–815
spiral stair 1–235
splash screen 1–17
splash.bmp file 1–17
splat map 2–1681
spline dynamics 2–833, 2–856, 2–980, 3–1015
spline IK
  animating with spline IK solver 2–473
  control modifier 1–830, 2–473
  spline IK solver dialog 2–477
  spline IK solver rollouts 2–478
spline select modifier 1–831
spline-based deformation in physique 2–1083
splines 1–266, 1–842
  adding 1–303, 1–308
  arc 1–274
  attach 1–297, 1–308
  circle 1–273
  cleaning up segments 1–308
  copying (outline) 1–308
  deleting 1–308
  donut 1–276
  editable splines 1–289, 1–295, 1–297, 1–303, 1–308
  ellipse 1–274
  explode 1–308
  glossary 3–1015
  helix 1–281
  line 1–270
  making coincident 1–842
  mirror 1–308
  NGon 1–277
  rectangle 1–272
  star 1–277
  text 1–278
  used with physique 2–1082
split mesh 1–1011
split scan lines 3–197
split tests (particle flow)
  split amount 2–234
  split selected 2–235
  split source 2–236
splitting particle stream 2–234 to 2–236
spotlights
  parameters 2–1338
  spotlight distribution (photometric lights) 2–1324
spray 2–244
spring 2–727
spring back - setting (IK) 2–466
spring controller 2–375
spring dynamics object 1–400
spring options (flex modifier) 1–700
squash 1–442
squeeze modifier 1–833
SSS materials (mental ray) 2–1583
stack 1–502, 1–504, 1–508, 3–973
stack updates 2–1104, 2–1108
stack.*see* modifier stack 3–760
stairs 1–210, 1–231
  l-type 1–232

spiral 1–235
straight 1–239
u-type 1–243
standard flow operator 2–209
standard helpers 2–16
compass 2–27
dummy 2–16
expose transform 2–17
exposetm 2–17
grid 2–20
point 2–23
protractor 2–26
tape 2–24
standard lights 2–1272, 2–1288
skylight 2–1296
standard material 2–1465
standard materials
shaders (glossary) 3–1010
standard primitives 1–170
box 1–171
cone 1–172
cylinder 1–177
GeoSphere 1–176
plane 1–185
pyramid 1–182
sphere 1–174
teapot 1–183
torus 1–180
tube 1–179
standard snaps 2–11
standard user grids 2–20
star 1–277
lens effects 3–246
starfield filter (video post) 3–347
start after last footstep 2–992, 2–995, 2–997
start at current frame 2–992, 2–995, 2–997
start frame
scripts 2–1048
transition editor 2–1051
start left 2–992, 2–995, 2–997
start position x
scripts 2–1048
start position y
scripts 2–1048
start position z
scripts 2–1048
start right 2–992, 2–995, 2–997
start rotation
scripts 2–1048
starting
manager and server (network rendering) 3–182
network rendering 3–182
startup files 1–17
startup layout - return to 3–807
startup screen 1–17
startup script (glossary) 3–1015
startup scripts
path for additional 3–813
startup.ms file 1–17
state dialog 2–1207
state filters 2–965
state panel 2–1248
state transition dialog 2–1208
statistics 3–861
statistics rollout (radiosity) 3–75
status bar controls
main window 3–698, 3–701
Track View 2–588
video post 3–313
step files 2–924, 2–936
step update scripts (particle flow) 2–139
steps
and editable patch 1–986
and editable spline 1–289
stereolithography (STL) 3–586
stick 2–1070
stick figures
transition editor 2–1051
stitch tool dialog 1–918
stitch UVWs 1–897
STL
exporting files 3–588
importing files 3–586
STL check modifier 1–834
stop animation playback 3–723
STP files 2–919, 2–924, 2–1263
loading 2–942
saving 2–882, 2–941
straight stair 1–239
Strauss basic parameters rollout 2–1483
streak
render effect 3–250
strength 2–1130
stretch 2–1091, 2–1136, 2–1147
stretch bias 2–1147
stretch modifier 1–836
stretcher 1–909
strokes 3–868
defining 3–863
preferences 3–862, 3–867
reviewing and editing strokes 3–865
viewport preferences 3–821
structure rollout 2–982, 2–984
stucco map 2–1682
styles
and Architectural Desktop objects 3–461

