3ds max® 6
discreet®

new features guide

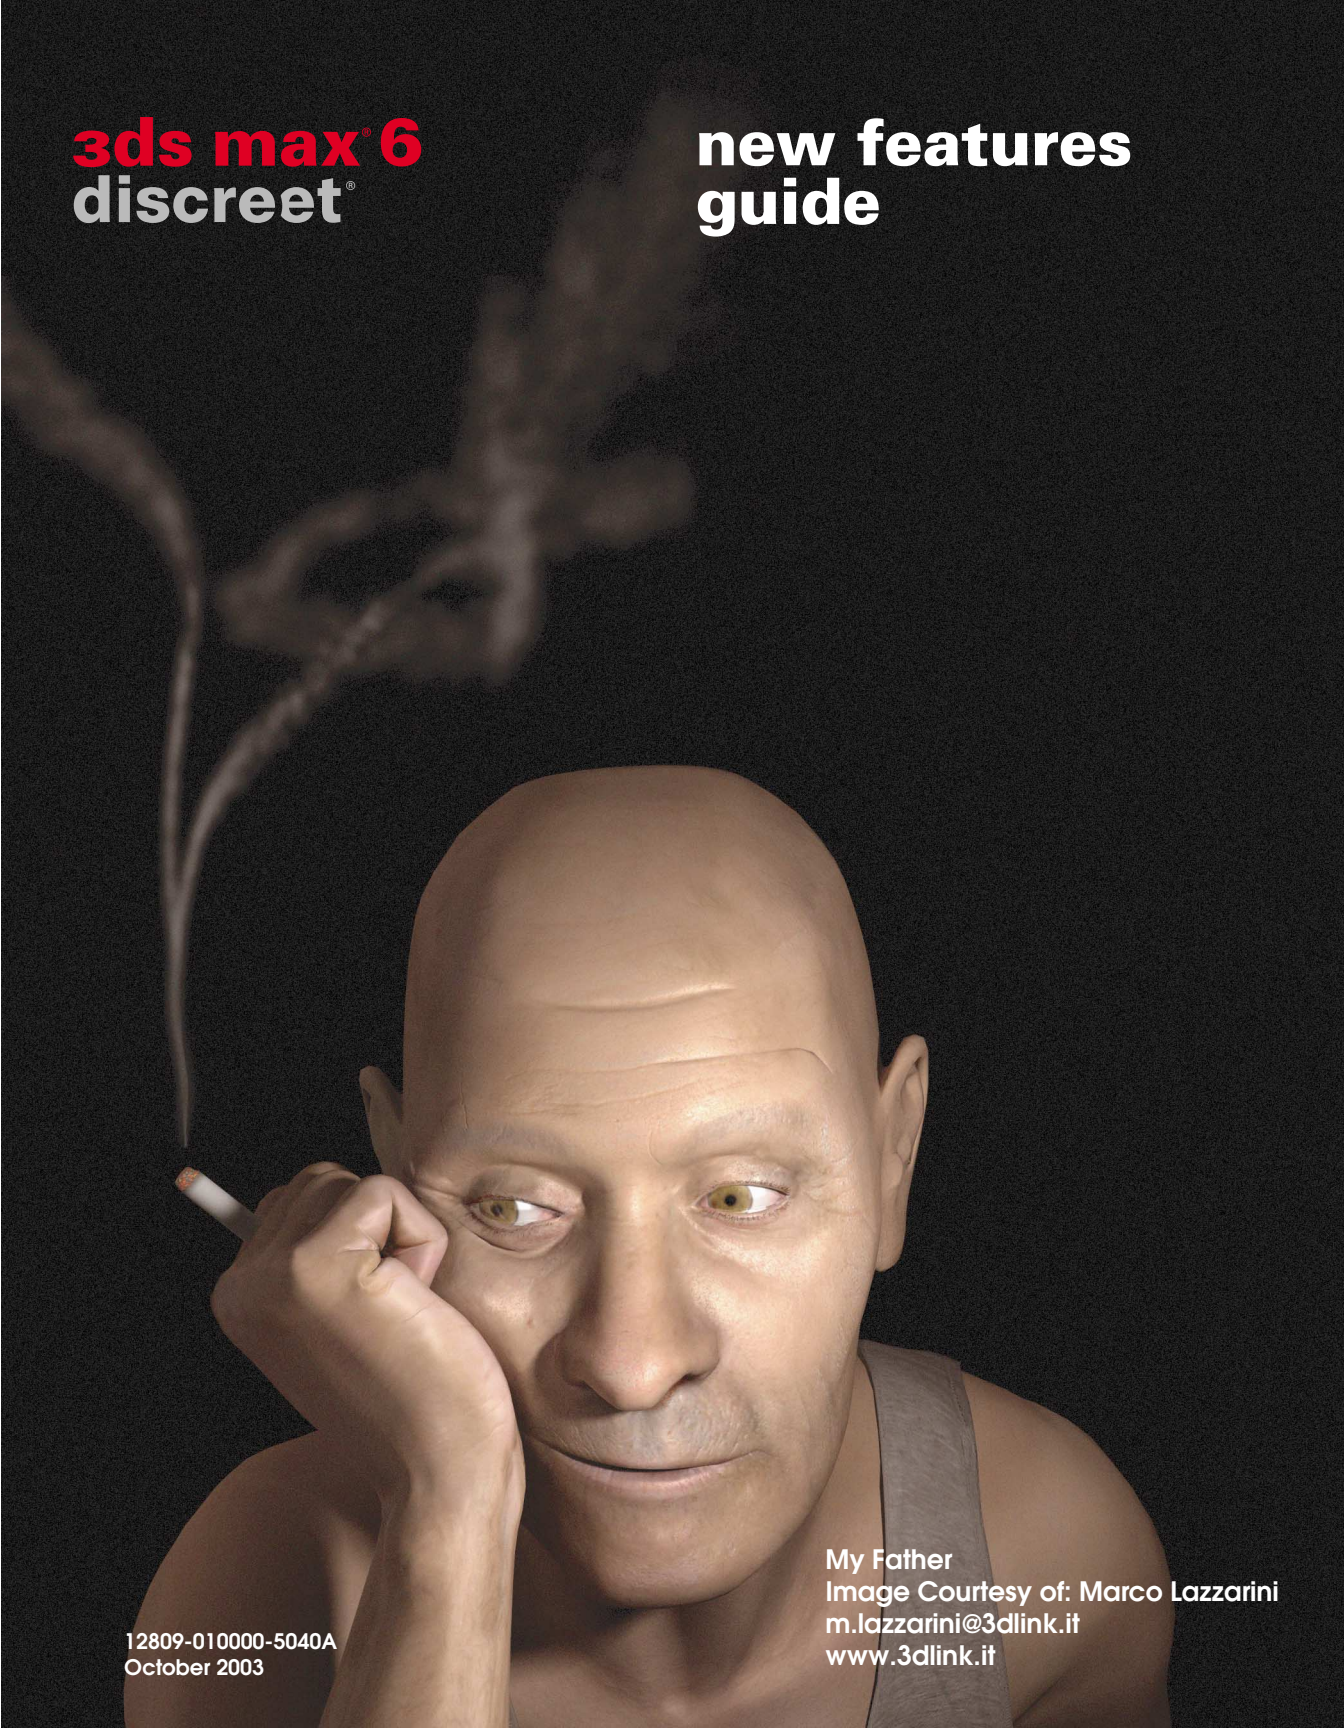

My Father
Image Courtesy of: Marco Lazzarini
m.lazzarini@3dlink.it
www.3dlink.it

12809-010000-5040A
October 2003

Introduction

The *3ds max 6 New Features Guide* is intended to guide current users of **3ds max**® through major new features added in this version of the software.

This guide attempts to explain concepts and features in a graphic and immediate way; it assumes a good, prior knowledge of the software.

Each feature is explained so that users can start using it as soon as possible. Some details and some minor features might not be present in this guide. After reading this guide, you will be able to explore the new features and start using them quickly.

This guide is not a substitute for the online reference and the tutorials. The *3ds max 6 User Reference* documents all feature details; if you have no previous knowledge of the software, you'll want to start learning the program by following the tutorials.

User Interface

1

User Interface Changes

In this first chapter, you will find a summary of the most important changes to the user interface. Some common tools will be introduced as well.

A new toolbar and a new menu have been added with all the reactor® 2 functionality. See *reactor 2* later in this guide.

On the menu bar, the Create and Modifiers menus have been updated. See more about menus in the following pages.

The Layer Manager dialog has been improved: it is now modeless, and can be started directly from the main toolbar.

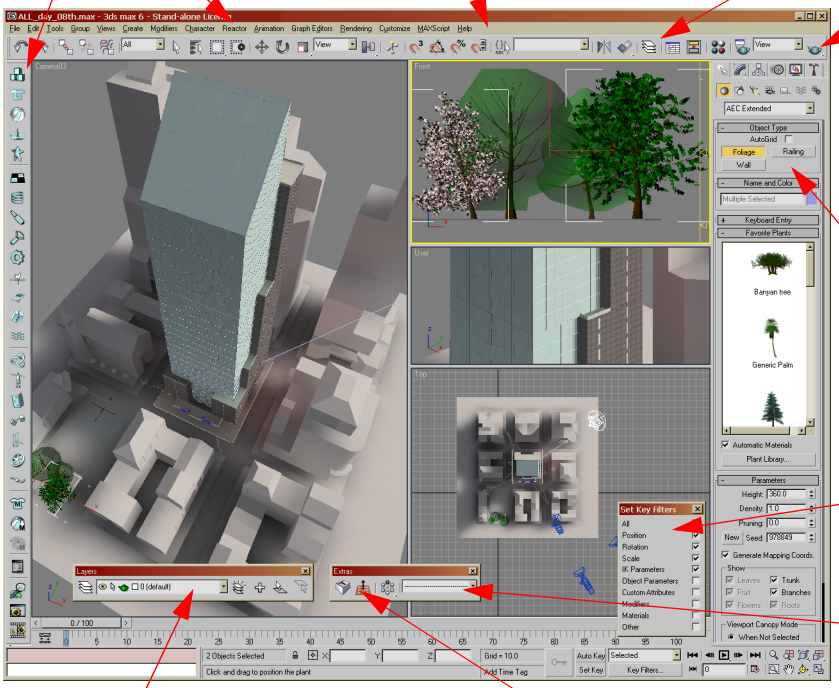

The Draft option has been removed from the rendering flyout. Render Presets replace and enhance that functionality. See the *Rendering Dialogs* topic.

Many new architectural objects and functionality have been added to provide full compatibility with other Autodesk® architectural applications.

Two new Set Key filters have been added: All and Other.

Render Presets are also available on the Extras toolbar. They let you switch settings quickly without opening the Render Scene dialog.

The Layers toolbar has been simplified because of improvements to the layers workflow and the Layer Manager. See the *Layer Manager* topic.

The new AutoGrid button in the Extras toolbar lets you enable AutoGrid while dragging i-drop® objects, like luminaires, from Web sites.

DirectX 9

Support for DirectX® 9 has been added. This enables real-time viewport preview of DirectX 9 shaders. The *.fx* shader file format is supported and FX files can be loaded directly. See the online reference for more information.

Menus

The most important changes to the main menus are listed here, as well as some of the tools they access. Be sure to follow references to other topics in this guide for additional information.

File Menu

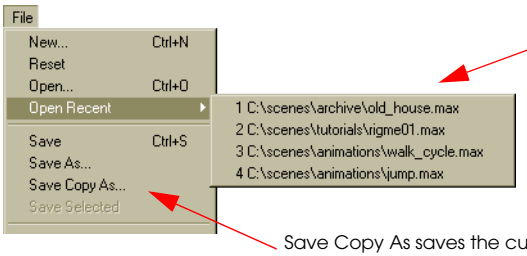

The "history list" of recently files loaded has been moved to a submenu, so more entries can be handled without confusion. You can set the maximum number of list entries in the Preferences dialog.

Save Copy As saves the current scene to a new file with a different name, but without changing the name of the current scene, This is useful for creating snapshots of the current scene.

Tools Menu

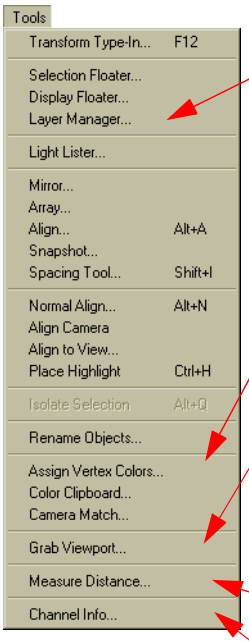

The new modeless Layer Manager dialog can be started from here, too. See the *Layer Manager* topic in this guide.

These three utilities, previously accessible only from the Utilities panel, have been added to the menu for easier access.

Grab Viewport captures the current viewport and lets you add a label. The captured image is then displayed and can be saved in any image file format.

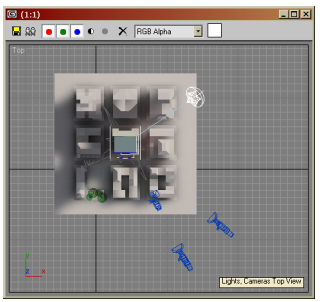

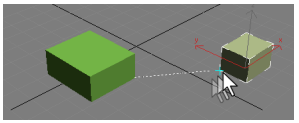

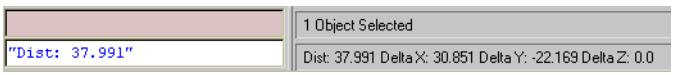

Measure Distance is a quick measurement tool that reports the distance between two points in the scene. The result is displayed in the Mini Listener, at the lower left of the main window, and in the prompt line. Use Snaps to measure a distance in 3D space.

Channel Info is a powerful new tool for channels manipulation. For more information, see *Channel Info* in this guide.

Group Menu

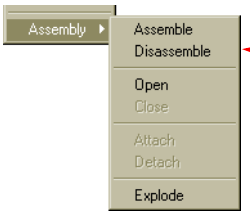

At the end of the Group menu, a new submenu has been added for Assembly objects. Assemblies are groups of objects with a special header that can hold custom attributes usable by the assembly components.

The Character Assembly is based on this technology and has its own menu.

For architectural use, we provide a Luminaire assembly. Luminaires can be used to assemble lights and geometry and to provide controls; for example, to dim all the lights in the assembly. Luminaires were introduced in Autodesk VIZ 4.

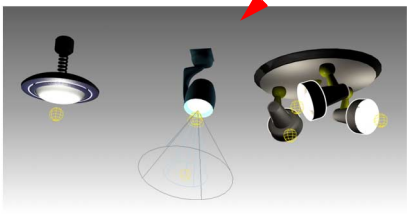

IK can be used to orient the lights, and assemblies can be inserted into a scene from Web pages using the i-drop technology.

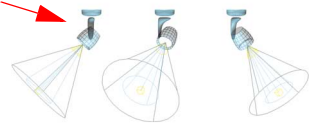

Views Menu

There's been only one change to the Views menu: Create Camera From View replaces the Match Camera To View command. This saves you the step of creating a camera first. It also automatically sets the viewport to the new camera view.

If you choose this tool with a camera already selected, this works as in previous versions of the software, repositioning the selected camera rather than creating a new one.

Graph Editors Menu

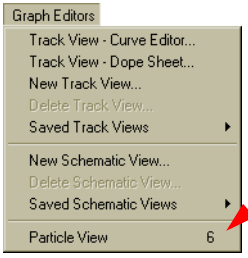

The new Particle Flow system has an editor specifically for particle behavior. You can access this editor directly from here without needing to select the particle-system gizmo and start it from the Modify panel.

For information on this new particle system, see the *Particle Flow* topic in this guide.

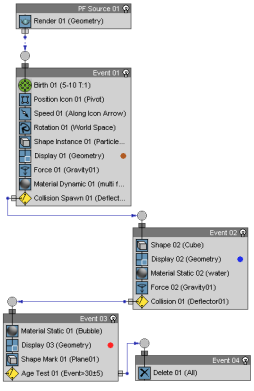

Rendering Menu

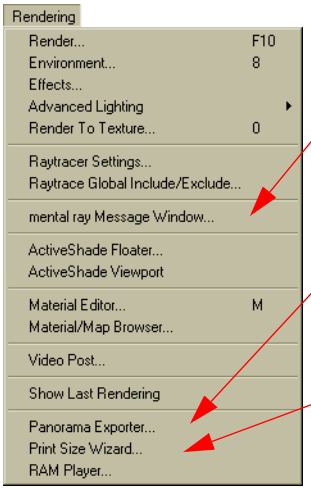

The mental ray® messages window appears automatically when error messages are issued during a mental ray rendering. By default, it doesn't open when warnings occur. To open it and check for warnings, just choose this command.

Panorama Exporter is a utility that creates 360-degree views of the current scene. For more information, see the *Panorama Exporter* topic in this guide.

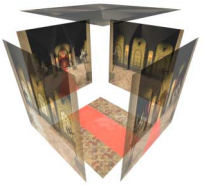

Print Size Wizard helps you calculate the size, in pixels of renderings, based on the desired resolution and print size. See the *Print Size Wizard* topic in this guide.

The Rendered Frame Window

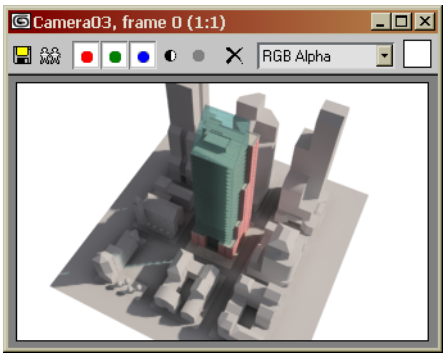

Note that the Virtual Frame Buffer has been renamed *Rendered Frame Window*.

No functionality has been changed, only the name. Be aware of this change when reading the documentation.

Other Changes

Preferences Dialog
The Advanced Lighting panel is replaced by the Radiosity panel. Some of the options have changed. A mental ray panel has been added. It contains options to enable extended features and to control the Messages Window.

Configure Paths Dialog
The Bitmaps/Photometric panel has been renamed *External Files*.

The Object Properties Dialog

The Object Properties dialog has been rearranged, with a the addition of a few new options.

The mental ray panel is not new, but in **3ds max 6**, mental ray is available to all users, by default. It contains options for setting an object's behavior with respect to global illumination and caustic photons.

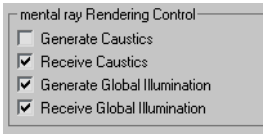

No changes to the Object Information group.

Objects can now be hidden or frozen here.

You can now display the Map Channel colors. When you chose this option, you can set the Map Channel number to one of the 99 channels available.
For more information on channels and how you can use this feature, see the *Channel Info* topic.

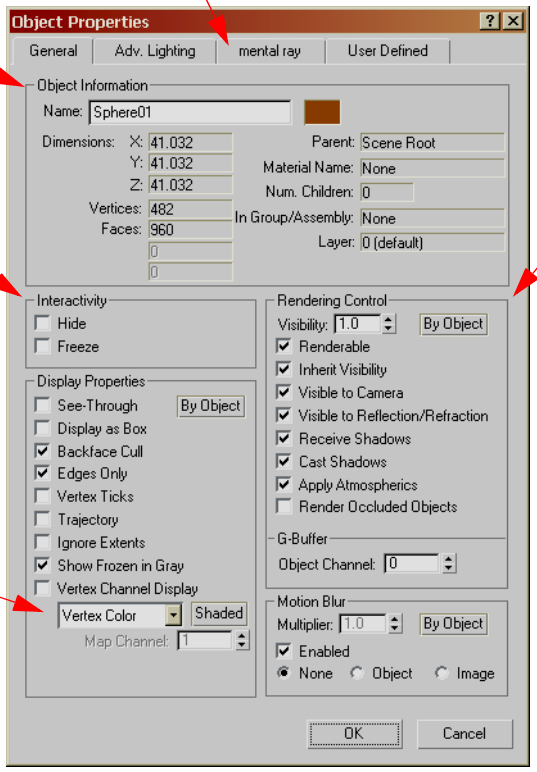

Note that the Display Properties, Rendering Control, and Motion Blur groups have been moved.

The Rendering Dialog

The new Render Scene and Environment And Effects dialogs reorganize the many dialogs needed to set all the rendering options.