propagation 2–1432
styling hair
hair and fur modifier 1–518, 1–526
styling rollout
hair and fur modifier 1–526
sub-materials 3–815
sub-object
chamfer curve (NURBS) 1–1161
common controls 1–1122
glossary 3–1017
material assignment 2–1424
selection 1–74, 1–506, 1–508, 1–998, 1–1084
sub-objects
Physique 2–1129
subanim controller 2–896
subdivide 1–986, 1–1011, 1–1019
subdivide modifier 1–839
subdivide modifier (world space) 1–555
subdivision displacement rollout
editable poly 1–1063
subdivision surface rollout
editable poly 1–1060
subdivision surfaces 1–701, 1–963
substeps 2–710
substitute modifier 1–840
subsurface scattering (SSS) materials (mental ray) 2–1583
subtractive opacity (glossary) 3–1017
subtree - modifying (Track View) 2–528
summary info 3–499
sun 2–1309
sunlight 1–418
sunlight (glossary) 3–1018
super black 3–826, 3–1018
super spray 2–249
superimposed material 2–1597
supersampling 2–1459, 3–1018
support period 3–1018
surf point 1–1222
surface approximation 1–1239, 1–1245 to 1–1246
surface arrive behavior 2–1223, 3–1018
surface constraint 2–396
surface follow behavior 2–1226, 3–1019
surface joints 2–485 to 2–486
surfaces
NURBS surfaces 1–1101
surface approximation (NURBS) 1–1239, 1–1245 to 1–1246, 1–1252
surface deform (SurfDeform) 1–557
surface edge curve 1–1177
surface mapper (world space) 1–556
surface modifier 1–623, 1–842
surface offset curve 1–1167
surface parameters (loft objects) 1–354
surface point 1–1222
surface properties rollout (editable objects) 1–308, 1–981, 1–1001, 1–1003, 1–1006, 1–1009
surface sub-objects - creating 1–1177
surface tools 1–623, 1–842
surface trimming 1–1080
surface-curve intersection point 1–1224
surface-surface intersection curve 1–1166
SurfDeform modifiers 1–848
SurfDeform modifiers 1–557, 1–848
swap
colors 2–1452
events (video post) 3–325
maps 2–1451
sweep modifier 1–848
extract 1–858
merge from file 1–859
pick shape 1–857
swirl map 2–1656
swivel angle 2–449, 2–472
symmetrical tracks 2–945
symmetry modifier 1–861
synchronizing animated bitmap with the scene 2–1450
synthesis control dialog 2–1179
synthesis dialog 2–1179
motion clips panel 2–1246
state panel 2–1248
synthesis panel 2–1250
synthesis panel 2–1250
synthesis/synthesize 3–1019
system paths 3–813
system unit
setup dialog 3–850
units mismatch dialog 3–852
system units and scaling 2–1099
systems 1–404
bones 1–404
daylight 1–418
ring array 1–415
sunlight 1–418

T

tags (time) 3–710
tails
adding 2–846
talent definition area 2–1070
talent figure mode 2–1061, 2–1065, 3–1019
tangent handles 1–297
tangent types 2–305
tangents
glossary 3–1019
locking 2–583
type of 2–310
tape helper object 2–24

taper modifier 1–863
targa files (glossary) 3–633, 3–997
target
and particles 2–218
camera 3–746
lights 2–1289, 2–1292, 2–1303 to 2–1305, 2–1307
target area light 2–1307
target camera 2–1371
target distance 3–90
target linear light 2–1305
target map slot 3–150
target objects - look at controller 2–344
target point light 2–1303
TCB 2–957
controllers 2–377
glossary 3–1020
TCB (biped) 3–1019
tcb rotation
controller 2–891
teapot 1–183
techniques
cloning objects 1–474
NURBS 1–1094
tee 1–287
teeter deformation 1–365
temp
path for 3–813
temporary
buffer 1–95
IGES files 3–560
tendon display options dialog 2–1128
tendons 2–1076, 2–1108, 2–1111, 2–1113, 2–1147, 3–1020
adding 2–1096
adjusting 2–1147
and fixed attach points 2–1147
attach points 2–1147
attached links 2–1147
attaching to another link 2–1096
attaching to link 2–1096
boundary conditions 2–1147
creating 2–1096
cross sections 2–1147
deleting 2–1096
inserting 2–1096
overview 2–1096
using 2–1096
workflow 2–1147
tendons sub-object 2–1128
tension/continuity/bias 2–956, 2–1135
tension/continuity/bias (glossary) 3–1020
tension/continuity/bias in biped 3–1020
terminating chains 2–471
terminators 2–437
terminology (inverse kinematics) 2–437
terrain 1–347
creating effects with noise modifier 1–744
glossary 3–1021
tessellate
and displace space warp 2–76
faces 1–1011
tessellate modifier 1–865
tessellate selection dialog 1–1077
test outputs (particle flow) 3–987
test time frames 2–141
tests (particle flow) 2–210, 3–1021
acceleration 2–233
age 2–211
circular travel 2–233
collision 2–212
collision spawn 2–215
distance from target 2–218
find target 2–218
go to rotation 2–224
scale 2–227
script 2–229
send out 2–230
size 2–227
spawn 2–230
speed 2–233
split amount 2–234
split selected 2–235
split source 2–236
time 2–211
velocity 2–233
texel 3–1021
text 1–278
texture
baked elements 3–146
baking 3–144
rendering to 3–144, 3–146, 3–150
target map slot 3–150
texture baking - shell material 2–1600
textures
and animated NURBS models 1–1099
and imported mask bitmaps 3–530
and material properties 1–1149
and NURBS models 1–1099
baking 3–156
disable texture map display 3–853
pick texture option (edit UVWs dialog) 1–888
pinning 1–878
TGA files (glossary) 3–633
thin wall refraction map 2–1703, 3–83
three-DOF limb 2–950
threshold 1–167, 1–828
adaptive control 2–1534, 2–1698

and HD IK solver 2–463
color 2–1681 to 2–1682
edge visibility 1–1006
error (camera tracker) 2–677
explode angle 1–1011
LOD 1–1253
motion capture samples 2–655
noise 2–1650, 2–1652, 2–1674, 3–282, 3–288
optimize 1–748
planar 1–719, 1–996
position/rotation (IK) 2–463
reduce keys 2–572
safe video 2–1434
sound 2–386
super black 3–826
supersampling 2–1459
use secondary (IK) 3–830
weld 1–297, 1–842, 1–888, 1–920, 1–1011, 3–542, 3–586
thumbnails
open file 3–387
viewport image 3–819
ticks (glossary) 3–1021
TIFF files 3–303, 3–634
tile/mirror (glossary) 3–1022
tiles map 2–1658
time 2–955
controlling 2–285
copy (Track View) 2–568
cutting 2–567
deleting 2–567
editing 2–566
fitting into 2–570
in particle flow 2–119
insert (Track View) 2–570
moving through 2–287
paste (Track View) 2–568
removing 2–570
rescaling active time segment 2–286
reverse (Track View) 2–569
scale (Track View) 2–570
selecting 2–566
setting time segments 2–286
specifying active time segment 2–286
time configuration button 3–725
time ruler (Track View) 2–510
time slider 2–510, 2–556, 3–701
time tags 3–710 to 3–711
time controls 3–716
time frames 2–141
time in the air 2–878
time menu, Track View 2–526
time paste (Track View) 2–568
time reverse (Track View) 2–569
time ruler (Track View) 2–510
time steps 2–709
time to next footstep 2–992, 2–995, 2–997
time warps 2–620
TimeSensor (VRML97 helpers) 3–604
timing parameters 2–988
tips
adjusting radiosity 2–1540
animation and textures (NURBS) 1–1099
camera correction 2–1393
copying keys between frames 3–703
flipping face normals 1–1010
magnifying camera adjustment 3–747
maintaining consistent camera lens size 2–1374
NURBS 1–1094
output size and rendering speed 3–28
playing animations in all viewports 3–723
propagating layer properties 3–657
radiosity and walkthroughs 3–60
selecting faces to hide 1–1001
testing radiosity 3–60
textures (NURBS) 1–1099
turning off material propagation 3–770
updating information in light lister 2–1285
using file link manager 3–419
VRML97 3–595
toes option 2–846
toggles
angle snap 2–37
animation mode 2–278, 3–717
auto key mode 2–278, 3–717
auto material propagation 2–1432
degradation override 1–34
enable ease or multiplier curve 2–585
full screen 3–738
key mode 3–724
maximize viewport 3–738
percent snap 2–38
selection lock 3–707
shortcut keys 3–872
spinner snap 2–38
window/crossing 1–93
toggling
events (particle view) 2–131
operators (particle view) 2–131
toggling dialogs 3–670
tolerance 2–1070
toolbars 3–685, 3–803
axis constraints 3–687
brush presets 3–690
controller toolbar 2–540
curve editor 2–535

displaying toolbars 3–787
dope sheet 2–538
extras 3–688
extras dope sheet toolbar 2–541
HTML help viewer 3–878
icon scheme 3–806
layers 3–688
main 3–686
ranges toolbar 2–541
reactor 2–706, 3–688
render shortcuts 3–689
Schematic View 3–649
snaps 3–690
toolbars panel (customize UI) 3–794
troubleshooting when missing 3–893
video post 3–323
toolbox (NURBS) 1–1083
tools
for low-polygon modeling 1–1252
Material Editor 2–1427
precision 2–1
rename objects 1–128
tools menu 3–674
align 1–462
align camera 1–468
align to view 1–468
array 1–450
camera match 2–1387
clone and align tool 1–459
color clipboard 1–165
display floater 3–775
floaters 3–775
grab viewport 1–35
isolate selection 1–73
light lister 2–1285
measure distance 2–15
mirror 1–448
normal align 1–465
open 1–109
place highlight 1–467
quick align 1–465
rename objects 1–128
selection floater 1–79
snapshot 1–453
spacing tool 1–455
tools rollout
hair and fur modifier 1–523
tooltips 2–133, 3–699
tooltips in viewports
preferences 3–815
toon shader 2–1605
top/bottom material 2–1599
topology (glossary) 3–1022
topology dependent modifier 3–1023
torus 1–180
torus knot 1–189
total statistics 1–1253
touch 2–936, 2–965, 3–1023
leg state 2–867
touch dynamics 2–878
TouchSensor (VRML97 helpers) 3–603
toy car 2–766
trace depth 3–106, 3–116, 3–994
track
copying 2–568
glossary 3–1024
note 2–552 to 2–553
track bar 3–703
track selection 2–945
using motion-capture filtering 2–1061
track selection in workbench 2–1016
track selection rollout 2–888
track set list 2–535, 2–538
track sets editor 2–591
track sets list 2–590
Track View 2–1000, 2–1002
assign controller 2–546
biped colored keys 2–1005
concepts 2–503
controller menu 2–521
controller toolbar 2–540
controller window 2–504, 2–512
curve editor 2–501, 2–507
curves menu 2–525
customization 2–599
delete Track View 2–598
dope sheet 2–501
editing biped keys 2–875
extras dope sheet toolbar 2–541
glossary 3–1023
hierarchy 2–512
hierarchy icons (glossary) 3–1024
hierarchy of biped objects 2–886
key time display 2–594
keys menu 2–524
keys window 2–504
menu bar 2–521
modes menu 2–521
new Track View 2–597
opening 2–886
Options menu 2–526
pan 2–595
pasting time 2–568
pick dialog 2–1252
pick dialog (block controller) 2–392
properties 2–560