The many changes are covered in the *Rendering Dialogs* topic in this guide.

Scene Management

2

Support for Autodesk Products

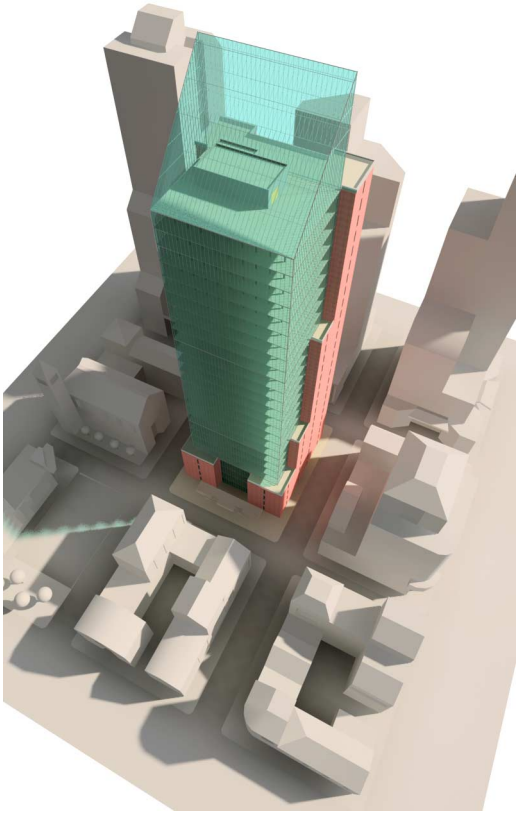

This version of **3ds max** supports most of the formats and tools previously available in Autodesk VIZ 4, plus the ability to import VIZ Render files created from Autodesk Architectural Desktop 2004.

Autodesk VIZ 4

Most of the specific objects and features of Autodesk VIZ 4 have been included, and you will be able to do everything in **3ds max 6** that was possible in Autodesk VIZ 4.

VIZ Render

Files saved from VIZ Render, the rendering add-on of Autodesk Architectural Desktop 2004, can be loaded for enhanced scene manipulation. VIZ Render files have the extension *.drf*, and a specific file type for them is available in the File Open dialog.

The image on the left was modeled in Autodesk Architectural Desktop 2004. Materials and lighting were applied in VIZ Render. Final touches and rendering were done in **3ds max**.

Importing DWG and DXF Files

The import and export of DWG and DXF files have been completely revamped. This version takes advantage of the newly written code for VIZ Render, and guarantees better integration with the new family of the AutoCAD® 2004 family of products.

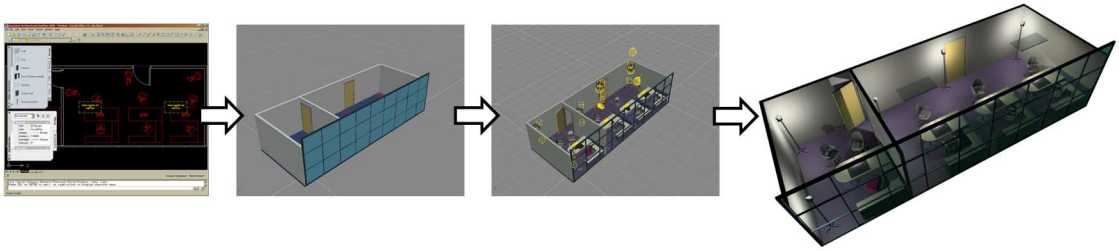

While DWG and DXF files can be imported and exported as usual, the user interface and options are different, and some elements are handled differently. For more details, refer to the online reference.

Improving Renderings

Bringing your architectural work into **3ds max**, either by loading a DRF file, or importing a DWG or DXF file, gives you the ability to use all **3ds max** features on your models.

You can add animation to all kinds of objects, and merge all kinds of entourage objects. You can use particle systems to add effects like water in a fountain. Atmospheric effects can add fog and fires. Rendering effects can add glows and lens flares.

mental ray adds a new level of realism to the existing global illumination systems, Light Tracer and Radiosity.

Be sure to check the other features related to architectural renderings and modeling in this guide: AEC Objects, Architectural material, Layer Manager, Particle Flow, HDRI files, Camera Correction modifier, and mental ray.

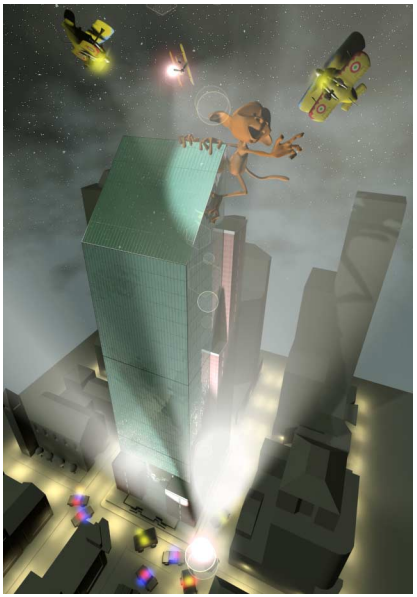

Adding effects and characters.

Material Propagation

To help manage materials assigned to blocks, especially when you have many instances of the same block, a new option has been added to the Material Editor.

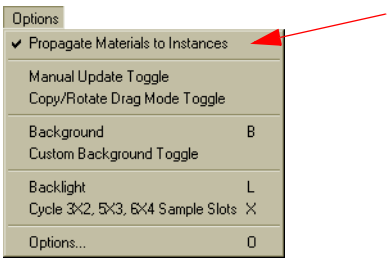

With propagation on, when a material is assigned to a block component, all the matching components in all the instances of the block will receive the material.

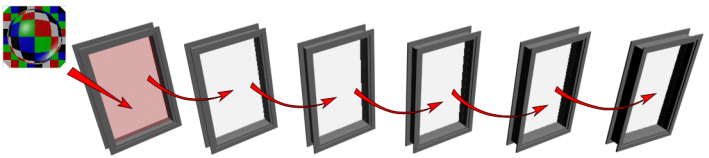

Assigning a material to the glass component of a window object propagates the assignment to all the instances of the same window.

While this was originally introduced as a VIZ Render feature for imported AutoCAD-style blocks; it now works also for standard instances of objects created in **3ds max**.

DWG Blocks

Blocks are imported and handled as instances. Changes to one block will propagate to all the others blocks of the same kind.

Block components are selectable in the viewport like any other object. To select the block header, press PAGE UP, or select it using the Object Selection dialog. Changing the position of a component changes the position of all instances. To move an individual block, you have to select the header.

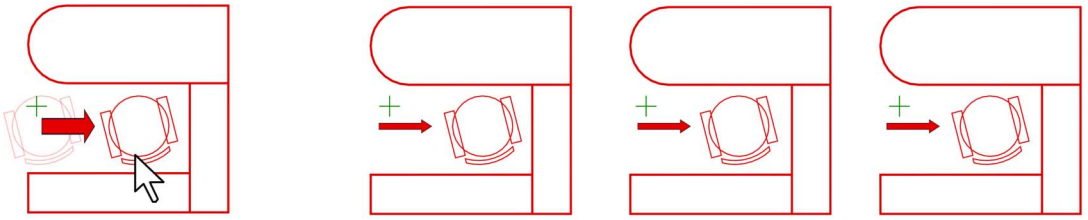

Selecting a block and moving it causes all instances of the block to move. This is because you are actually moving the object representing the block relative to the block header, which is not visible (represented by green crosses); the transform is propagated to all the instances.

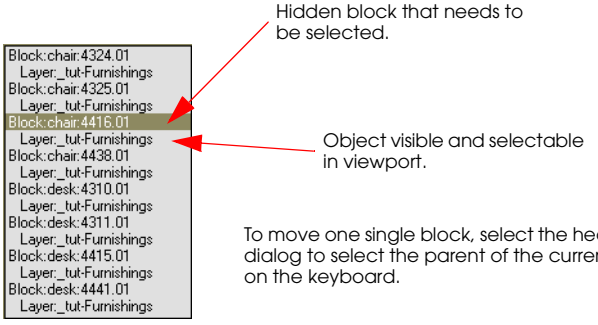

To move one single block, select the header by using the Select Object dialog to select the parent of the current object, or by pressing PgUp on the keyboard.

Lighting Data Export Utility

Found in the Utilities list on the Utilities panel, the Lighting Data Export utility exports data collected by the exposure controls. You have to have one exposure control enabled for it to work.

Exposure controls capture the whole image lighting range, and then scale it down to within the range of the output device. This utility exports the data before scaling occurs.

The output of this utility is PIC or TIF files containing the Luminance and Illuminance unclamped values (48 bits). The utility performs a rendering and records the values before exporting.

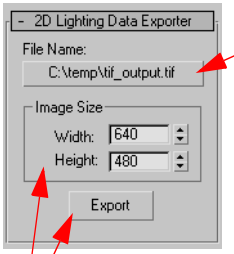

Specify the output file name before exporting. If you select a PIC file, two files will be created, one for Luminance, the other for Illuminance. If you select a TIF file, a single file with two layers is created.

Set the image size and select the desired viewport before clicking Export. The rendering will not be visible during the processing.

Note: The samples above have been adjusted from the actual output, since the high dynamic range of the image was not suitable for print or display.

Material XML Exporter Utility

This utility exports a material definition to an XML file. XML files can be shared by other **3ds max** users by dragging the file onto an object in a viewport from Windows® Explorer or a Web page.

Choose here how you want to select the material to export when the process is started.

You can choose to export as an XML file or as an Autodesk Tool Catalog for the AutoCAD 2004 Content Browser. Specify XSLT lets you add an XSL transform to the XML output.

You can choose to export the material and to add a thumbnail. If your material selection is from an object, you can also export the mapping applied.

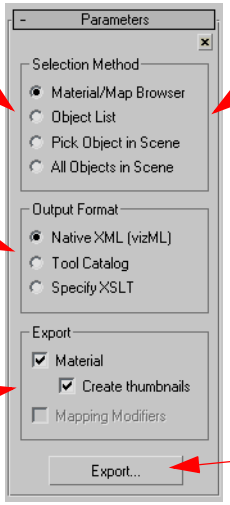

When you start the export process, you will be prompted to select the material based on your selection method. The Material/Map Browser will open, or you will have to select an object in the scene.

You may need to specify paths for maps and thumbnails during the export process.

Once all your settings are final, start the export process and follow the prompts.

Layer Manager

The way layers are managed has been completely overhauled in **3ds max 6**. The common tasks have been simplified, and the interface is now more natural. The smaller toolbar and improved workflow give you the same functionality, but we hope will make for less thinking and fewer mistakes.

The Layers Toolbar

Many functions have been removed because of the new workflow. Only the basic functions are now in this toolbar; the major layer-handling tasks are now the responsibility of the new Layer Manager window described on the next page.

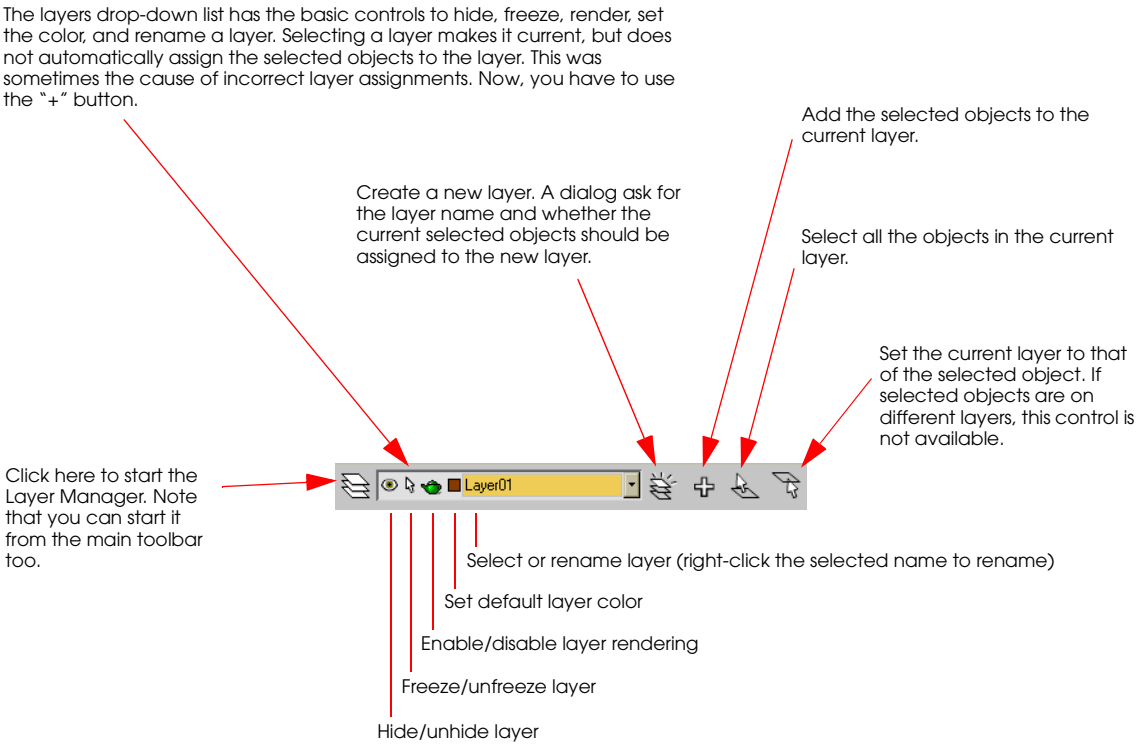

These are the tools you need for the basic layer tasks in **3ds max 6**. The toolbar and the Layer Manager do not both need to be open at the same time, as previously. Use the toolbar if you have small layers tasks to perform, or open just the Layer Manager for complete control.

The Layer Manager

This dialog replaces and enhances the old Layer Properties dialog. The Layer Manager adds many functions, and it's now modeless and accessible from the main toolbar. You can keep the dialog open, resize it, and perform all layer adjustments at any time while working on the scene.

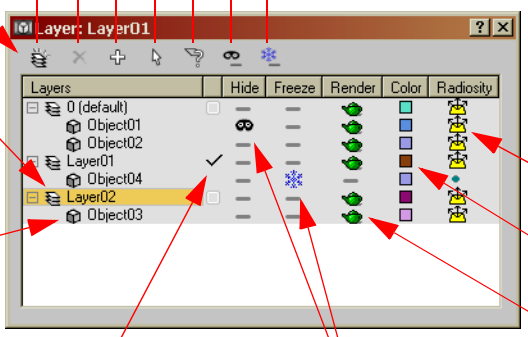

Create a new layer.

Delete the highlighted layer, if empty.

Add the selected objects to the highlighted layer.

Select all the layer objects if a layer is highlighted, or just the highlighted objects.

Highlight the layers of the selected objects.

Hide all layers.

Freeze all layers.

Tools on the Layers toolbar are replicated here. No need to have the toolbar on-screen when Layer Manager is open.

Click the column headers to sort the list as desired.

Click the layer icon to open the Layer Properties dialog.

Enable/disable radiosity.

Click the object icon to bring up the Object Properties dialog.

Set the default layer color.

Enable/disable rendering.

You can select multiple layers and objects, and use the right-click menu to cut and paste them into different layers.

These two columns hide and freeze the layers or objects. The behavior is different from the previous version, where adjusting hide or freeze settings at the object level caused the ByLayer attribute not to work. Now, objects are hidden or frozen if any of the settings for the layer or the objects are checked.

The check mark indicates the current layer, where new objects are created. Click the small squares to change current layer.

Schematic View

Schematic View has been redesigned to improve workflow and usability. Many of the past limitations are gone, and Schematic View is now an effective scene-management tool, where hierarchies, materials, modifiers, and controllers can be edited and wired.

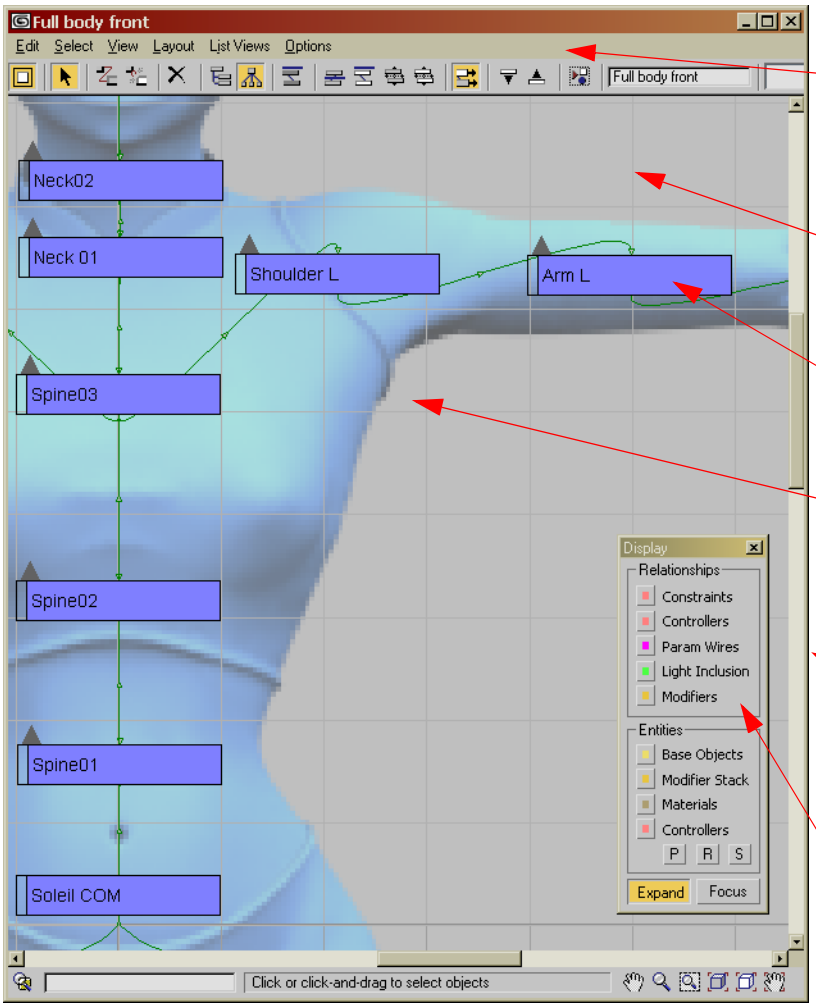

More options and more controls over layout and display from the menu and toolbar.

A visible custom grid you can snap to helps in the layout.

Easier-to-read nodes at any zoom level.

Background images can be used as reference.

Node layout is saved and maintained between sessions. You can have multiple Schematic View windows with different layouts and backgrounds.

The display floater quickly lets you filter or display various relationships and entities.

Node navigation is improved, and layouts can be saved and restored. Icons are more readable and background images can be used to lay out the nodes based on their usage within a complex hierarchy.

Schematic View is now an effective tool to analyze, diagnose, and fix existing rigs, giving you an organized way to see and modify the relationships between the nodes.

Layout Improvements

Nodes are added to the workspace starting from the origin, which can be identified by thicker lines in the new background grid.

The origin is visible as thicker lines in the grid.

All nodes for new elements in the scene are initially added to the right of the origin.

The lower-right quadrant grid is the default working area, and is a slightly darker gray.

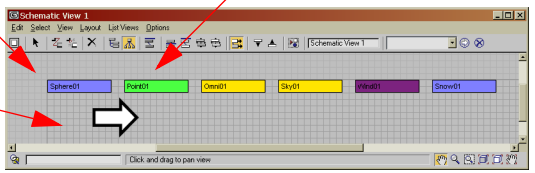

Background Images

Adding background images helps you with a visual reference to the node, rather than just the name. In complex rigs, you can have multiple views for parts of the model.

A global view of the model can be used for most of the nodes, while details like hands and feet can be hidden.

Different backgrounds and layouts can be created and saved for specific parts of the rig.

Unimportant nodes can be shrunk to clean up the display, but they remain accessible.