ranges toolbar 2–541
selecting with 1–69
shortcuts 2–510
sound options 2–520
status bar/view controls 2–588
time menu 2–526
tracks menu 2–524
utilities 2–561 to 2–564
utilities menu 2–533
value display 2–594
view menu 2–532
working with 2–503
workspace 2–504
zoom 2–596
zoom horizontal extents 2–595
zoom region 2–597
zoom selected object 2–588
zoom value extents 2–596
Track View utilities
current value editor 2–565
tracker gizmo 2–671
trackgroup
filter 2–645
trackgroup filter dialog 2–631, 2–645
trackgroups
adjusting balance 2–622
creating and filtering 2–612
menu 2–631
tracks
adding to motion mixer 2–607
copying and pasting 2–966
menu 2–632
selecting workbench 2–1011
tracks menu, Track View 2–524
trajectories 2–957
display 2–853, 2–931
glossary 3–1025
motion panel 2–301
trajectories (biped) 3–1025
trajectory key editing 2–914
transform coordinates and coordinate center 1–442, 1–447
transform gizmo 3–1025
transform gizmos
preferences 3–832
show transform gizmo 1–45
using 1–426
transform script controller 2–379
transform tools 1–448
transform type-in 1–431, 3–709
transformation axis coordinate system list 1–443
transforms
adjusting 2–432
and envelopes 2–1086
and light objects 2–1282
and mesh sub-objects 1–998
and modifiers 1–499
animating 1–432
applying 1–423
commands 1–438
controllers (glossary) 3–909
curve 1–1157
curve sub-object 1–1157
glossary 3–1026
locking 2–433
locking axes 2–500
managers 1–433
resetting AutoCAD objects 3–442
surface 1–1182
surface sub-object 1–1182
transform tools 1–448
using 1–424
viewing and copying keys 2–283
transition 2–1028
clip 2–1048
create 2–1028
create all 2–1028
customize 2–1034
edit 2–1028, 2–1034, 2–1048
focus 2–616
motion mixer 2–616
optimize 2–616, 2–641
optimize transition 2–1058
random motion 2–1035
transition editor 2–1028
angle 2–1051
ease in 2–1051
ease out 2–1051
fixed 2–1051
frame 2–1051
length 2–1051
mixer 2–636, 2–638
optimize 2–1051
probability 2–1051
rolling 2–1051
start frame 2–1051
transition optimization dialog 2–1058
motion mixer 2–641
transition track 2–607
motion mixer 2–616
transitions
menu 2–633
state transition dialog 2–1208
translate only 3–124
translation file (specifying name and path) 3–124
translator options rollout 3–119

translucency 2–1491, 2–1538
glossary 3–1027
translucent highlights 2–1496
translucent shader 2–1484
transmittance display 2–1430
transparency 2–1537
TRC
convert into CSM 2–665, 3–577
importing 3–577
tri patch 1–995
triangle count 1–1253, 3–861
triangle pelvis 2–984
and physique 2–846
trigonometric functions 1–150
trim clips
motion mixer 2–615
trim overlapping segments 1–308
trim/extend modifier 1–866
troubleshooting 3–883
assertion failed errors 3–883
basic troubleshooting start point 3–896
Boolean objects 3–885
camera match 2–1387
camera tracker 2–685
creases or ridges in Boolean objects 3–885
Direct3D failed to initialize message 3–896
Direct3D reports a memory warning 3–896
dual monitor configuration 3–896
garment maker errors 1–622
large font problems 3–893
lost dialogs 3–893
merging corrupt files 3–883
missing command panel 3–893
missing gizmos 3–893
multiple or missing buttons on toolbars 3–893
network rendering 3–183
normal bump maps 3–151
objects disappear when the camera gets close 3–891
reactor 2–823
remember back up files 3–883
slow file opening 3–889
slow response to open or drag dialogs 3–889
slow startup time 3–889
sluggish command response 3–889
spanning across monitors 3–896
splines and Boolean operations 3–885
tips for successful Boolean operations 3–885
unit scale and movement resolution relationship 3–891
video post 3–314
viewport transparency 3–896
truck camera 3–748
truck light 3–755
true/false, setting test results (particle view) 2–132
truecolor 3–826, 3–1027
tube 1–179
turbosmooth modifier 1–868
turn to mesh modifier 1–871
turn to patch modifier 1–873
turn to poly modifier 1–874
turning on/off
actions, events (particle view) 2–133
particle system 2–136
tweens 2–276
twist 2–950, 2–1136
deformation 1–364
modifier 1–876
twist individual mode 2–952
twist links 2–855, 2–984
twist links mode 2–895, 2–952, 3–1028
twist parameters (links) 2–1091
twist poses 2–855
two-DOF limb 2–950
two-point perspective 2–1392
two-sided 3–855
type-in weights 2–1150
types of
dynamics objects 1–395
space warps 2–55
transforms 1–424

U

U and V iso curves 1–1168
U loft surface and sub-object 1–1196
u-type stair 1–243
UDeflector space warp 2–89
UDynaFlect space warp 2–86
UI 3–683, 3–785, 3–788, 3–792, 3–805 to 3–807, 3–853, 3–856
customizing 3–804
UI files 3–804
unbinding
objects 2–461, 2–491
vertices 1–297
UNC 3–1028
understanding
crowds 2–838
motion flow 2–837
motion mixer 2–604
workbench 2–837, 2–1008
understanding crowd behaviors 2–1159
undo 1–12, 1–36, 1–94
undo levels 3–815
unexpected particle spawning 2–122
unfold mapping 1–898, 1–919
unfreezing objects 1–70, 3–775
ungroup 1–106
unhide 1–53, 3–951

unhide all 2–1150
unified motion
create 2–1038
uniform scale 1–441
unify normals 1–166, 3–429
units
and display of mouse position 3–709
automatic unit conversion 3–387
file load units mismatch 3–852
mismatch 3–852
setup 3–848
synchronizing between programs 3–422
system setup 3–850
troubleshooting problems with 3–891
unit scale preference 3–815
using 2–2
universal deflector 2–89
universal naming convention (UNC) 3–1028
unlink selection 2–422
unlock
character 1–115
unlock assignments 2–1150
unlock interior edges (of selected patches) 1–968
unwrap UVW 1–878, 1–900
automatic mapping 1–898
Edit UVWs dialog 1–888
options dialog 1–920
pack UVs dialog 1–909
relax tool dialog 1–912
UOmniFlect space warp 2–85
up vector 2–1070
update 2–119
ActiveShade 3–904
background image 1–44
background while playing 3–821
during spinner drag 1–51
particle shape 2–182
scene materials 2–1457
update types (particle flow) 2–130
upper bound 2–1147
use center flyout 1–445
use pivot point center 1–446
use selection center 1–447
use transform coordinate center 1–447
use dual planes (viewport preference) 3–821
use key reduction 2–1070
use large toolbar buttons preference 3–815
use pivot point center 1–446
use pivot points 1–509
use planes (viewport preference) 3–821
use secondary threshold (IK) 3–830
use selection center 1–447
use soft select 2–525
use transform coordinate center 1–447
UseEnvironAlpha setting 3–934
user grids 2–20, 2–51
user interface
cloth modifier 1–582
customizing 3–785
garment maker modifier 1–613
hair and fur modifier 1–521
introduction 3–669
menu bar 3–672
problems and recovery 3–893
user reference 3–873
user views 1–24
user-defined object properties 1–127
using
assemblies 1–98
asset browser 1–17
auto key button 2–278
axis constraints 1–437
batch rendering 3–202
bend links mode 2–895
bipeds with crowd delegates 2–1172
clipping planes to exclude geometry 2–1379
configure paths 3–189
create panel 1–154
default joint precedence 2–468
dummy objects 2–429
grid objects 2–5
grids 2–4
groups 1–96, 3–674
help 3–873
home grid 2–4
horizon to match perspective 2–1380
HTML help viewer 3–874
IK keyframe parameters 2–900
in place mode to adjust keyframes 2–930
interparticle collision 2–243
layers 2–913
lights 2–1274
mapped materials with Particle Systems 2–240
maps to enhance a material 2–1403
materials 1–6
materials with particle array 2–239
modifier stack 1–502
modifier stack at sub-object level 1–508
modify panel 1–499
move and rotate to aim 2–1379
multi/sub-object materials with particle systems 2–242
multiple computers 3–173
named selection sets 1–67
NURBS toolbox to create sub-objects 1–1083
online reference 3–873
Schematic View 3–640

select by name 1–67
selection filters 1–68
shapes 1–262
shift+clone 1–478
spawned particles 2–242
standard view navigation 1–29
transform gizmos 1–426
transforms 1–424
transforms to aim a camera 2–1379
units 2–2
using crowds
behaviors 2–1162
crowd helper 2–1157
delegate helper 2–1157
using props 2–898
using the track sets list 2–590
utilities
animation 2–653
asset browser 3–504
assign vertex colors 2–1734
camera match 2–1387
camera tracker 2–667
channel info 2–1738
clean multimaterial 2–1742
collapse 1–966
color clipboard 1–165
create out-of-range keys (Track View) 2–562
dialog 3–779
dynamics 2–686
3ds Max file finder 3–510
filter selected euler tracks (Track View) 2–564
fix ambient 3–512
follow/bank 2–653
IFL manager 3–619
instance duplicate maps 2–1744
level of detail 1–1253
lighting data exporter 3–303
Lightscape Materials 3–574
link inheritance (selected) 2–435
list of 3–778
LOD 1–1253
MACUtilities 2–665
material xml exporter 2–1407
MAXScript 3–684, 3–780
measure 2–52
motion capture 2–655
object display culling 1–58
panorama exporter 3–170
randomize keys (Track View) 2–562
rescale world units 2–53
reset XForm (transform) 1–438
resource collector 3–512
select keys by time (Track View) 2–563
shape check 1–265
skin utilities 2–700
strokes 3–868
surface approximation 1–1245
Track View 2–561
utilities menu, Track View 2–533
utilities panel 3–778
UVW remove utility 2–1408
visual MAXScript 3–783
utils rollout 2–813
UV
coordinates 2–1405
loft surface 1–1200
sample UV tiling 2–1433
uv coordinate shader (mental ray) 2–1728
uv generator shader (mental ray) 2–1724
uv generator shader parameters rollout (mental ray) 2–1725
uv generator shaders rollout (mental ray) 2–1727
UVW
coordinates 2–1405
coordinates (glossary) 3–1028
edit UVWs dialog menu bar 1–895
map modifier 1–922
mapping in AutoCAD Architecture objects 3–447
mapping in Revit objects 3–455
remove utility 2–1408
XForm modifier 1–934
UVW mapping add modifier 1–933
UVW mapping clear modifier 1–933
UVW mapping paste modifier 1–934