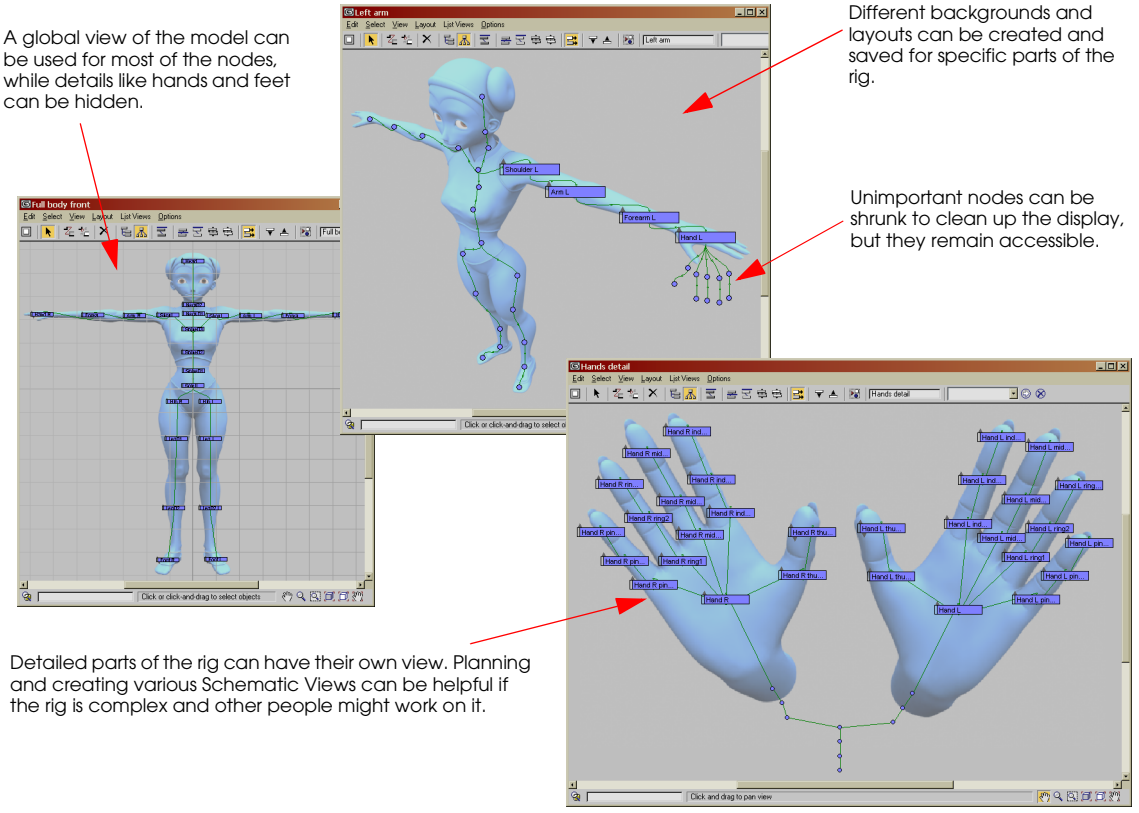

Detailed parts of the rig can have their own view. Planning and creating various Schematic Views can be helpful if the rig is complex and other people might work on it.

Bookmarks

With Bookmarks, you can record the current view of the layout, which can be recalled for easier and quicker navigation of big hierarchies.

In the Schematic View toolbar, when you want to save the current view as a Bookmark, simply enter a name and press ENTER.

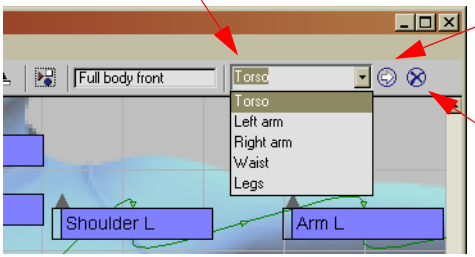

To recall a saved Bookmark, choose it from the drop-down list, then click the arrow button. The current view will be replaced with the saved one.

To delete the selected bookmark, press the delete button. The current view will not change.

Hiding and Shrinking

When some of the nodes are not necessary in a layout, you have a choice of hiding or shrinking them.

Select the nodes, then chose Hide or Shrink from the right-click quad menu, or from the Layout menu.

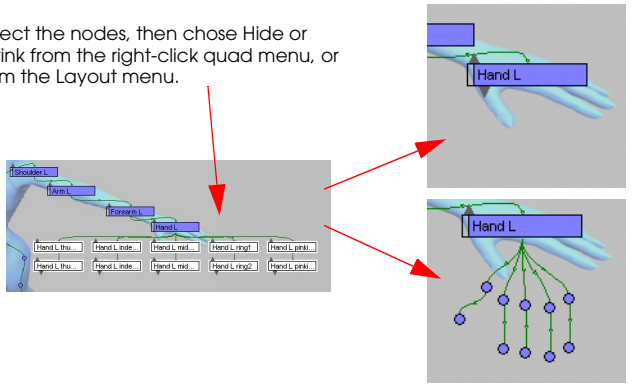

Hidden nodes will disappear from the view. Children of the selected node will be hidden, too.

Shrunken nodes appear as little dots, and while they help clean up the layout, they still remain accessible.

Free Nodes

Nodes can be flagged as Free, and will not be auto-arranged when in this state.

A free node is identified by a small empty area on the left side.

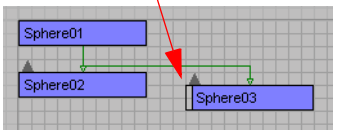

To mark a node as free, you can use the menu, the quad menu, the toolbar, or you can simply move the node manually.

Related buttons on the Schematic View toolbar.

Free Selected

Free All

Arrange Selected

Arrange Children

Always Arrange, or any of the arrange buttons here, will reset the Free flag from the nodes.

The Display Floater

The Display floater is visible by default, and can be turned on or off. You can control the display of relationships between nodes (like constraints or wired parameters), and filter some entities.

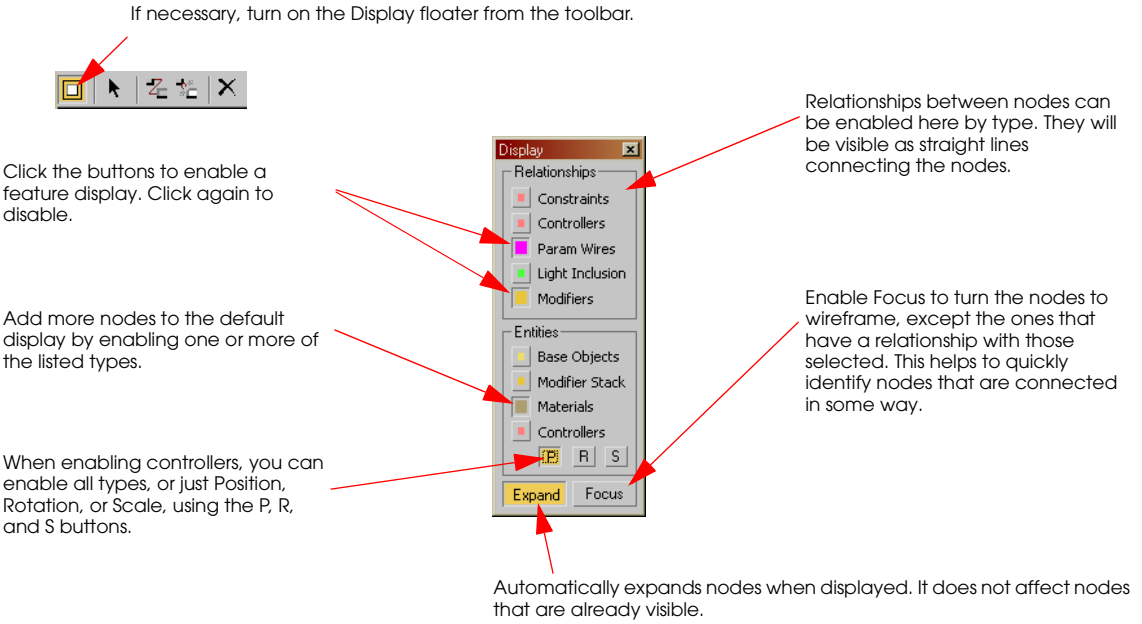

Workflow Improvements

Many things are now possible in the Schematic View that make working on scene objects much easier and more powerful.

The Link tool has been replaced by the Connect tool, which has more flexibility. When using Connect with two geometry nodes, it will link them in a hierarchy like the Link tool did, but when used on controllers, for example, it will copy or instance them.

Selecting a material in the Material Editor highlights the nodes to which it's assigned. Double-clicking a map in Schematic View opens and selects it in the Material Editor. Using the Connect tool, you can assign a material by connecting the material node to the object node.

There are a lot of little things that you will discover while working and experimenting with Schematic View. For more details on all the improvements, refer to the online reference and the tutorials.

Next, we will take a look at two other important improvements: creating and managing relationships, and list views.

Working With Relationships

Relationships are shown as straight lines connecting nodes. In a complex rig, these lines can be confusing. This is where preparing some clean Schematic Views for various parts or details of the rig comes in handy.

You can create relationships, as well as check and modify them. For example, using the quad menu, you can wire parameters between nodes in a way similar to wiring them in a viewport.

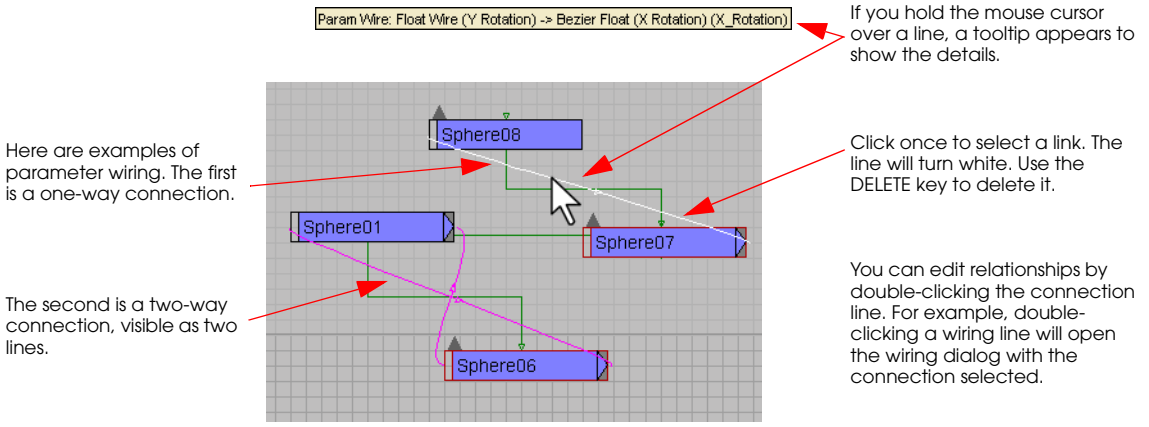

List Views

A new menu gives you access to a set of diagnostic tools called *list views*.

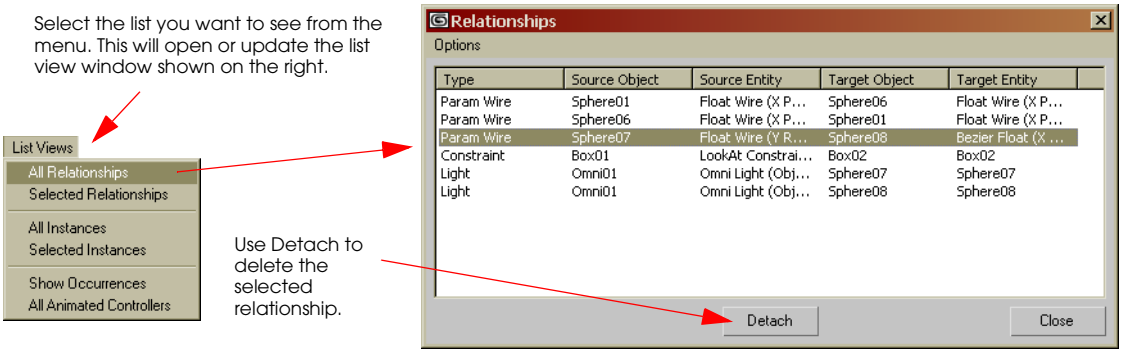

The list view window shows you details of all objects or selected objects in the scene. There are different columns and options for different views. You can always sort the list can by clicking the column headers.

For example, in the Relationships view, you can double-click an entry to open the specific relationship editing dialog. In the Instances views, the Detach button is replaced by a Make Unique button.

Use the Options menu to enable Sync Selection to the Schematic View working area, to see where the relationship is. For big hierarchies, you can also automatically zoom and pan to the selection.

Preferences Dialog

Unselected items from this list are not included in the Schematic View list of objects, and will not be accessible in the work area. Disabling some of them will gray out some items in the display floater.

Grid can be enabled or disabled. You can also set the space between grid lines, and enable snap.
Snap affects each node when moved, and will force the node to a grid crossing.

These are filters that you can enable to have only selected, visible, or animated objects displayed, in a way similar to the Track View filters.

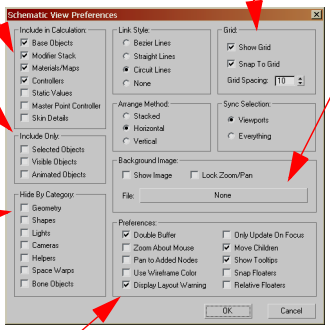

Here, you can specify a file name to use as a background image.

Select the file, then enable Show Image, and close the Preferences dialog.

In the work area, zoom and pan the nodes to set an approximate position. Make sure the node size is appropriate for the background.

When done, reopen the Preferences dialog and enable Lock Zoom/Pan.

You can now go back to the work area and start arranging the nodes.

Hide By Category hides the objects of a specific type, in a way similar to what you can do in viewports from the Display Panel.

Here are common display and navigation options. Of interest are Snap Floaters and Relative Floaters, which affect the Display floater and the list views.
Snap Floaters forces floaters to the edge or corners of the work area when dragged close to them, and keeps them there even if the window is resized.
Relative Floaters makes the floaters keep their position relative to the closest corner of the work area.
By default, floaters are positioned independently from the Schematic View window.

HDRI Files

HDRI (High Dynamic Range Image) files are now supported. They can be read like any other bitmap format, and used for environment and reflection maps, but full advantage of their extensive light range is achieved when using them as Radiance maps for skylights.

Devil statue on a snowy field

When used as a map for a skylight, the scene can be lighted by the map alone. Lights in the HDRI map will act as real lights and cast shadows.

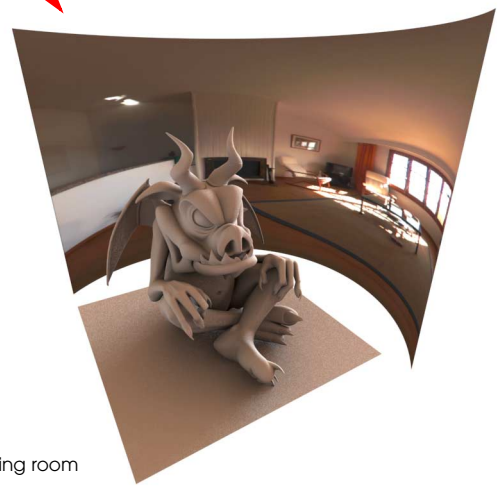

Devil statue in a living room

HDRI files are usually large files. High-resolution files might be necessary if used as environment or reflection maps. When used for skylights, the map is not visible and is used only to calculate the sample ray hits. In this case, a smaller map, sometimes even blurred if light sources in the map are small, will produce better results with a smaller impact to system memory.

HDRI files used as maps for reflections and environment need to be high-resolution because they are visible.

HDRI files used for a skylight can be smaller and blurred for better sampling and lower memory usage.

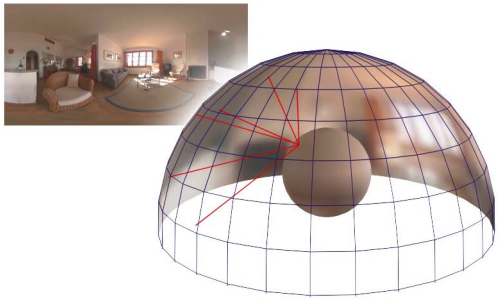

HDR images © 2003, Sachform Technology: http://www.sachform.de.

Loading HDRI Files

HDRI files have *.hdr* file-name extension. When loading, you can use the settings dialog to remap the dynamic range. This is usually necessary when the map is visible in a final rendering that is not an HDRI file. For skylights, the full range can be used instead, since it's used as a Radiance map.

Use the histogram to see the dynamic range of the image. Most images use only a narrow part of the full range.

Adjust the black and white points to limit the range. This is necessary for fine-tuning the image's appearance within the scene. To adjust the Black Point, turn on the check box. Usually, only the white point needs adjusting.

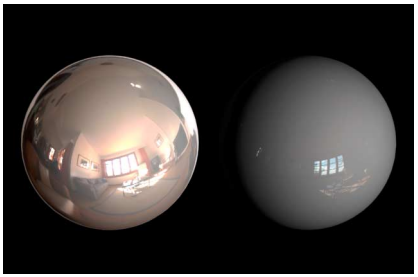

Above, the range of the HDRI reflection map has been adjusted to use only the brighter part of the lighting values.

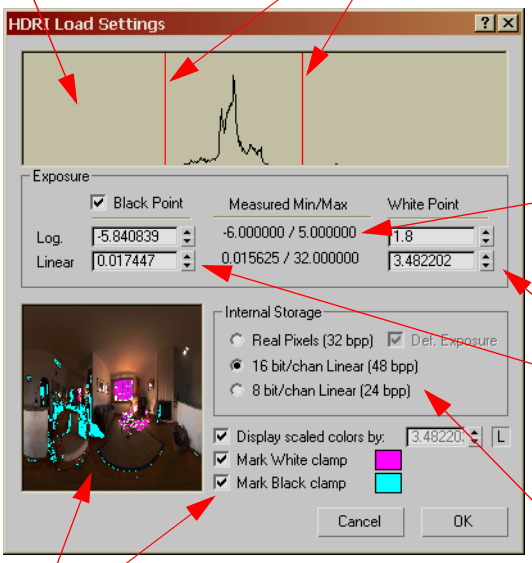

The Measured Min/Max values are the minimum and maximum range values found in the image. Use them as a reference.

Black and White points can be adjusted using a linear or a logarithmic value. Linear moves the points at constant speed. Logarithmic is very fast in the low ranges, and slow in the high ranges.
These values can be adjusted only when turning off Def. Exposure or using 16- or 8-bit-per-channel clamping.

These settings let you see the clamped areas, below the Black point (darkest color) and above the White point (brightest color), and adjust the thumbnail display. This is to aid you in fine-tuning the range for the specific image you are loading.

When using a clamped 16- or 8-bit-per-channel image, adjusting the Black and White points in the range histogram is necessary to get the best range.

Rendered images can also be saved as HDRI files. When doing so, you have the option to save the full range (Real Pixels, HDR) or a 16-bit-per-channel clamped image (Normal Pixels).

Shockwave 3D Export

Shockwave® 3D files can be exported, maintaining textures, lights, and animation. The exporter can be accessed from the File -> Export menu by selecting Shockwave 3D Scene Export (*.W3D).

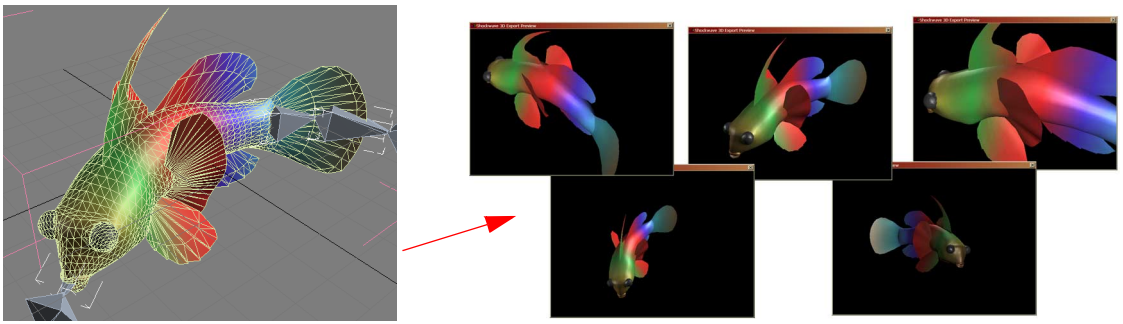

After selecting the output file name, you will be see a dialog with export options, where you can also preview the result as played by Shockwave 3D.

The Export Options Dialog

Deselect the elements you don't want to export. Textures can be demanding on memory resources. Animation can be disabled if all you want is the static object.

Select the viewport or camera you want as the initial point of view. Remember that if you select an orthographic camera, the exported scene will also be orthographic.

Use Author Check to check for problems that your scene may have on export. Click Analyze to review memory- and resource-use statistics.

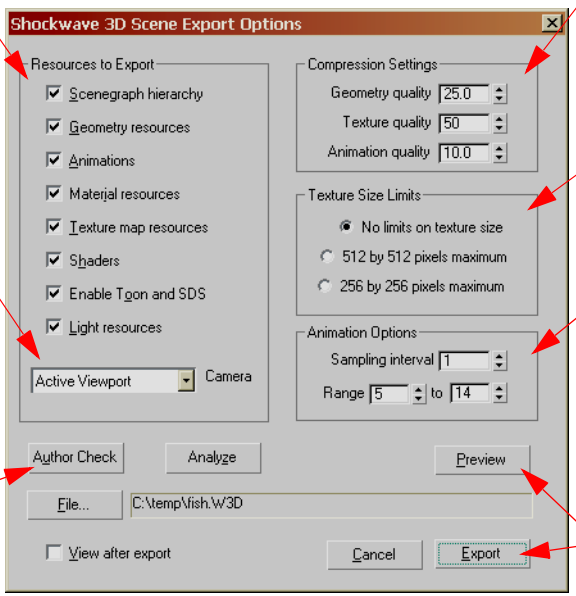

Optimize the quality of geometry, textures, and animation if you don't need all the details of the original scene.

To save memory, you can force the textures to be resized to smaller ones.

Limiting the animation sampling is another way to save memory on objects with a lot of animated parts.
Limit the range to create correct loops during playback.

Click Preview to check the result, and Export to create the final W3D file.

The Analysis Dialog

This dialog shows you some scene statistics that might give you hints on what to optimize.

Each type of resource is shown as KB used and as a percentage of the scene size. You can use these values to optimize your scene.

This value is the runtime memory used by the uncompressed textures.

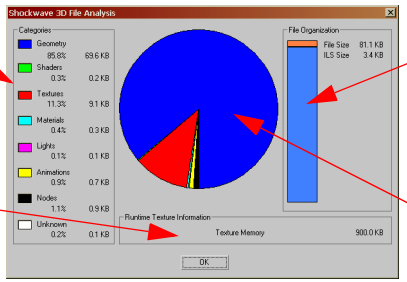

This bar chart shows the size of the Initial Load Segment (ILS), which is the amount of data Shockwave 3D must load before it can display the scene.

This pie chart shows the relative size of all the scene's resources.

Previewing the Scene

When running a preview, you can see the Shockwave 3D playback and check the result of the export. This preview might have elements that have been disabled in the Options dialog, and might be different from the final W3D export file.

If your scene is animated, the animation will loop continuously.

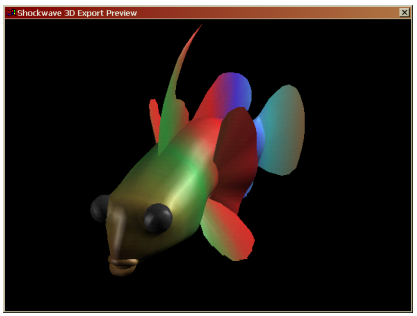

To rotate the scene, click and drag the mouse in the viewing area.

To dolly, use CTRL when you click the mouse, and scroll up and down to move closer or farther away.

To pan, press and hold SPACEBAR, then click and drag the mouse. Use SHIFT, as well, if you want to constrain the pan to the horizontal or vertical.

Restrictions

The conversion between **3ds max** and Shockwave 3D has some restrictions due to their differing technologies. You need to check some things in order to get the correct export. For example, characters using bones and skin need to be in a group, and materials have to use the Blinn shader to make sure they look the same in both programs.

You will find more information about conversion and other details in the online reference and the tutorials.

Particle Flow

Particle Flow is a new event-driven particle system that allows extreme customization of the particles' behavior and is able to create very complicated simulations. This flexibility is achieved using a schematic view of the particle behavior, where rules, tests, and events can be created and edited to visually describe each particle's behavior during its life.

Basically, it's like writing a little program that will govern the life existence of particles, by influencing their motion, changing their attributes and testing the particles' interaction with other objects in the scene, deciding each time a on new state and a new behavior for the particle.

The life of a particle

To better understand how it works, you should look at it from the particle's point of view.

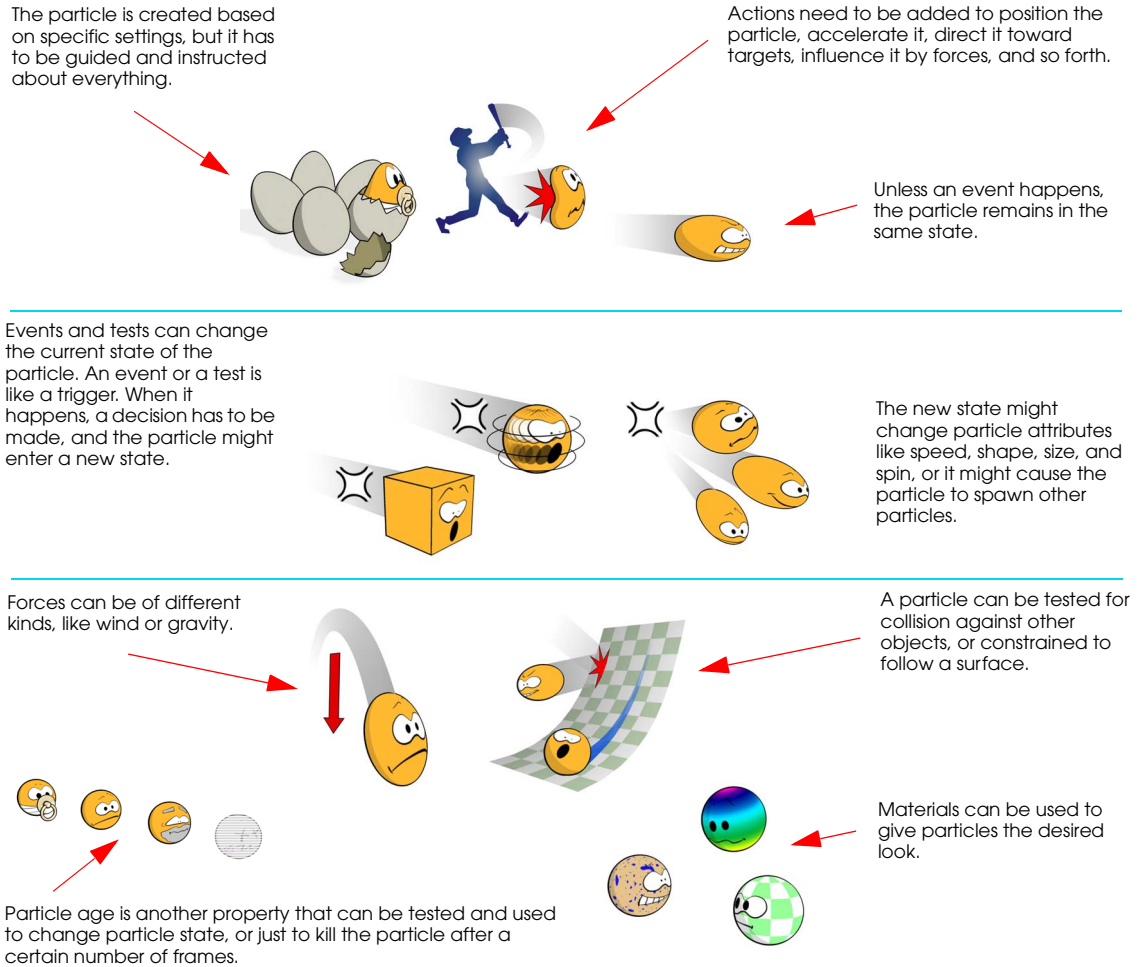

Terminology and Concepts

It's easier to understand how things work by defining a few terms and concepts. A particle system is an organized world of settings that describe particle behavior, and is logically organized in a hierarchical way.

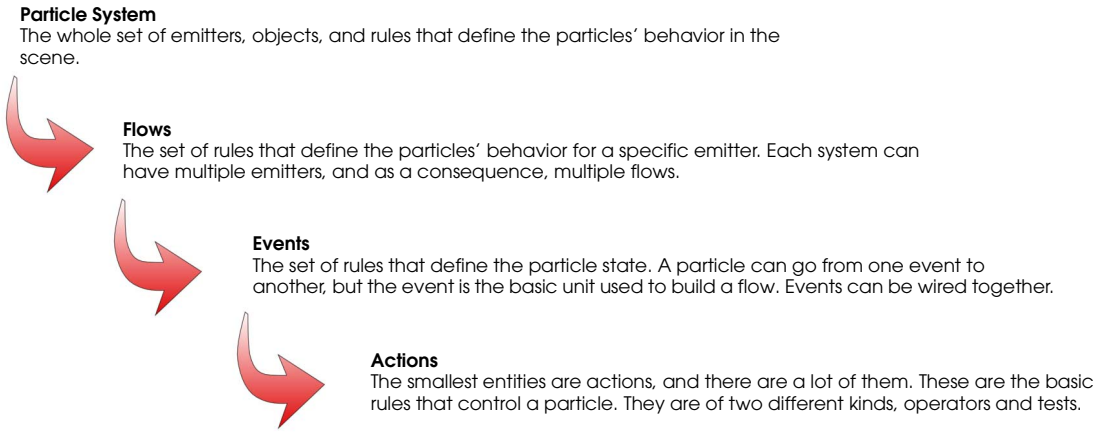

Reversing the order of the hierarchy might give you a better understanding. You can select the actions that define a specific state of a particle and group the actions into *events*. Then, you use tests to send particles from one event to the next, and to create a *flow* for the specific emitter.

Developing the flow starting from events, then adding actions and events as necessary is a more common approach, due to the sequential nature of the particles' behavior.

Actions in Detail

Now that you have an overall understanding of how Particle Flow works, you can look at the actions; the possibilities of the system will soon become clear. Actions are divided into two categories: Operators and Tests.

Operators
Operators define individual particle attributes like size, shape, color, or speed. You use them at the beginning of any event to define the particles' appearance for that state. For example, after a particle is created, you want it to assume a specific color or material, a specific size, or to start moving while being affected by gravity.

On the next page, we list all the operators, and give a brief description of each one. Every operator has an exhaustive set of parameters, and you will find them in the descriptions in the online *User Reference*.

	Primary Operators
	**Birth:** Enables the creation of particles and sets the creation attributes.
	**Birth Script:** Allows the use of a script for custom particle creation.
	**Delete:** Deletes (kills) a particle.
	**Force:** Adds a space warp force that influences the particle's trajectory (wind, gravity, vortex).
	**Keep Apart:** Applies a force to particles that get too close, trying to avoid collisions.
	**Mapping:** Sets the color of particles based on map coordinates.
	**Material Dynamic:** Changes material IDs for individual particles over time.
	**Material Frequency:** Assigns sub-materials to particles based on a specific frequency.
	**Material Static:** Assigns static material IDs to particles.
	**Position Icon:** Sets the initial position of particles on the emitter at birth.
	**Position Object:** Allows particles to be emitted from any geometric object.
	**Rotation:** Lets you set and animate the particles' initial orientation.
	**Spin:** Applies a spin to the particles.
	**Scale:** Changes the size of particles by scaling them.
	**Script:** Lets you create new behaviors by writing a script.
	**Shape:** Sets a default shape for particles, like cube, sphere, tetra, or vertex.
	**Shape Facing:** Sets the shape of particles to be a rectangle always facing the camera or an object.
	**Shape Instance:** Lets you specify a scene object as your particle.
	**Shape Mark:** Creates mapped rectangle-shaped objects that can be used as marks on collisions.
	**Speed:** Defines speed and orientation of particles, based on the emitter.
	**Speed By Icon:** Uses special icons in the scene to define speed and orientation of the particles.
	**Speed By Surface:** Uses an object's surface to define speed and orientation of the particles.
	Utility Operators
	**Cache:** Stores the particle state to memory for faster viewport operations.
	**Display:** Specifies how particles appears in the viewports; can be used to speed up feedback
	**Notes:** Adds comment text to an event
	**Render:** Provides rendering controls for the particles

Most of the attributes of these operators are animatable.

The Birth operators can be used only at the beginning of a flow.

Tests

Tests are continuously evaluated; each test returns a value of True or False. Tests can check for a lot of things, like speed, collisions, age, or scale. You set your conditions, and if the test returns True, you can send the particle to another event. For example, you can test the particle speed, and if it tests higher than a specific value, you can then forward it on to another event, where the particle can be slowed down or stopped.

If the test returns False, the particle remains in the same event, unless other tests return True.


Age: Checks if the particle, event, or absolute age is greater than or lesser than a specific value.


Collision: Tests for collisions with deflectors and related events.


Collision Spawn: Tests for collisions and spawns a number of particles if the collision is detected.


Find target: Creates a target object and directs particles to it. When the target is reached, the test returns True.


Go To Rotation: Performs a smooth rotation to a specific orientation. Returns True when rotation is complete.


Scale: Checks the scaling amount and branches accordingly.


Script: Allows you to use a script as a test operator.


Send Out: A pseudo test that always returns True or False.


Spawn: Creates a set of new particles and sends them on to another event.


Speed Test: Performs various tests related to speed, acceleration, and steering.


Split Amount: Returns True only for a specific amount of the tested particles.


Split Selected: Returns True only for the particles that are selected in sub-object mode.


Split Source: Returns True only for the particles coming from a specific emitter.

Predefined Flows

Two special items available from the same action list are predefined flows that you can use as a starting point for your custom flows.


Empty Flow: Creates a single event flow and an emitter, but no generation of particles.


Standard Flow: Creates a default particle emitter generating particles, ready to be customized.

Some actions automatically create helper objects in the scene, such as emitters, target objects, or other icons necessary as references. Deleting the objects removes the actions, and deleting the actions removes the objects.

The Particle View Window

Particle View is the editing environment for flows and events. When you open Particle View, you enter a graphical editor, similar to the Schematic View. In this topic, you'll see how the parts are visualized graphically, and how things works together. After learning the basics, this editor should be easy to understand and manage.

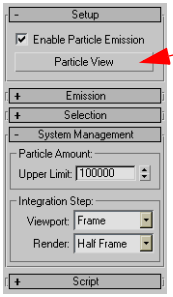

Click the Particle View button on the command panel to open the Particle View window. There are other general settings on the command panel, mainly to adjust performance of complex simulations. The Selection rollout will be explained later when we talk about sub-objects.

The menu provides some common editing tools, most of them also available as drag-and-drop interactive operations or as right-click menus.

A special source event connected to the emitter is always placed at the beginning of the flow. The particles' birth and the initial emission settings are in the second event of the flow.

Options for the selected action are displayed on the parameters panel. Here, a Speed test is selected and the options show what tests can be performed.

This is the editing work area, where flows are built, using a schematic-like view. You can zoom and pan this window.

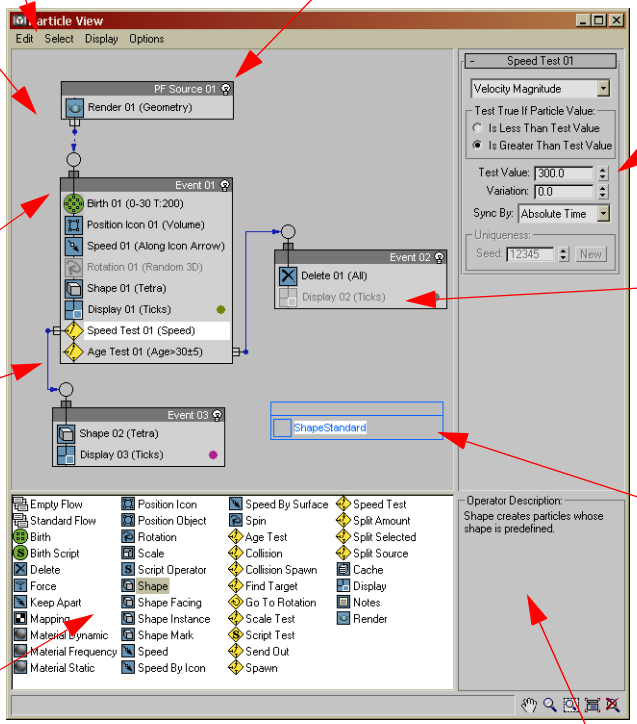

Events are groups of actions connected together to create flows. Click the little light bulb to enable or disable an event.

Individual operators can be enabled or disabled by clicking their icons. When disabled, the icons are unavailable.

Build a connection between the output of a test operator to the input of an event just by dragging a line between the two connection points. Drag a test output to the other side to move the anchor point and make the wiring more readable, as shown in the Age test here.

The easiest way to create a new event is to drag and drop a new action in an unused area. The event is created and the new action added.

The depot lists all the available actions (operators and tests). Select one and drag it to the location where you want to insert it. A blue bar shows you the insertion point above, below, or as a replacement (strikethrough) for an existing action.

The description panel displays a short explanation of the currently selected action in the depot.

A Sample Flow

Shown here is the flow for the little animation seen on the film strip, shown at the left, below. Only one particle is used, but that's how the flow works, one particle at the time. This sample is very basic, but already shows the potential of Particle Flow.

A cannon is ready to shoot a ball. The emitter is set to emit a single particle; it is hidden inside the barrel. The whole cannonball animation is controlled by the flow on this page.

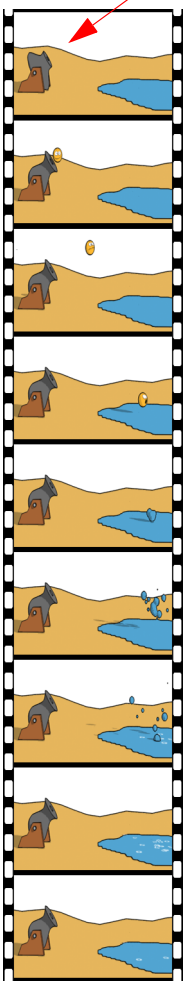

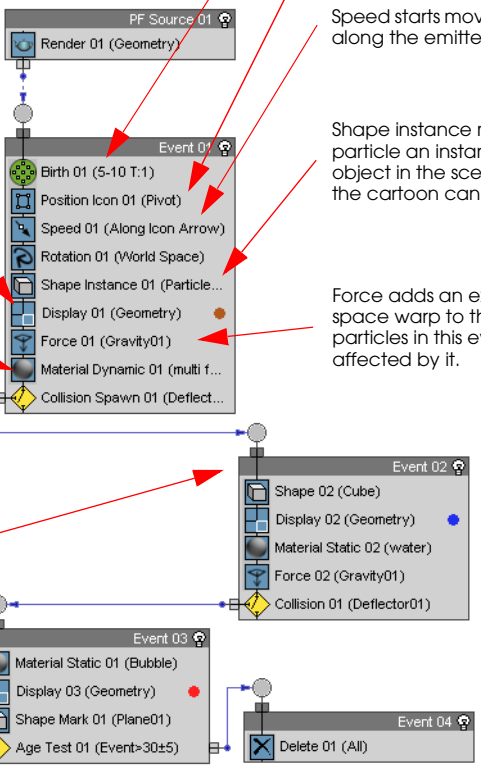

The birth operator emits from frame 5 to 10, for a total of 1 particle.

Position sets the particle's initial position to the emitter icon's pivot point.

Speed starts moving the particle along the emitter arrow.

Shape instance makes the particle an instance of another object in the scene, in this case, the cartoon cannonball.

Force adds an existing gravity space warp to the system; particles in this event will be affected by it.

Display enables the geometry of the particle (the cannonball) to be visible in the viewport. You can choose to see just a tick or other placeholders.

Materials dynamics changes the particle material over time. In our case, the face changes expression three times during the trajectory.

The Collision Spawn test creates and sends spawned particles when the first event particles collide with a deflector on the ground. Their shape is changed to cubes, their material to water, and the gravity force is added for them.

When the water particles collide with the ground again, they are sent to a third event, where a Shape Mark operator creates bubbles where they fall. While in this third event, a BlobMesh compound object is used to blend the particles together. See *BlobMesh* in this guide.