V

v command-line option 3–672
-v command-line option 3–672
value display 2–594
Vault 3–487
setting working folder 3–488
VDA targa files (glossary) 3–633
vector
editable patch handle sub-object 1–979
handles and editable patch vertex sub-objects 1–986
introduction 1–151
vector handles (glossary) 3–1029
vector projected curve 1–1171
vector field 3–1029
vector field space warp 2–1241, 3–1029
create method rollout 2–1242
lattice parameters rollout 2–1242
obstacle parameters rollout 2–1242
velocity element parameters rollout 3–142
velocity interpolation 3–1030
verbosity (messages) 3–124
versioning 3–487

vertex 2–1089
operations 2–1150
settings 2–1111, 2–1113
vertex sub-object 2–1150
vertex - definition 3–1030
vertex alpha 1–938
vertex color 1–936
rendering 2–1693
vertex color map 2–1693
vertex count 1–1253, 3–861
vertex display size 3–822
vertex illumination 1–938
vertex normals (scaling) 1–996
vertex to link assignment
initialization 2–1123
vertex type 1–975
vertex weld modifier 1–935
vertex-link assignments 2–1111, 2–1113
vertexpaint modifier 1–936
adjust color dialog 1–949
brush options 1–960
color palette 1–950
paintbox 1–941
palette 1–950
vertical (move key) 2–579
vertical motion
dynamics of 2–878
vertices
adding 1–297, 1–303
align 1–1011
attach/detach 1–1011
breaking 1–1011
changing type 2–1089
checking assignments 2–1089
checking for alignment in loft objects 1–374
choosing type 2–1089
creating 1–1011
deleting 1–297, 1–1011
inserting 1–295, 1–308
make planar 1–1011
making rigid 2–1089
manually assigning deformable blended 2–1089
manually overriding assignments 2–1089
painting 1–936
reassigning manually 2–1089
removing deformable 2–1089
selecting by color 1–652, 1–1029
slice 1–1011
weld 1–1003, 1–1011
working with 2–1089
video color check 2–1434
video driver and display problems 3–896
video post 3–311, 3–315
abut 3–329
add external event 3–340
add image input event 3–332
add layer event 3–337
add loop event 3–342
add output event 3–339
add scene event 3–329
align left 3–328
align right 3–328
alpha compositor 3–381
alpha filter 3–344
animating lens effects 3–349
automatic secondary flare parameters 3–356
common procedures 3–315
composite image sequences 3–315
composite scene over image sequence 3–315
configure presets 3–327
contrast filter 3–343
create animation from still images 3–315
create starfield 3–315
cross fade compositor 3–381
edit current event 3–324
edit external event 3–340
edit image input event 3–332
edit layer event 3–337
edit loop event 3–342
edit output event 3–339
edit range bar 3–327
edit scene event 3–329
execute sequence 3–325
fade filter 3–344
flare glow parameters 3–355
flare inferno parameters 3–360
flare lens effect 3–350
flare preferences 3–353
flare ray parameters 3–358
flare ring parameters 3–355
flare star parameters 3–359
flare streak parameters 3–360
focus lens effect 3–362
glow inferno 3–368
glow lens effect 3–364
glow preferences 3–367
glow properties 3–365
highlight geometry 3–374
highlight lens effect 3–370
highlight preferences 3–376
highlight properties 3–371
image input event 3–334
image input options 3–334
join two animations 3–315
lens effects 3–349