A Particle Age test keeps the bubble particles alive for approximately 30 frames, and then sends them to a final event where they are deleted, making the bubbles gradually disappear until the water surface is quiet again.

Selecting Particles

One special test, Split Selected, uses the sub-object feature of Particle Flow. When you enter sub-object mode, you can select particles individually or by event. Let's see an example of how to use this test.

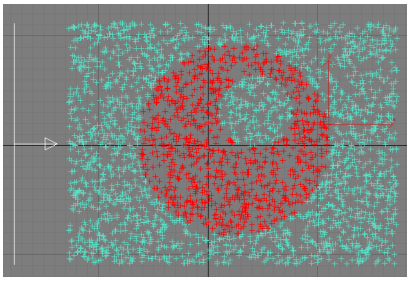

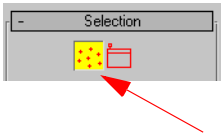

Set up your flow so particles are emitted as desired. Then, go to a frame where the particles are visible, enter sub-object mode (Particles), and select the particles you want to be affected. Particle Flow remembers the selection, even when the particles are not yet created.

Before selecting, make sure that in the Emission rollout on the command panel, the Quantity Multipliers are both set to 100% for viewport and rendering; otherwise, your selection will not be on the total number of particles.

Add the Split Selected test, and send them to an event where they are changed. The Delete operator has an option to act only on selected particles too, since this is a common option.

To correct the selection, temporarily disable the Split Selected test so you can select or deselect from all the particles again. See below how to make the test always return False.

Debugging Tests

Tests can be temporarily set to always return True or False. This is useful for debugging complex flows, or for forcing different behaviors from the same flow.

Click the left side of the test icon to set it to always return True. The icon turns into a green light bulb.
Click the icon again to restore normal functionality.

Click the right side of the test icon to set it to always return False. The icon turns into a red light bulb.

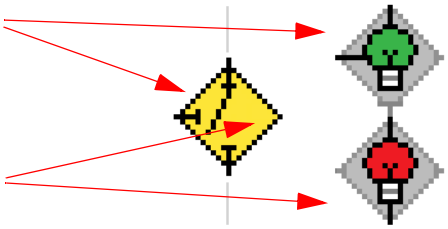

Particle Flow Scene Icons

Two tests rely on special helper objects in the scene to work. The helper icons define points in space that can't otherwise be specified in the Particle View window.

Each of these two tests, when added to an event, automatically creates a helper in the scene that needs to be moved appropriately.

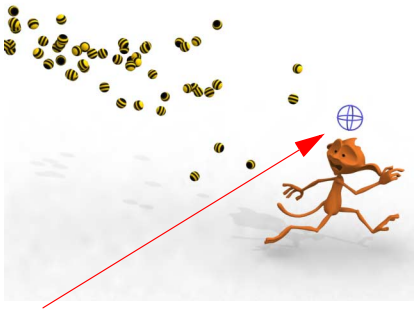

Find Target creates this helper. Particles in the event move toward this object. Move the object as desired in the scene, or link it to an existing animated object.

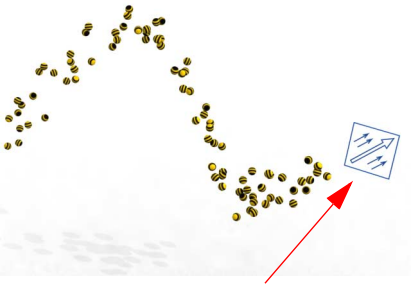

Speed by Icon uses the trajectory of its animated helper object to create a path for the particles to follow.

Selecting the helpers in the scene lets you change some particle parameters directly on the Modify panel without the need to open Particle View.

More Information

As we said, Particle Flow is a flexible and complex environment. Experimenting with all the features and looking at each action parameter are the next steps to get to know it better. Use the online reference as your companion for this stage, and you should be able to master it quickly. Be sure to complete the new Particle Flow tutorials as well.

Particle Flow is able to do much more than the few samples we used in these pages.

Channel Info

Channel Info is a low-level editing tool that gives you the power to manipulate object channels directly. This utility is found on the Tools menu, and is mainly targeted to technical directors or very experienced users.

Every geometric object has channels. They are used to define various surface characteristics of objects, from the mesh itself (for mesh objects), to the UVW mappings, and vertex selections.

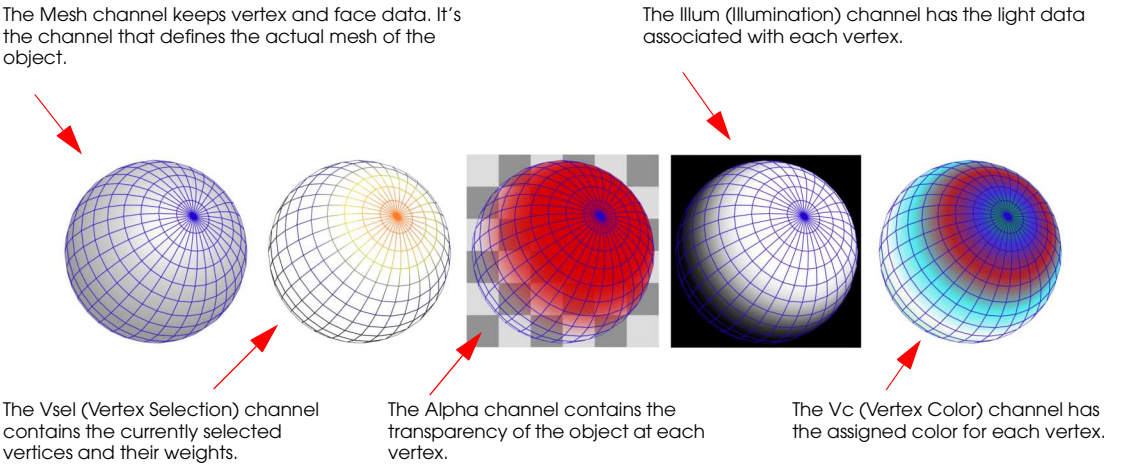

Channels are like data arrays, and the data they contain can be of different kinds. With this tool, you can copy, paste, and manipulate channels in various ways. However, pasting the data from one channel into a channel of a different type (different topology) can cause unexpected results.

Channels other than the basic ones (Mesh and Vsel) are numbered. Channel -2 is Alpha, -1 is Illum, 0 is Vc, and 1 to 99 are map channels.

This tool can be useful to analyze map-channel usage in models done by other artists, or to delete channels in collapsed meshes. Also, if you have two identical models, you can copy and paste mapping channels between them.

Another interesting trick that can be done with this tool is the ability to store multiple vertex selections, and then manage them with the help of the new Select By Channel modifier. You'll see this in an example later.

Channels are added, pasted, and cleared using special modifiers in the stack. You will see them added automatically. You can manually change them if necessary. Collapse the mesh when done.

Channel Info User Interface

Select a channel from the list, then select the operation you want to perform from the toolbar. The right-click menu also has some of the toolbar operations.

Copy and paste channel data from a source to a destination channel. Pay attention when pasting into channels with different topology. Paste is disabled when the destination channel is not compatible, but in some cases it's still allowed.

Name renames the channel. Names are used to give a meaning to a channel, and are also used by the Select By Channel modifier.

SubComp opens channels with sub-channels like X,Y,Z and R,G,B, and gives access to the individual components.

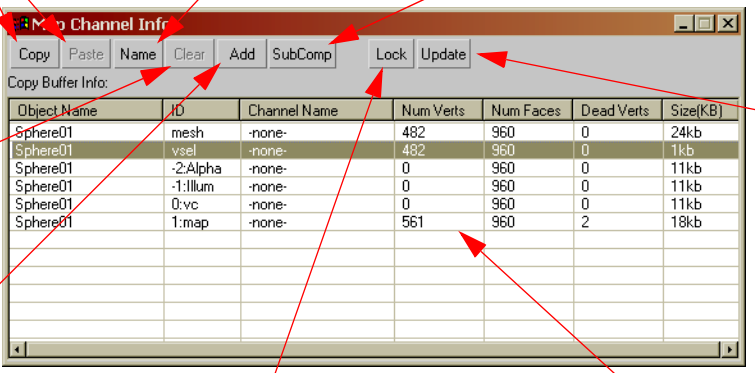

Clear deletes the data for a channel. You can't delete mesh and vsel channels.

Update refreshes the list when changes to the object geometry or topology are done from outside the tool, like from the modifier stack.

Add appends a new mapping channel after the last one. Mapping channels need to be consecutive. If you create map channel 10, empty channels 1 to 9 will be created as well.

Lock keeps the list focused on the current object, so you can select other objects, and keep the same channel list visible. Note that the list can show the channels of multiple selected objects.

The number of vertices in the channel can give you an indication if the channel is used, and if it's compatible for a paste operation with another channel.

When selecting SubComp, the channel list grows and shows all the individual channels' components. You have more choices for copy and paste combinations.

With SubComp enabled, you can see the X,Y,Z and R,G,B components of channels.

Sphere01	mesh:X	-none-	482	960	0	1kb
Sphere01	mesh:Y	-none-	482	960	0	1kb
Sphere01	mesh:Z	-none-	482	960	0	1kb
Sphere01	vsel	-none-	482	960	0	1kb
Sphere01	-2:Alpha...	-none-	0	960	0	0kb
Sphere01	-2:Alpha...	-none-	0	960	0	0kb
Sphere01	-2:Alpha...	-none-	0	960	0	0kb
Sphere01	-1:Illum:R	-none-	0	960	0	0kb
Sphere01	-1:Illum:G	-none-	0	960	0	0kb
Sphere01	-1:Illum:B	-none-	0	960	0	0kb
Sphere01	0:vc:R	-none-	0	960	0	0kb
Sphere01	0:vc:G	-none-	0	960	0	0kb
Sphere01	0:vc:B	-none-	0	960	0	0kb
Sphere01	1:map:X	-none-	561	960	2	2kb
Sphere01	1:map:Y	-none-	561	960	2	2kb
Sphere01	1:map:Z	-none-	561	960	2	2kb

An Example Using Vertex Selections

Some copy and paste operations appear obvious, like copy and pasting map channels from one to another, but there is more that can be done.

For example, X,Y, and Z components of a map channel can be used to store vertex selections. When you do this, the map channel loses any meaning for mapping, of course, so don't use it for that purpose.

First, let's prepare the object. For this example, we use a plane converted to an Editable Poly. In the Channel Info window, we add one extra map channel by clicking Add. We will use this extra channel for vertex selection storage only.

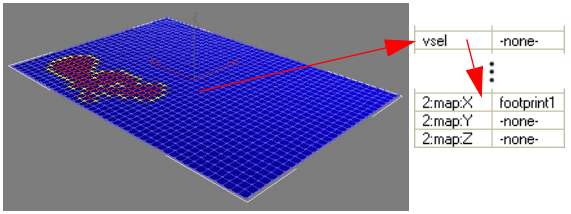

Select a set of vertices, using soft selection. These vertices are stored in the Vsel channel. When done, copy the Vsel channel, enable SubComp, and paste it into the X component of the map channel 2 we added. Name it appropriately when requested.

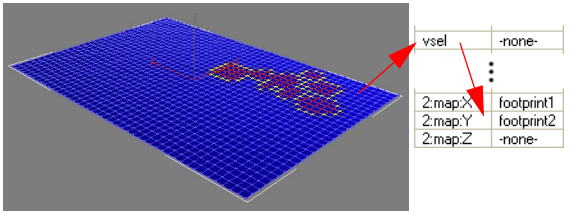

Go back and select another set of vertices. Copy and paste the Vsel channel into the Y component of the map channel 2, and give it an appropriate name.

The object is now ready. On the modifier stack, you will see a few modifiers added by the Add and Paste operations we did. You can collapse the stack and go back to an Editable Poly basic object; channels will not be lost. Now, let's see how we can use this setup.

On the modifier stack, we add a couple of Select By Channel and Push modifiers.

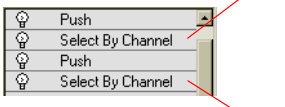

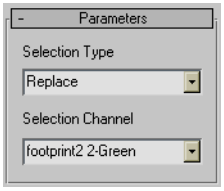

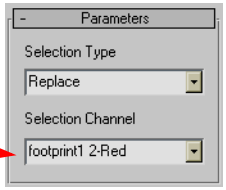

Select By Channel lets you select the saved selections from the list of all selections available.

With this setup, you will be able to animate each Push modifier independently, and even edit and add selections and other modifiers, without the risk of losing the original settings.

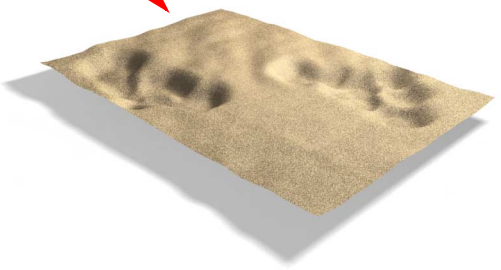

An interesting and important feature of selections saved this way is that they will survive a topology change if the original object is edited or other modifiers are added.

An Example of Blending Textures

The Vertex Color map can now select not just the current vertex colors, but also any copy saved to a map channel. You can also select any of the individual X,Y,Z or R,G,B subcomponents.

Copying a vertex selection to a map channel's X,Y,Z sub-components is also a way to create a vertex color channel from the vertex selection.

The Vertex Color map can then be used as a mask for other maps, or as an Opacity map in a composite material, using any of the channels available in the object map channels.

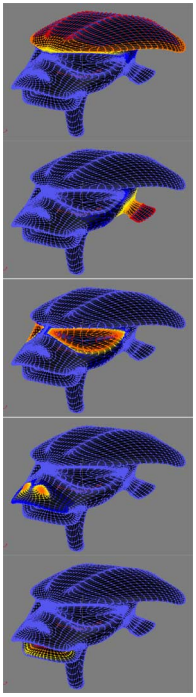

Here is an example of how selections can be used to create multi-material objects.
As seen before, select vertices and save the selections into map channel X,Y,Z components, and name them appropriately.

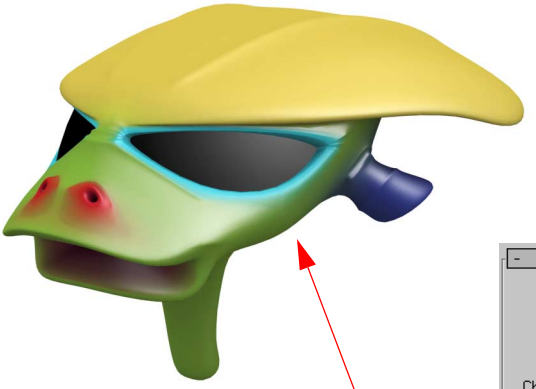

The sample shows how different materials blend into each other. A setup like this saves hours over preparing UVW-mapped masks.

Then, create a composite material, and add as many materials as you need. Each material will have an opacity map that uses the Vertex Color map to select the appropriate saved Vertex Selection channel.

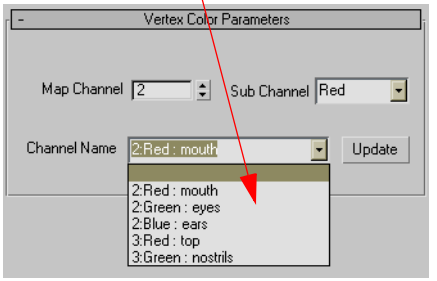

This technique can also be used to create masks for maps, rather than opacity for materials.

More Information

You can find more information in the online reference. Channel Info is a powerful tool, and probably requires a bit of experimentation and technical knowledge to be used to its fullest.

VertexPaint Modifier

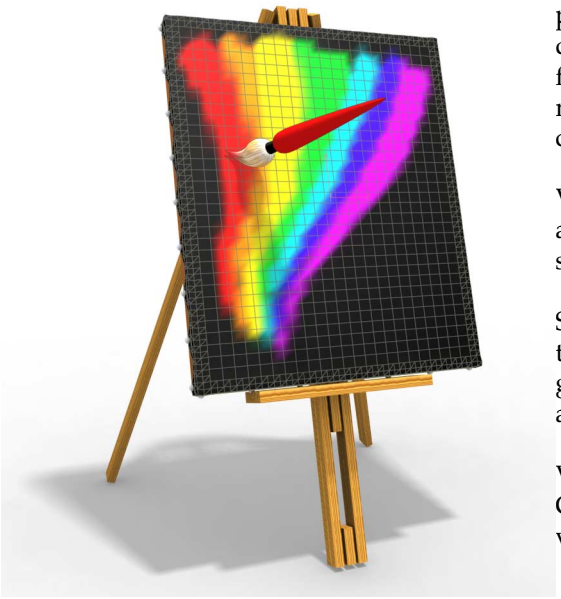

The VertexPaint modifier is a much more flexible and powerful tool in **3ds max 6**. Not only do the painting tools now have full pressure-sensitive device support, custom profile curves, and a floating palette with lots of new controls, but the modifier itself can be applied multiple times to create layers of paint that can be blended.

VertexPaint also supports painting on the 99 available map channels. For more information, see the new *Channel Info* utility.

Selections (including soft selections) can be used to limit the working surface of the object for greater control. Painting on multiple objects is also supported when they share the modifier.

VertexPaint, the Channel Info utility, the Vertex Color map, are now a complete set of tools for vertex color data manipulation.

Painting on Layers

When using multiple VertexPaint modifiers on the stack, you can consider each of them as a layer of paint, a concept common in painting applications. Each layer overlays the ones below, and has the ability to blend with them using different modes.

There are 15 different blending modes available, including screen, multiply, lighten, darken, hue, saturation, and color.

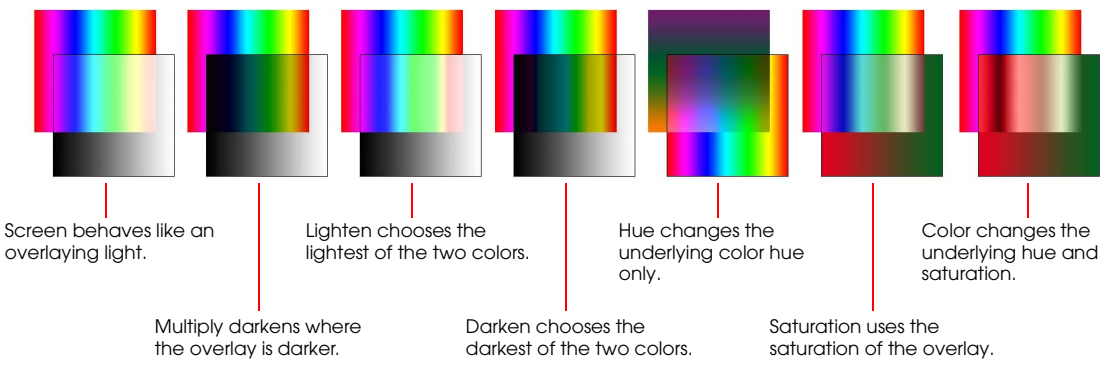

Brush Controls

VertexPaint brushes now use the same technology that weight painting uses in the Skin modifier. This means complete control over size and strength using a tablet (pressure-sensitive device) and customization of the brush falloff profile.

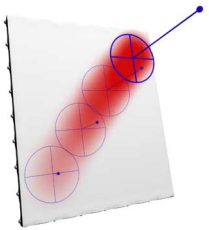

The input device pressure can control the strength of the brush. The length of the brush normal bar reflects the pressure applied.

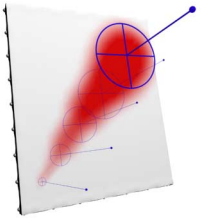

Alternatively, the input device pressure can control the size of the brush, as indicated by the brush ring.

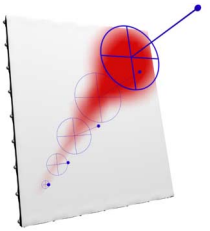

Size and strength can also be controlled together, for more natural brush feedback.

The Painter Options dialog is the same as used with the Skin Modifier. We summarize here the controls available; no new options have been added.

The Min and Max Strength and the Min and Max Size specify the weight and the size range of the brush. If used with a non-pressure-sensitive input device, only the Max values are used.

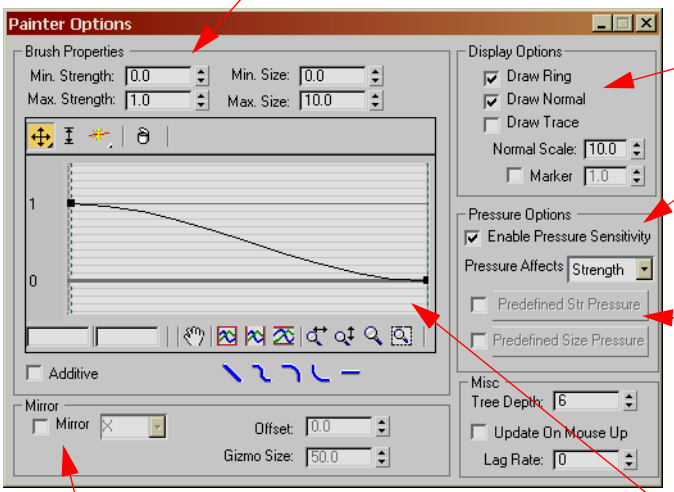

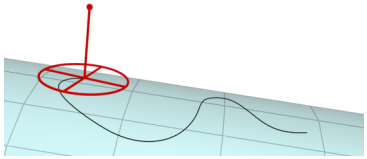

The brush gizmo Ring and Normal components can be disabled. The Trace is a temporary mark of the brush stroke on the surface.

With a pressure-sensitive device, like a tablet, you can use the pressure value to affect the brush size, strength, or both.

If you don't have a pressure-sensitive device, you can simulate the effect by defining a curve for the size or strength values over a brush stroke.

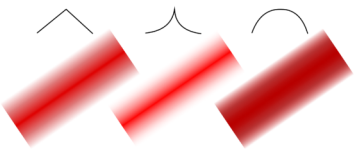

The brush strength falloff curve can be customized. You can also use one of the pre-defined curve shapes.

On symmetrical models, you might find it useful to mirror the brush on one axis, and paint on both sides.

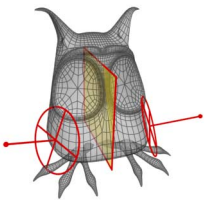

The Command Panel Parameters

The command panel rollouts have been changed. Some options have been moved to the floating Vertex Paintbox, and new options are available. Some options have been removed and replaced because of the workflow change. Overall, the modifier's usability has been made much more intuitive and flexible.

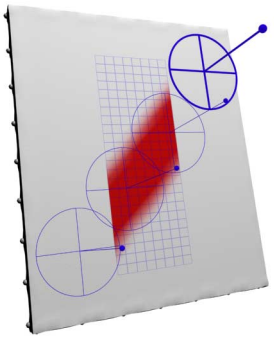

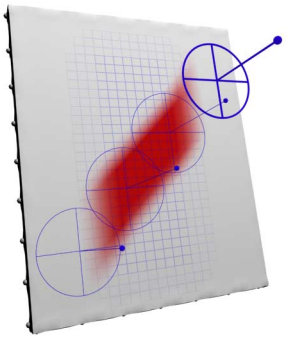

Note that the selection tools are also duplicated in the floating Vertex Paintbox. Since it's a modeless dialog, it allows for stack management without the need to use the Modify panel.

Selection enables sub-object selections of vertices, faces, and elements. When in sub-object mode, only the selected elements are affected by the brush strokes.

Soft selection is supported. This allows for painting selected areas with smooth blending.

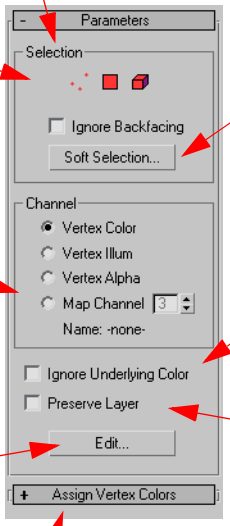

This option turns off the blending mode for this instance of the modifier. This can be used to temporarily see and fix the actual color data rather than the blended result.

Select the channel you want to paint on. You can also paint on 99 extra map channels. See the *Channel Info* utility for more information.

Preserve Layer avoids this instance of the modifier being deleted when the VertexPaint layers are collapsed. See *The Vertex Paintbox*, on the next page.

Edit opens the Vertex Paintbox. The palette opens automatically when the modifier is applied or selected. If you close it, you can re-open it here.

The Assign Vertex Color utility is duplicated here for easy access. This utility is useful to bake lighting and provide a new layer of paint, or a starting point for vertex color editing. For more information, see the *Assign Vertex Color Utility* topic.

Unlike in previous versions of **3ds max**, the command panels are no longer where you will find the painting tools. All the painting work is now done on the floating Vertex Paintbox.

The Vertex Paintbox

This modeless dialog is the main user interface for all vertex painting work. It's a single workplace where, in addition to painting, stack (paint layers) management and sub-object selection can be performed.

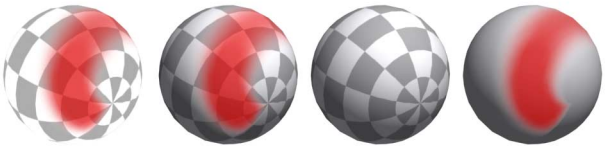

These buttons control the object display: vertex map unshaded, shaded, on/off, and maps off. They are toggle switches, so more combinations are possible than the ones shown on the left.

Here, you set which channel you want to display. The first three flyout options are the default ones (Vertex Color, Illumination, Alpha). The fourth option instead allows you to specify one of the 99 map channels.
The Lock button locks the display and paint channels so you are not looking at a different channel than the one you are painting on.

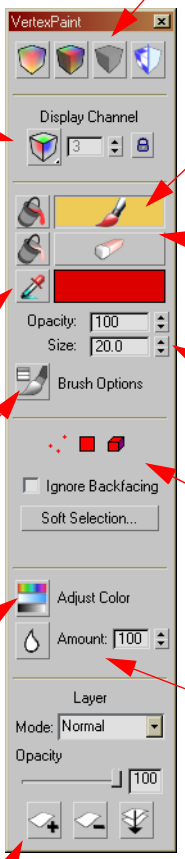

Click here to enable the brush and start painting. Click the button on the left with the bucket to paint everything (or just the selection if in sub-object mode). The bucket is useful to set the starting background color.

Click here to begin erasing the vertex colors. This tool uses the same brush as for painting. Click the button on the left to erase everything (or just the sub-object selection).

Choose the paint color using the color picker, or the Color Selector when you click the color swatch.

Sets the opacity and the size of the paintbrush.

Here, you can enter sub-object mode and select the vertices, faces, or elements to limit the painting area. For smooth blending, soft selection is supported as well.

Click here to open the Painter Options dialog discussed previously.

Opens the Adjust Color dialog (below), where you can change Hue, Saturation, Value, or RGB intensity.

Blurring works by mixing the colors of adjacent vertices to smooth irregular brush strokes or hard edges when sub-object selection is used. It acts on the whole object or the current sub-selection, as shown below.

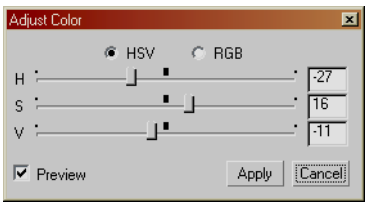

These are the layer controls. See the next page for a full explanation.

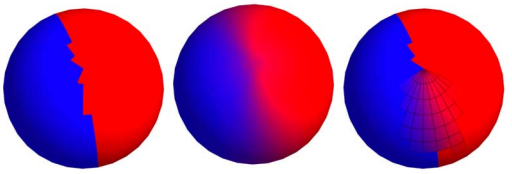

Layer Controls

The ability to paint on layers is achieved using multiple VertexPaint modifiers on the stack. To avoid using the stack in the Modify panel, the floating Vertex Paintbox gives you basic control of layers in order to make modifier management more intuitive.

Selects the blending mode for the layer (modifier instance). The layer blends with the underlying ones. If it's the first layer, it blends with the object's vertex color channel. At right is the list of the available blending modes,

Opacity makes the layer more or less transparent; lower opacity values reduce the effect of the layer colors in the blending evaluation.

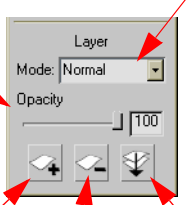

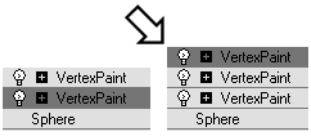

Adds a layer by putting a new VertexPaint modifier at the top of the stack. An alert dialog asks what channel the new layer will affect.

Deletes the current layer by removing the corresponding VertexPaint modifier from the stack.

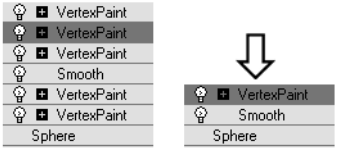

Collapses the stack by consolidating all the VertexPaint modifiers into a new one that will be added to the stack. Other modifiers are not affected.
If a VertexPaint modifier has Preserve Layer enabled, it is not deleted.

Assign Vertex Color Utility

This utility has been enhanced with new options. It also uses the new VertexPaint modifier instead of the basic Vertex Color, adding a whole new set of possibilities for blending using multiple assignments, and vertex color adjustment.

Two other basic object channels (Illumination and Alpha) are available for assignment, not just the vertex color one.

You can also use any of the 99 map channels available to store the vertex color information. See the *Channel Info* utility for more information on channels and names. The Vertex Color map will also retrieve color information from the extra map channels.

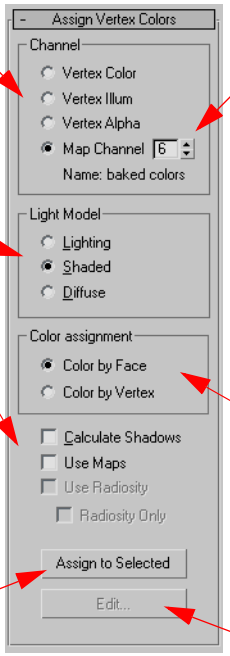

These options are unchanged from the previous version.

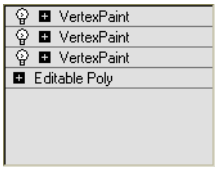

Color by face assigns the colors to the vertices on a per-face basis.

Assign to Selected does the vertex assignment, adding a Vertex Paint modifier to the objects with the calculated colors. Multiple assignments will result in multiple Vertex Color modifiers that can be later adjusted or blended.

Edit will bring you directly to the newly assigned Vertex Paint modifier for further editing.

The Skin Modifier's Mirror Mode

When modeling characters, it's quite common that you'll use symmetrical modeling. The Skin modifier can now take advantage of symmetry and save you the time to set up envelopes and adjust vertex weights for half of the model.

Mirror mode is available from the Mirror Parameters rollout in the Modify panel for Skin. You enter Mirror mode by clicking the Mirror Mode button; to exit, click it again.

Mirror mode is available only when in Envelope sub-object mode.

The Basics

When entering Mirror mode, the display of envelopes and vertices changes. An orange gizmo defines the mirror plane. Vertices and envelopes on one side are colored blue, and on the other side green. Unmatched vertices and envelopes are displayed in red.

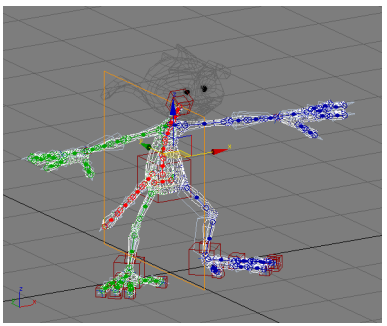

Note: since the viewport display might be confusing for these illustrations, we will use conceptual images to be clearer.

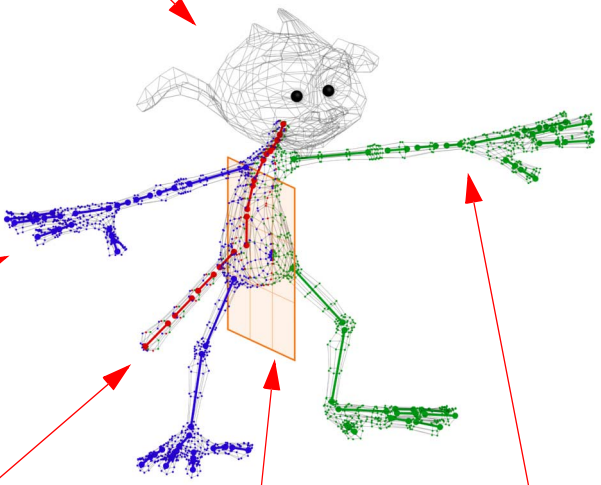

Blue envelopes and vertices are on one side. When colored blue or green, a match between objects on both sides of the mirror plane has been made.

Red envelopes and vertices do not have a match on the other side of the mirror. In this case it's not a problem, since these objects are common middle objects of our model.

The mirror plane can be oriented in X, Y, and Z, and positioned to match the mid-line of your symmetrical model.

Green envelopes and vertices are on the other side of the mirror plane.

Symmetrical Items Matching and Colors

When in Mirror mode, the software attempts to match envelopes and vertices between the two sides of the mirror plane by searching for a match within a certain volume.

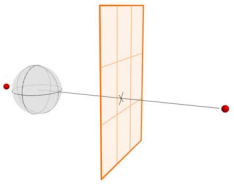

For each vertex and envelope on one side of the mirror plane, the software looks for a match on the other side. If nothing is found within a certain threshold, no match is done; in this example, the unmatched vertices are shown in red.

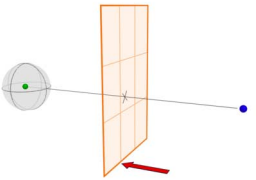

Adjusting the mirror plane, a match is found within the threshold, and the vertices are colored blue and green.
You can set the search threshold, but If you modeled a perfectly symmetrical object, the match is never a problem.

Selected bones or vertices are highlighted in yellow. When you make a selection, the matching elements on the other side are also highlighted. You can select objects on both sides of the mirror plane. This software can transfer data in either direction.

Using Mirror Mode

Mirror mode only transfers envelope and vertex data between the two sides of the mirror plane. All your envelope and vertex weighting has to be done beforehand, as in earlier versions of the software. In **3ds max 6**, you need only do it on half of your model, and rely on Mirror mode to complete the skinning job.

Step 1

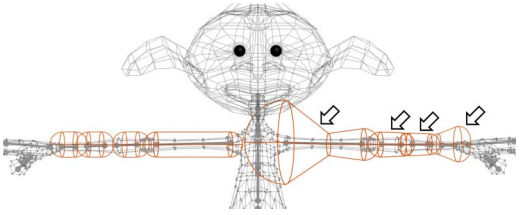

Edit envelopes and adjust vertex weights for half of your model. You can leave the other half with the default envelopes and weights.
In this example, we show four adjusted envelopes on the right side, and four default ones on the left.

Step 2

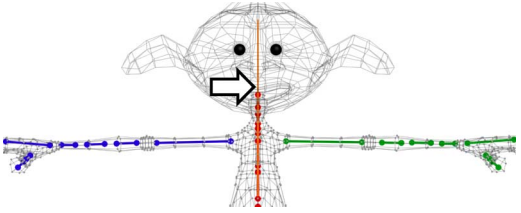

Once ready, enter Mirror mode, and make sure the mirror plane is in the right position, orientation, and that all vertices and envelopes match. The mirror plane is initially created at the pivot point of your mesh.

Step 3

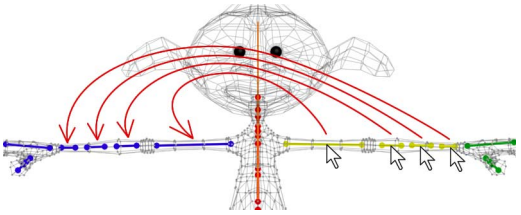

Select the envelopes or vertices you want to transfer, then use Mirror Paste to transfer the skin data to the other side. You can use the selection filters for vertices and envelopes while selecting.

Step 4

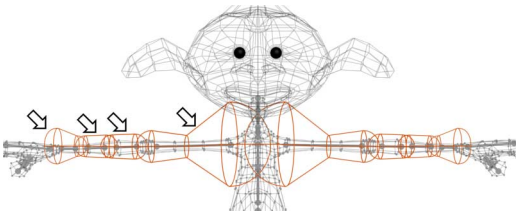

Exit Mirror mode, and the other side of the model will now match the one you adjusted. Test the skinning, and if it needs fixes, adjust as necessary, and then repeat the Mirror mode operations.

While the example in the illustrations above uses only envelopes, the procedure is the same as when mirroring vertices.

To transfer everything from one half to the other there are special options available. See below, in the *User Interface* section.

Display Projections

Projections are helpful to set the mirror plane in the right position and to troubleshoot meshes that are not perfectly symmetrical.

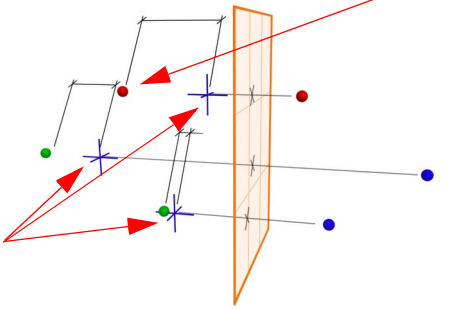

Distances between projected crosses and vertices can also highlight trouble areas on the original mesh, where the symmetry might not be as good as expected.

You can project the vertices of one side to the other. They will be displayed as crosses. Zooming in, you can see how close they are to the vertices on the other side, and you can adjust the mirror plane offset to try and overlap them as closely as possible.

Mirror Mode User Interface

This button enters Mirror mode. It's enabled only in Envelope sub-object mode. To exit Mirror mode, click the button again.

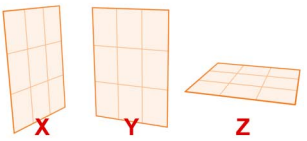

You can align the orientation of the mirror plane to the local X, Y, or Z axis.

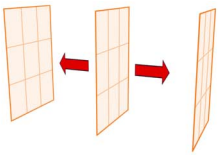

This offsets the mirror plane toward one side or the other.

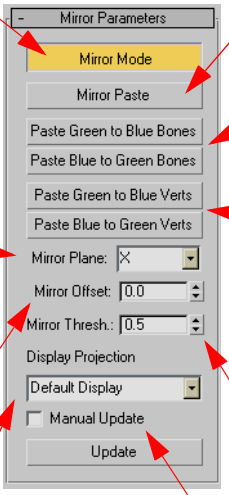

Mirror Paste transfers the selected envelope and vertex skin data to the other half.

These two paste buttons transfer all the *envelope* data from one side to the other.

These two paste buttons transfer all the *vertex* data from one side to the other.

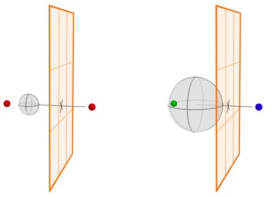

Here, you can set the threshold for the matching of the items on the two sides of the mirror plane.

Default Display shows the green and blue vertices; when one vertex is selected, the matching one is highlighted, as well. Positive Direction and Negative Direction project the vertices from one side to the other and vice versa, as crosses. The last option, None, will not display any vertex.

Manual Update is useful for complex meshes, when feedback is too slow. Enable it, and click Update when you want to refresh the display.

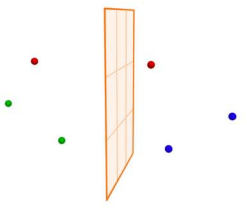

Default Display

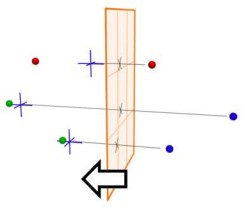

Positive Direction

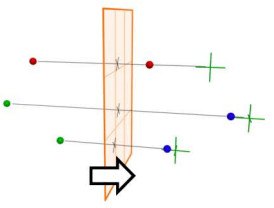

Negative Direction

BlobMesh Compound Object

A new compound object, BlobMesh, can create metaballs from geometry or particles. Metaballs are spherical entities that blend together when close to each other. They can be used for modeling organic surfaces or liquid effects.