lens effects filters 3–345
lens effects gradient colors 3–381
lens effects gradient options 3–378
lens effects gradient types 3–379
make an object glow 3–315
make same size 3–328
manual secondary flare parameters 3–357
negative filter 3–345
new sequence 3–323
open sequence 3–323
pseudo alpha compositor 3–382
pseudo alpha filter 3–346
queue 3–312
render in reverse 3–315
resize images 3–315
save sequence 3–324
simple additive compositor 3–383
simple cross fade 3–315
simple wipe compositor 3–383
simple wipe filter 3–347
starfield filter 3–347
status bar 3–313
swap events 3–325
switch views 3–315
toolbar 3–323
troubleshooting 3–314
view controls 3–313
video post compositors
alpha compositor 3–381
cross fade compositor 3–381
pseudo alpha compositor 3–382
simple additive compositor 3–383
simple wipe compositor 3–383
video post filters
alpha filter 3–344
animating lens effects 3–349
automatic secondary flare parameters 3–356
contrast filter 3–343
fade filter 3–344
flare glow parameters 3–355
flare inferno parameters 3–360
flare lens effect 3–350
flare preferences 3–353
flare ray parameters 3–358
flare ring parameters 3–355
flare star parameters 3–359
flare streak parameters 3–360
focus lens effect 3–362
glow inferno 3–368
glow lens effect 3–364
glow preferences 3–367
glow properties 3–365
highlight geometry 3–374
highlight lens effect 3–370
highlight preferences 3–376
highlight properties 3–371
lens effects 3–345, 3–349
lens effects gradient colors 3–381
lens effects gradient options 3–378
lens effects gradient types 3–379
manuarl secondary flare parameters 3–357
negative filter 3–345
pseudo alpha filter 3–346
simple wipe filter 3–347
starfield filter 3–347
video safe frame (glossary) 3–1030
view
align to view button 1–468
axonometric views 1–24
camera views 1–24
light views 1–24
navigation 1–29
perspective views 1–24
preset 1–24
view file dialog 3–502
view image file (file menu) 3–502
view preview (rendering menu) 3–170
view change 1–36
view controls
video post 3–313
view image file 3–502
view menu
Track View 2–532
view menu (Schematic View) 3–644
view samples 3–1005
view steps 1–991
view-handling commands 1–35
viewing
3D space 1–21
and changing normals 1–166
grid objects 2–6
portions of scene 2–1379
smoothing 1–167
transform keys 2–283
viewing sample biped animations 2–920
viewport clipping 2–1379, 3–731, 3–853
viewport configuration 3–853
adaptive degradation 3–859
layout 3–856
regions 3–860
rendering method 3–853
safe frames 3–857
viewport configuration dialog 3–853
viewport controls 3–735
viewport display, particles 2–202

viewport navigation
walkthrough 1–30, 3–738
viewport properties menu 3–731
viewport renderer (glossary) 3–1030
viewport rendering 3–853
viewport shaders 2–1464
lightmap 2–1614
metal bump 2–1614
viewport shading 3–853
viewports 1–22, 3–729
and display of modifier effect 3–760
arc rotate snap angle preference 3–821
background 1–38
configuring 3–853
controlling rendering 1–27
create snapshot of 1–35
DirectX manager rollout 2–1464
general concepts 1–22
grab 1–35
layout 3–856
navigating 3–735
preferences 3–821, 3–838, 3–840
prompt line 3–699
redraw all views 1–50
reset layout 3–785
right-click menu 3–731
setting layout 1–26
status line 3–701
tooltips 3–729
tooltips preference 3–815
viewport controls 3–735, 3–738, 3–745, 3–750
viewport renderer (glossary) 3–1030
views menu commands and 3–675
views menu 3–675
activate all maps 1–50
adaptive degradation toggle 3–859
add default lights to scene 1–49
create camera from view 1–48
deactivate all maps 1–50
expert mode 1–51
grids 2–33
redraw all views 1–50
reset background transform 1–45
restore active view 1–37
save active view 1–37
shade selected 1–47
show dependencies 1–47
show ghosting 1–46
show key times 1–46
show transform gizmo 1–45
undo/redo 1–36
update background image 1–44
update during spinner drag 1–51
viewport image 1–44 to 1–45
virtual viewport 3–860
visibility tracks 2–549, 2–556
visible after/before 2–979
visible/invisible 1–1006
visual MAXScript utility 3–783
VIZ files
linked geometry 3–525
VIZBlock 3–525
VIZ Render files 3–527, 3–529
Linked Geometry 3–529
VIZBlock 3–1031
VIZBlocks
selecting when file linking 3–440
volume fog environment effect 3–284
volume light environment effect 3–288
volume select modifier 1–952
volume shading
mental ray 3–95
vortex space warp 2–63
voxel 2–1531
voxel size 3–129
VPX files 3–1031
VRML format 3–591
VRML97 3–595
exporting to 3–591, 3–594
helpers 3–597 to 3–608
specification 3–597
VST targa files (glossary) 3–633
VUE file
glossary 3–1031
renderer 3–130

W

W3D files
analyzing 3–585
exporting 3–580 to 3–581
previewing 3–585
walk 2–936, 2–988
footstep 2–988, 2–992
walking gait 3–1033
walking parameters 2–861
walkthrough button 1–30
walkthrough flyout 3–738
walkthrough navigation 1–30, 3–738
wall 1–210, 1–223
editing wall objects 1–228
wall behaviors
repel 2–1227, 3–1033
seek 2–1229, 3–1033
wall seek behavior 2–1162
wander behavior 2–1231, 3–1033
warning messages 3–819