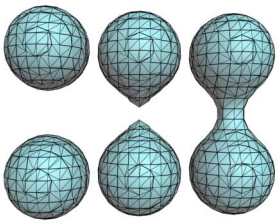

First, create a BlobMesh object somewhere in the scene, then switch to the Modify panel. Click Pick in the Blob Object group, and select one or more objects or a particle system you want to use as the source.

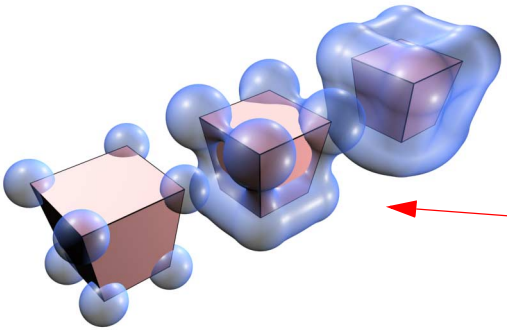

When you use BlobMesh with geometry, metaballs are placed at each vertex of the mesh. You specify the size of the metaballs in the BlobMesh rollout. Increase or decrease the size to adjust the metaballs and the blending.

When used with a particle system, BlobMesh combines the particles into a liquid-like surface. For particle systems, the size is not specified by BlobMesh, but it's taken from the particle size. Adjustments in that case have to be done in the particle system's parameter rollouts.
If you are using Particle Flow, you can also specify which events of the system will be used by BlobMesh. See the *Particle Flow* topic in this guide.

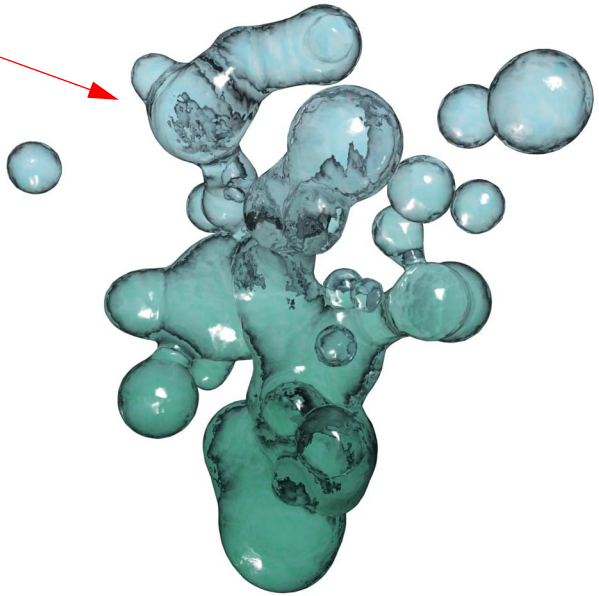

BlobMesh User Interface

Size sets the radius of the metaballs when meshes or splines are used as source objects. However, this value has no meaning when using particle systems as source objects, since the size is set by the particles.

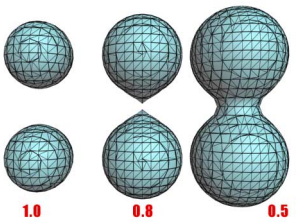

Tension makes the surface tighter or looser. Values under 1.0 relax the surface and create more extended surfaces.

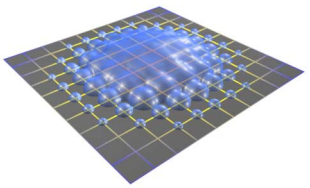

Using soft selection to select vertices on the source mesh using soft selection, makes the metaballs smaller along the selection edges. Set the minimum size appropriately to get the correct edge effect.

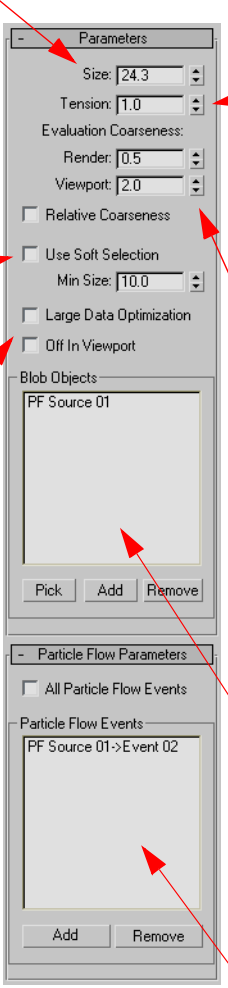

BlobMesh can be pretty heavy to calculate for complex surfaces. Large Data Optimization improves performance when the number of metaballs is high, which is likely, for example, when using particle systems. Off In Viewport is useful once the settings are final; you only need it when rendering.

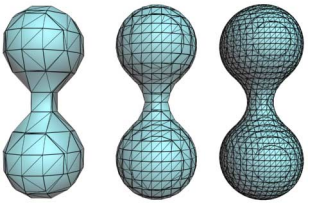

Coarseness sets the quality of the mesh subdivision for face sizes. Use different values for rendering and viewports, with lower viewport quality (higher values) for more responsive feedback.
Relative Coarseness changes the value's meaning to be relative to the metaball's size. In this case, higher values mean higher quality.

Add from a list or in the viewports, select the objects or particle systems you want BlobMesh to operate on.

Event 01

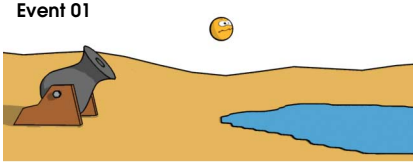

Event 02

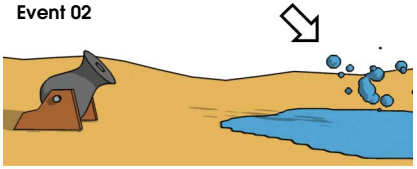

Event 03

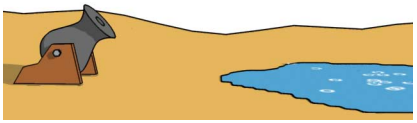

When using Particle Flow, you can choose which events BlobMesh should act on, or use all of them. In this sample, only the particles after the impact on the water use BlobMesh (Event 02), as specified in the list. See the *Particle Flow* topic in this guide.

Isoline Display

By default in **3ds max 6**, MeshSmooth Subdivision Surfaces does not display the edges of subdivided surfaces, but only the edges of the original editable control mesh. This will display a less confusing smoothed mesh, keeping only the edges that are important for the modeling process.

This option is available in the MeshSmooth modifier and the Editable Poly object, in the Subdivision Surface rollout.

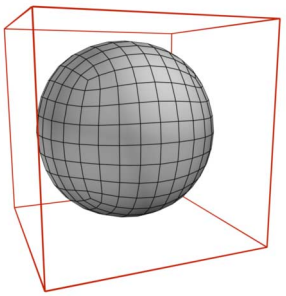

A mesh or poly model with MeshSmooth creates a lot of edges in wireframe mode, making details harder to see and adjust.

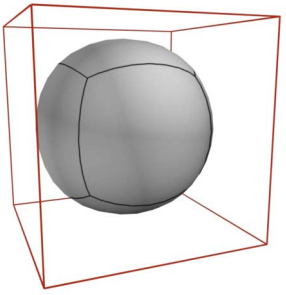

With Isoline Display on, the extra edges are not visible, but the object is still subdivided as before. Editing is now much easier, since the cage-control effect on the model is more clear.

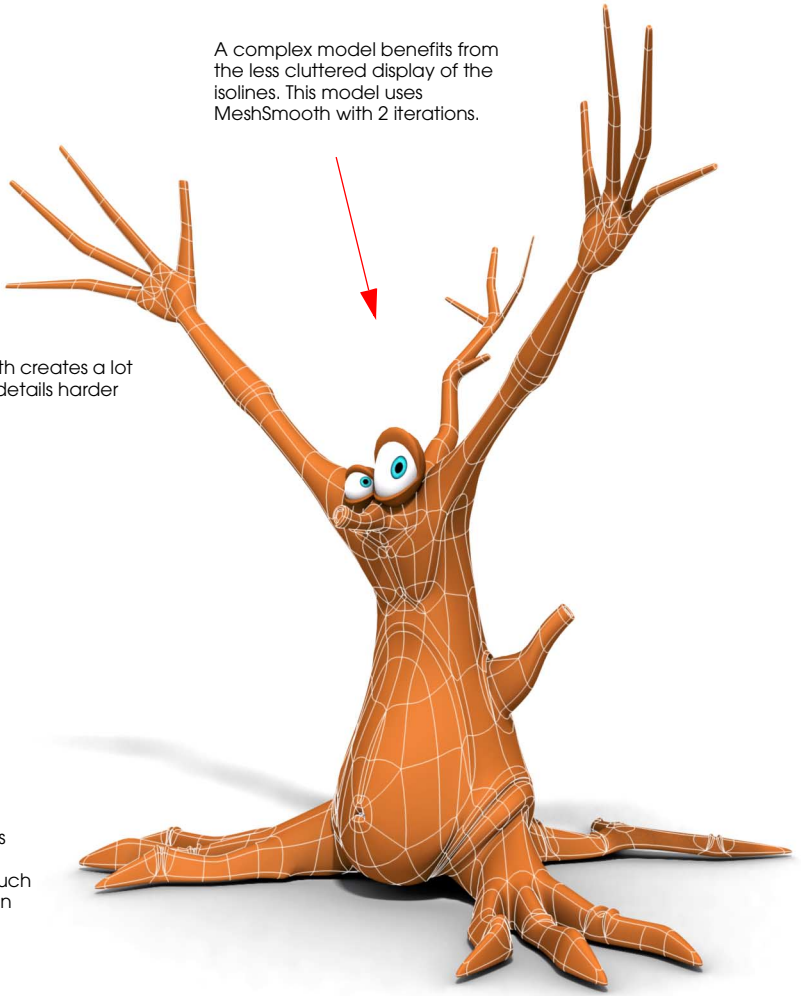

A complex model benefits from the less cluttered display of the isolines. This model uses MeshSmooth with 2 iterations.

In the MeshSmooth modifier, Isoline Display is valid only for NURMS and Quad Output subdivision options. It's not available for the Classic method.
Isoline Display is for viewports only. A wire material still renders each subdivided edge. Also, modifiers applied over the smoothed mesh might revert the display to a face-edged mesh.

Editable Splines

Improvements to the Editable Spline object and the Edit Spline modifier simplify the workflow for surface modeling, by adding automatic connection lines and cross-section functionality.

SHIFT+Copy Connection Line

When cloning a line segment or element, you can automatically generate the lines that connect the vertices of the two lines. Works with SHIFT+move, rotate, or scale.

SHIFT+moving, rotating, or scaling a segment or spline, automatically creates the vertices' connection lines. If SHIFT+cloning more than once, the connection lines will be joined.

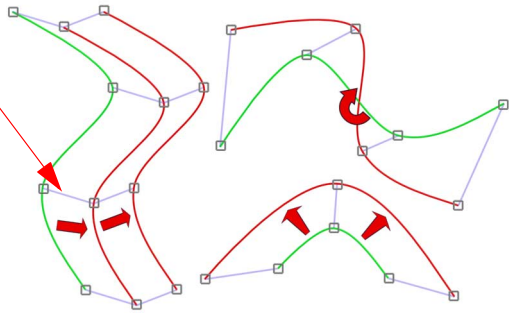

Turn on the new Connect option in the Editable Spline or Edit Spline modifier's Geometry rollout to use this feature. Threshold is the value used to find and extend connection lines when cloning more than once.

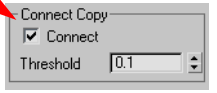

Cross Sections

Base Spline objects and the Edit Spline modifier now have the added functionality previously available only using the CrossSection modifier. Splines to be connected need to be part of the same spline object.

Before starting, set the vertex type you want for the cross-section lines.

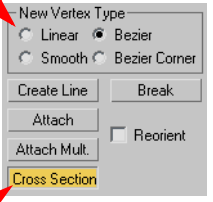

Turn on the Cross Section mode from the Geometry rollout. Click again to exit.

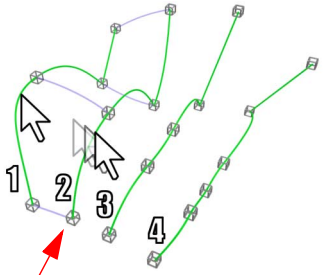

Select the lines to connect in the desired order. The cross-section lines will be created as you select.

When done, exit Cross Section mode. The lines created are now independent splines within the original object.

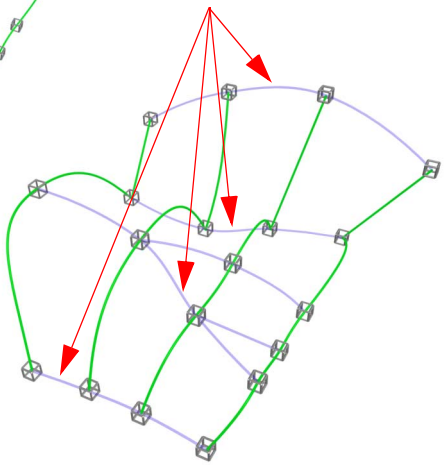

Copy and Paste Tangents

This feature lets you copy one tangent handle orientation, and then paste it to one or more other handles.

Click Copy to enter Copy mode. All the tangents in the spline will be displayed; you will be able to click the one you want to copy. Once selected, Copy mode closes automatically.

Click Paste to enter the Paste mode. All the tangents in the spline will be displayed. Choose the one(s) you want to match the tangent you just copied. You have to exit Paste mode manually, by right-clicking in the viewport or clicking the button again.

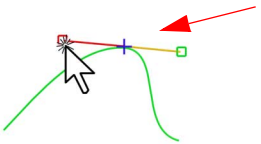

In Copy mode, select the tangent you want to copy. Note that you are actually selecting one half of it.

In Paste mode, click the tangents you want to align to the original one.

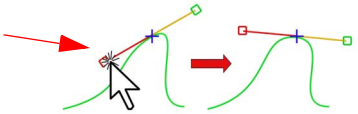

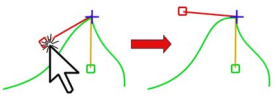

The other half of a Bezier tangent will move to stay aligned with the new orientation of the selected half. In a bezier Corner tangent, the other half is not affected.

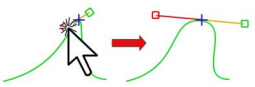

Use the Paste Length option to match the original tangent length as well. By default, only the orientation is matched.

Quad Menu

Two small improvements have been made to the quad menus. First, when you need to change the tangent type of one or more vertices, your needn't to be over one of them in order to activate the options. You can right-click anywhere, and the tangent options will be available, and will act on the selection.

Also, a reset option has been added: Reset Tangents restores the tangent handles to a default alignment position.

Editable Patches

A few enhancements to Editable Patch objects have been included to improve the workflow: copy, paste, and reset tangents, better soft-selection display, grow and shrink selections, smoothing, and the addition of surface creation to the Edit Patch modifier when applied to a spline object.

Handle Sub-Object Mode

A new Handle sub-object mode has been introduced to simplify tangent editing. In Vertex mode, you need to select the vertices to enable the handles. In Handle mode, instead, all handles become visible. Vertices are not selectable, making it easier to work on just the tangents.

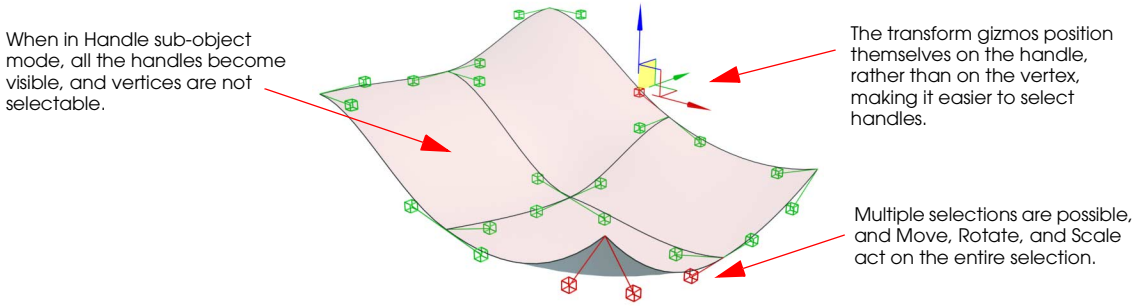

Also, the same ability to copy and paste tangent orientations and length available for Editable Splines is available for patch handles as well, and works in exactly the same way.
The Reset Tangents option on the quad menu is implemented here just as has been with Editable Splines.

Selection

The same selection tools available previously for Editable Poly are now available for Editable Patch.

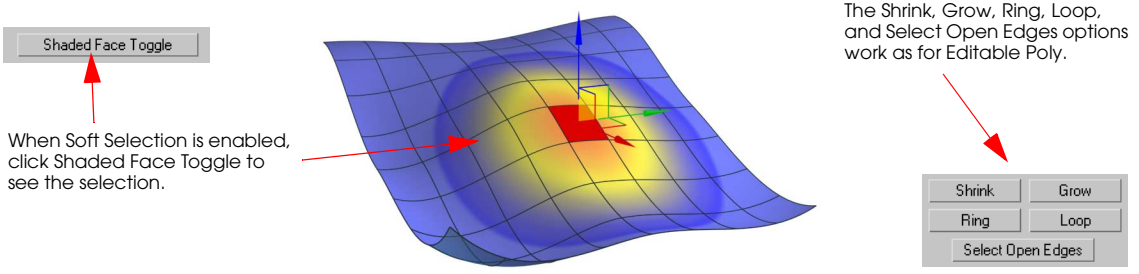

Patch Smooth

Patch Smooth is a new utility in the Miscellaneous group. It works on the entire patch surface or any sub-object selection, and aligns all the handles so that the resulting surface is as smooth as possible.

This example shows what Patch Smooth can do on a selected vertex whose handles create a corner on the surface.

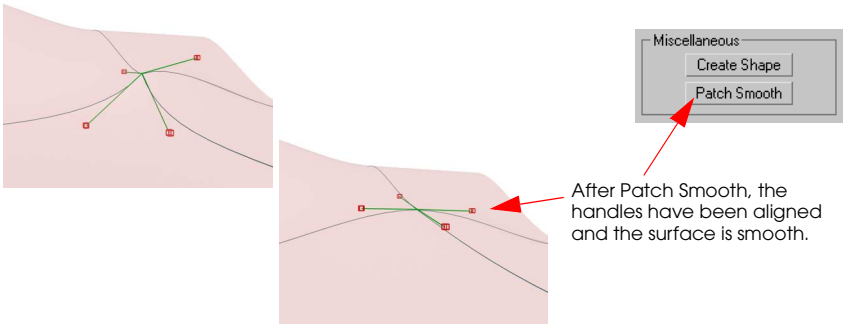

Handles are oriented based on the neighboring patches.

Edit Patch Modifier

The Edit Patch modifier has all the enhancements done for the Editable Patch object, plus the ability to turn a spline model into a surface. This same functionality was previously available with the Surface modifier.

When adding the Edit Patch modifier to a spline model, you have the same options to create the patch surface as you had in the Surface modifier. If Edit Patch is applied to an existing patch surface, these options will be disabled.

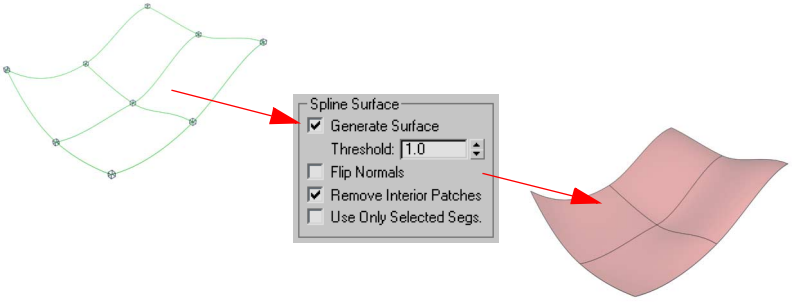

Now that editable splines have automatic connections and cross sections, you can create patch surface models with just two modifiers on the stack.

Previously, you might have needed up to three modifiers to create a patch model based on a spline.

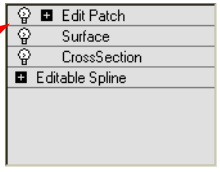

Now, cross section functionality is in editable spline, and surface functionality is in Edit Patch, and the workflow is greatly improved.