water 2–801
rendering 2–803
space warp 2–801
wave
wave modifier 1–957
waveform controller 2–381
wave space warp 2–100
wavefront files (obj, mtl) 3–588
Wavefront material files 3–590
Wavefront object files 3–589
waves map 2–1683
web distribution 2–1325 to 2–1326
web distribution (photometric lights) 2–1355
web parameters rollout 2–1355
web site
getting content from 3–504
weight 2–1114, 2–1141, 2–1150
assignments (Skin modifier) 1–807
assignments (vertex) 2–1092
vertices (Skin modifier) 1–791
weight table 1–810
weight tool dialog 1–807
weighted vertices 2–834, 2–1130
weld 1–935
editable mesh edges 1–1011
editable mesh vertices 1–1003, 1–1011
threshold 1–297, 1–842, 1–888, 1–920, 1–1011, 3–542, 3–586
vertices 1–297
weld vertices/edges dialog 1–1077
what you should know to use character studio 2–832
white paper, swivel angle and HI IK solver 2–449
wide flange 1–288
width 2–990
wind 2–803
wind space warp 2–75
window/crossing toggle 1–93
windows 1–210, 1–253
3ds Max 1–9
awning 1–256
casement 1–257
fixed 1–258
pivoted 1–259
projected 1–260
sample preview 2–1420
sliding 1–261
wire editor 2–412
wire parameters 2–411 to 2–412
expression techniques 1–146
wireframe color 3–757
wireframes 1–46, 1–52, 3–1034
wiring
particle view 2–134
wiring (particle flow) 2–134, 3–1033
wiring parameters 2–411
wiring tests to events (particle view) 2–131
wood map 2–1684
workbench 2–837, 2–1008
analyze panel 2–1017
analyzing curves 2–1011
animation 2–1012
curve view 3–925
filters panel 2–1023
fix panel 2–1020
fixing curves 2–1012
navigating 2–1010
select panel 2–1016
workflow 3–1035
and biped 2–843
applying physique 2–1083
creating bulges 2–1094
footstep animation 2–856
in character studio 2–839
motion capture 2–1061
motion flow 2–1043
procedures in this reference 2–1264
Revit 3–452
tendons 2–1147
workflows
designing materials 2–1395
edit poly 1–643
editable poly 1–1022
project workflow in 3ds Max 1–1
radiosity 3–57
set key 3–719
sub-object selection (NURBS) 1–1084
working folder 3–487 to 3–488
working with
crowd animation 2–1154
workbench 2–1008
working with AutoCAD, AutoCAD Architecture and Revit files 3–440
working with biped 2–843
working with crowd animation 2–1154
working with drawing files 3–417
world axis 1–424
world coordinate system (glossary) 3–1035
world rollout 2–808
world space 2–959, 3–1036
world space (biped) 3–1036
world-space modifier 1–512
camera map 1–513
displace mesh 1–514
displace NURBS 1–515
glossary 3–1036
LS colors 1–550

MapScaler 1–551
PatchDeform 1–552
PathDeform 1–552
subdivide 1–555
surface mapper 1–556
SurfDeform 1–557
world-space tripod 3–729
wrap cursor near spinner 3–815
wrectangle 1–284
WSM modifier 1–512, 1–550
camera map 1–513
displace mesh 1–514
displace NURBS 1–515
MapScaler 1–551
PatchDeform 1–552
PathDeform 1–552
subdivide 1–555
surface mapper 1–556
SurfDeform 1–557

X

XAF files 3–472
adding to motion mixer 2–609
adjust time in motion mixer 2–615
adjusting in motion mixer 2–611
combining with mixer 2–604
filtering in motion mixer 2–612
transitions in motion mixer 2–616
XForm modifier 1–959
XLI files 3–560
XLO files 3–560, 3–562
XMM files 3–472
xref 2–918
add offset 2–383
xref biped 2–918
xref controller 2–298, 2–383
xref ik chain 2–436
xref material 2–1616
XRef object
proxy object 3–414
XRef objects dialog
entities list right-click menu 3–405
files list right-click menu 3–403
XRef objects dialog 3–397
xref scenes
overlays 3–408, 3–412
xref systems 1–404
Xrefs
glossary 3–1036
resolving in file linking 3–439
resolving paths 3–431
XRefs
and paths 3–411, 3–415
glossary 3–1037
scenes 3–411
user path configuration 3–812
using XRefs 1–16
XRef merge dialog 3–406
XRef objects 3–393 to 3–394, 3–397, 3–414
XRef scenes 3–393, 3–407
XYZ controllers 2–317 to 2–318, 2–344, 2–356, 2–371
XYZ coordinate shader (mental ray) 2–1730
XYZ generator shader (mental ray) 2–1729 to 2–1730
xyz position 2–958
XYZ to UVW option (UVW map modifier) 1–922

Y

YUV file (glossary) 3–635

Z

-z command-line option 3–672
z element parameters rollout 3–143
zero all 2–952
zero twist 2–952
zoom
about mouse point (preference) 3–821
no zoom (particle view) 2–129
region zoom (particle view) 2–135
zoom 3–739
zoom (particle view) 2–129, 2–135
zoom (Track View) 2–596
zoom all 3–740
zoom extents (particle view) 2–129
zoom extents all/all selected 3–737
zoom extents/extents selected 3–740
zoom horizontal extents/extents keys (Track View) 2–595
zoom region (particle view) 2–129
zoom region (Track View) 2–597
zoom selected object (Track View) 2–588
zoom value extents (Track View) 2–596
zooming views 1–29
zoom selected object option (Track View) 2–886
zoom value extents (Track View) 2–596
zoom value extents range (Track View) 2–596
ZT file 3–1037