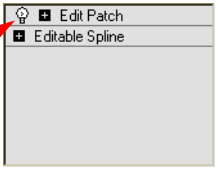

AEC Objects

AEC Objects are standard architectural assemblies like stairs, windows, doors, and walls. They are fully parametric pre-made objects, and give you a quick way to fill an architectural model with objects that otherwise would involve time-consuming modeling tasks.

Each object has several parameters for customizing look and dimensions. They already have material IDs assigned to their components for easier material assignments.

Windows
Windows styles include Fixed, Sliding, Awning, Projected, Pivoted, and Casement. You can change size and thickness of each component, including the frame.

Doors
There are different styles of doors, like Pivot, Sliding, and BiFold. You can have one or more panels, beveled or not, or glass panels. You can change the size and thickness of the door, the frame, and flip the hinge and the swing.

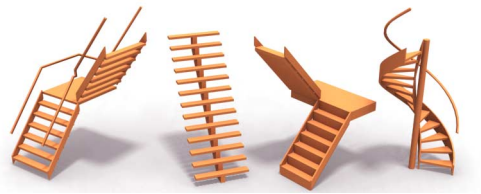

Stairs
You can create U- or L-shaped stairs, straight or spiral stairs. You can add components like stringers and open or solid-filled railings. You can change the angle of an L-shaped stair, the number of turns of a spiral one, and each stair has full control of overall height, number of risers, or riser height.

Railings
You can create a single straight railing, or pick a path and have a railing that can follow curved and off-plane lines. You can change the number of segments, customize the top rail size, and add pickets and lower rails.

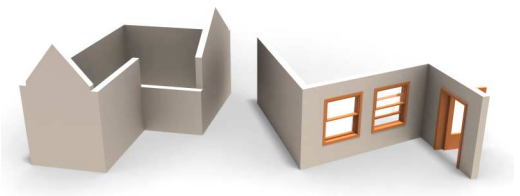

Walls
You an create walls almost as easily as you can create a spline, open or closed. A useful characteristic of walls, is that if you create and link doors and windows to them, holes will be created automatically and updated if the doors or windows are moved.

Foliage
You can add trees and plants from a library of different types. Each one has parameters for changing the appearance, like removing the leaves, or pruning the branches. Each plant is created randomly and it's always unique.

Camera Correction Modifier

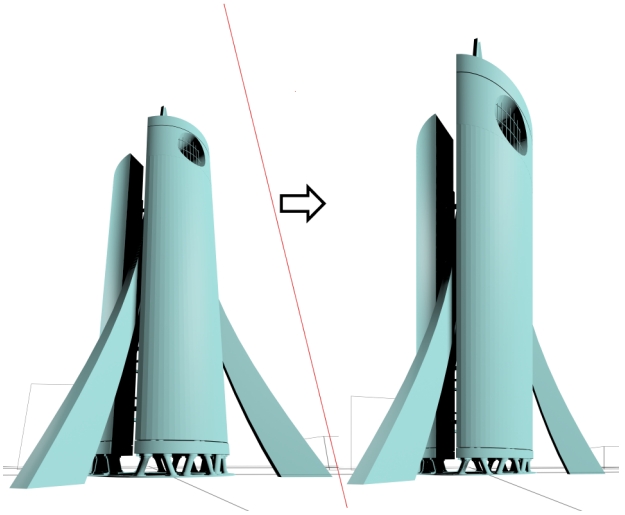

The Camera Correction modifier applies a two-point perspective to a camera. This is a common requirement for architectural visualization.

You can add the modifier to the selected camera from the Modifiers menu.

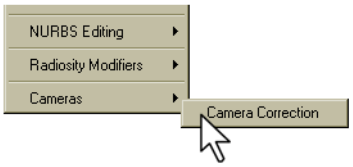

To change the modifier settings, go to the Modify panel and select the camera.

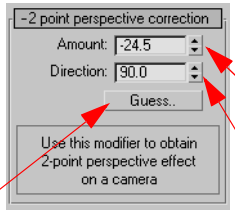

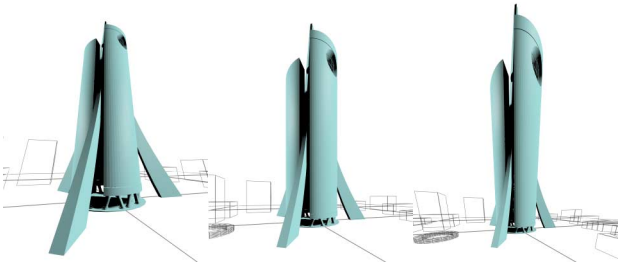

Amount is the strength of the correction. The camera will rotate to keep the center of the image in the same place. The camera cone gizmo will distort to give a feedback of the correction applied.

Guess estimates the amount of correction needed based on the camera rotation. You can then adjust the values manually. if necessary.

Direction is usually vertical (90 degrees), but needs to be adjusted if your camera is rotated or if you are looking for a different orientation of your perspective.

MapScaler Modifier

MapScaler was in previous versions of **3ds max**, but only as a World-Space Modifier (WSM). Now, it's also available as an Object-Space Modifier (OSM). They are both listed in the Modifiers drop-down list.

The behavior of the two versions is the same except for scaling. Both keep the map size unchanged if the object dimensions are changed, but the WSM version also maintains the size if the object is scaled, while the new OSM version will scale the map as well.

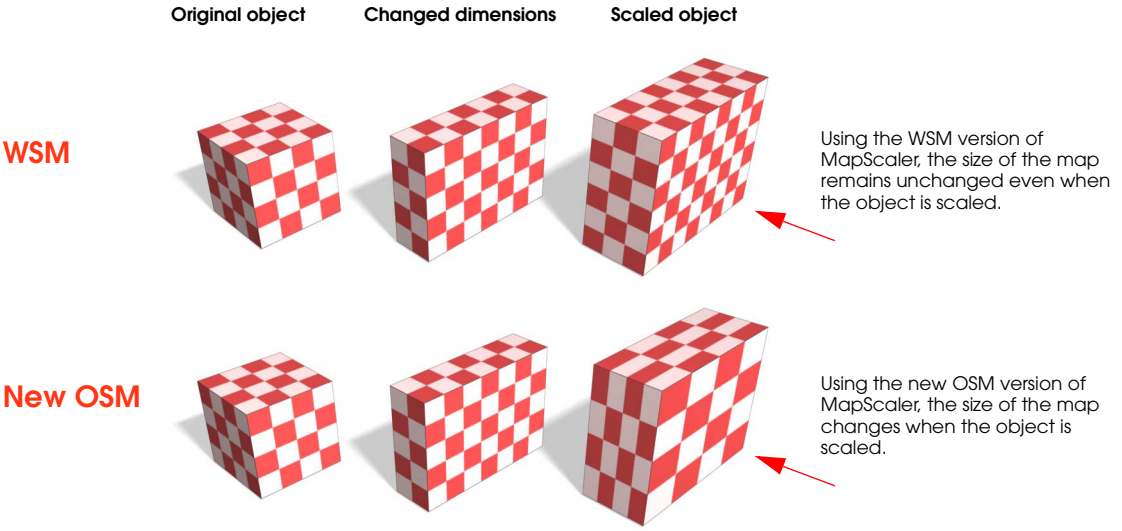

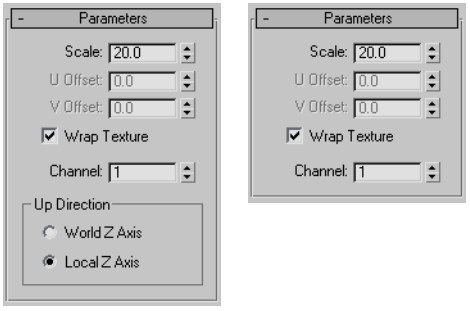

WSM New OSM

The user interface is practically identical between the two, except that Up-Direction selection is not necessary for the OSM version of MapScaler. The OSM version always uses the Local Z axis, while the WSM version allows you to chose between Local and World.

Unwrap UVW: Relax Tool

The new Relax Dialog command on the Tools menu of the Unwrap UWV editing window allows you to relax the UV vertices. This is similar to how the Relax modifier relaxes mesh vertices.

Relaxing UV vertices helps resolve overlapping or stretched areas of your model mapping. The way it works is better understood on a simple model.

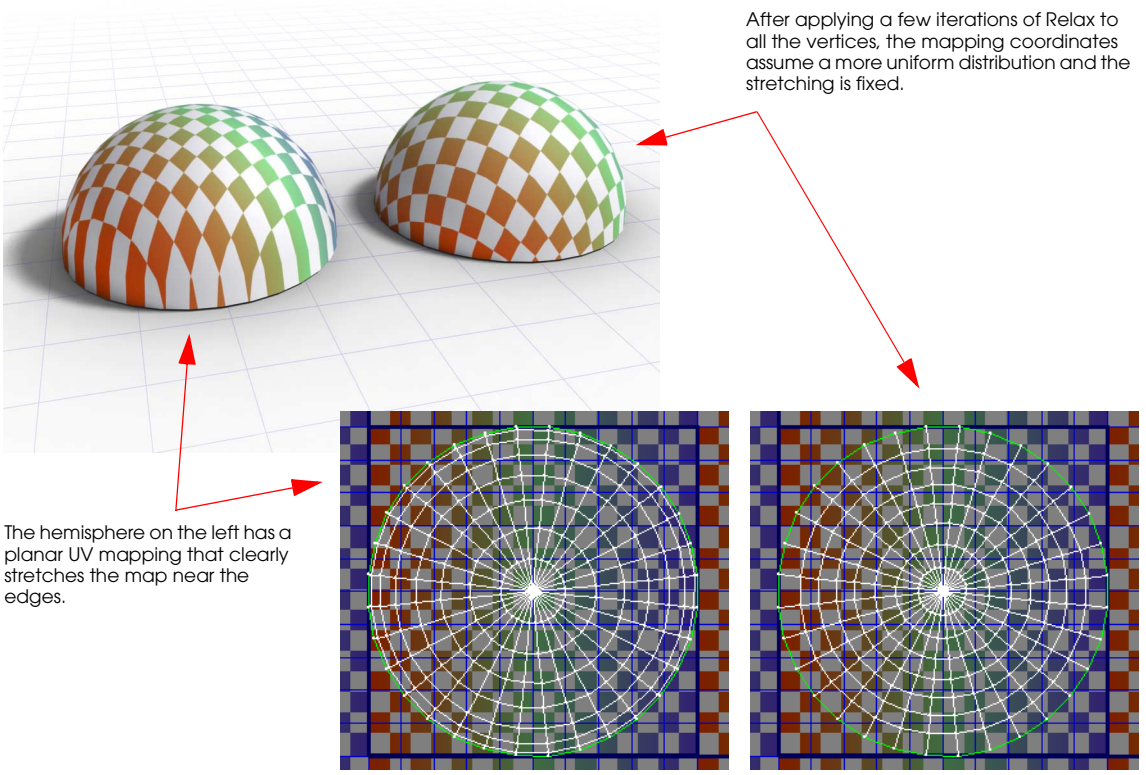

Relax works only on vertices. You can apply it to all vertices, or only to a selection. Depending on your model, Relax can be helpful for fixing specific areas by better redistributing the distances between vertices.

Best results are usually obtained with a mix of manual adjustment and applying Relax on isolated vertex selections.

Shell Modifier

This modifier add thickness to a surface. Surfaces with no thickness are often used to model hollow objects, like shells, where the interior is never seen. The Shell modifier extrudes the faces and gives a real thickness to the object if you need, for example, to show the object's surfaces as standalone objects in an exploded view. This modifier works on meshes. It will work on poly objects, patches, and NURBS surfaces, but converts them to meshes.

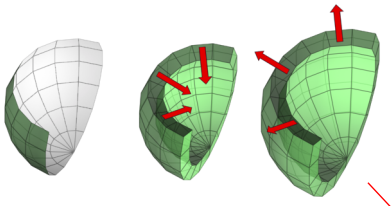

The inner and outer amounts extrude the surface along the face normals (outer) or opposite to them (inner).

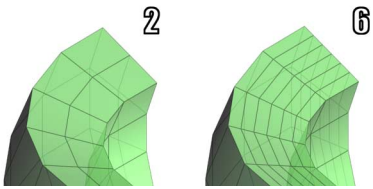

Segments is the number of extrusion intervals along the side faces

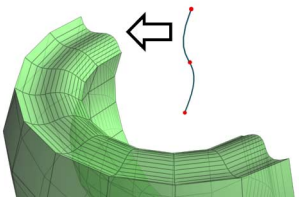

Bevel Edges lets you specify a spline to shape the extrusion.

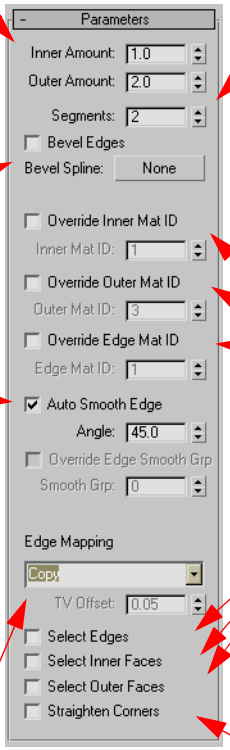

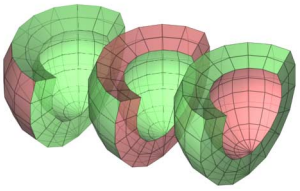

You can override the material IDs for the newly created surfaces.

Auto Smooth Edge uses automatic smoothing to smooth the side extruded faces.

The new surfaces can be selected as sub-objects and passed up the stack.

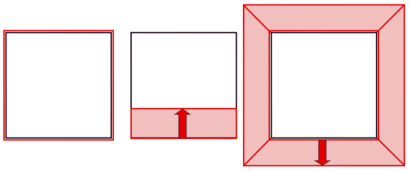

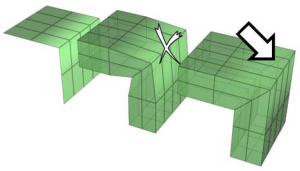

Edge Mapping applies predefined mapping techniques to the side faces: Copy, Strip, Interpolate, or None. Above is a representation of the UV mapping extensions applied to a rectangular surface by Copy, Strip, and Interpolate. TV Offset specifies the size of the mapping when Strip or Interpolate is used (arrows in the illustration).

Straighten Corners corrects the extrusion of corner vertices when their normals are averaged.

When using high thickness values, some vertices might overlap and further editing of the resulting mesh might become necessary.

Materials

4

Architectural Material

This new material type gives you quick and high-quality, real-world materials. It has a simplified user interface that works very well for architectural surfaces and the use of Radiosity or mental ray for global illumination effects.

This material is physically based and, with minimum settings, will provide highly realistic results in your architectural studies. It gives you automatic ray-traced reflections and applies the correct falloff for reflections and refractions.

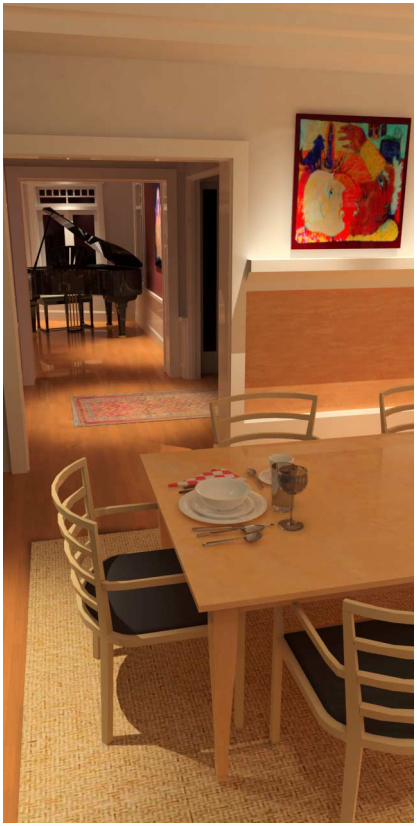

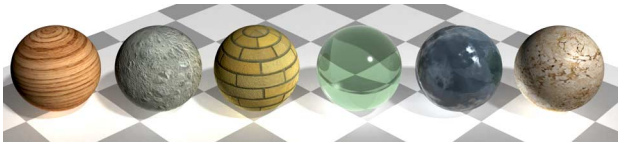

This sample scene uses only Architectural materials. Practically every real material can be replicated by this new material type. Templates give you the basic material settings for the most common ones, like wood, stone, masonry, glass, water, polished stone, and many more.

Materials can have any type of supported diffuse textures.

Transparency, refractions, and reflections are automatically handled, based on the template used as the basic material.

Note that reflections are stronger when objects are seen at an angle, as happens in reality for glass or varnish (Fresnel reflections).

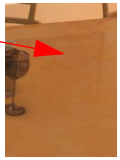

Radiosity and global illumination properties are automatically set, based on the material texture, color, and reflectivity.

mental ray® can be used with excellent results with this material, but it will ignore the Emit Energy setting and the Sampling Parameters, using its own sampling instead.

If you are familiar with **VIZ Render** and Autodesk **Architectural Desktop 2004**, you will recognize the Architectural material as the one used by **VIZ Render**. This material is used when you import DRF files.

User Interface

Templates can be selected from this drop-down list. The full list is on the left. Each template sets the basic material properties. You can then add textures and refine the settings to customize the material.

As for standard materials, you can use texture maps for many of the material parameters.

The diffuse color is specified here. If the material has a diffuse texture, you can automatically calculate an average color based on the texture by clicking the arrow button.

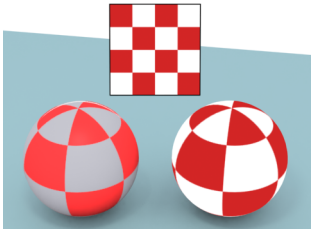

Raw Diffuse Texture excludes the material from illumination and exposure-control processing. The texture will match the original image or color. Useful, for example, to keep self-illuminated materials fully white when Logarithmic Exposure Control is used.

Index of Refraction, Transparency, Translucency, and Shininess are set by the template, but can be adjusted manually.

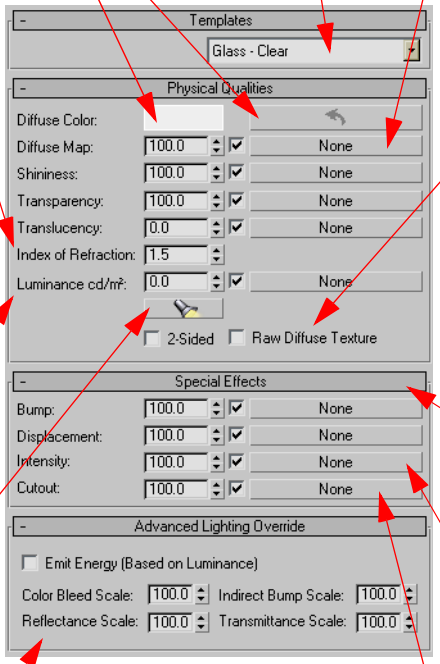

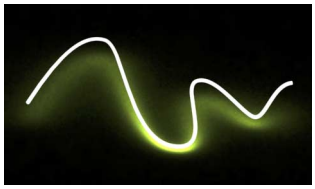

Set a Luminance value if you want the material to emit light.

Special Effects maps add the flexibility of standard material features.

Click this button, then select a photometric light in the scene to set the luminance at the same level as the light. This is a quick way to make an object emit as much light as a known light.

Intensity modulates the brightness of a material. Using a Noise map increases realism and avoids the flat computer look.

The settings in this rollout are similar to those found in the Advanced Lighting Override material, and are included in the base Architectural material for convenience.

Cutout is similar to standard opacity mapping, but where the material is transparent, no reflections, refractions, or shininess will be visible. It effectively "cuts" the material, it does not just make it transparent.

Multi/Sub-Object Material

Reference to multiple sub-object materials on meshes was sometimes confusing in earlier versions of the **3ds max** software, because of the use of numeric values. Objects with many material IDs, or objects created by attaching two meshes, both using multi sub-object materials, were very hard to figure out.

In this version, editable objects have the ability to select sub-objects assigned to specific materials by using names instead of numbers. The Material group in the appropriate sub-object modes has been expanded.

The base objects using this new feature are Editable Mesh, Editable Poly, Editable Spline, and Editable Patch, plus objects using Edit Mesh, Edit Spline, or Edit Patch modifiers.

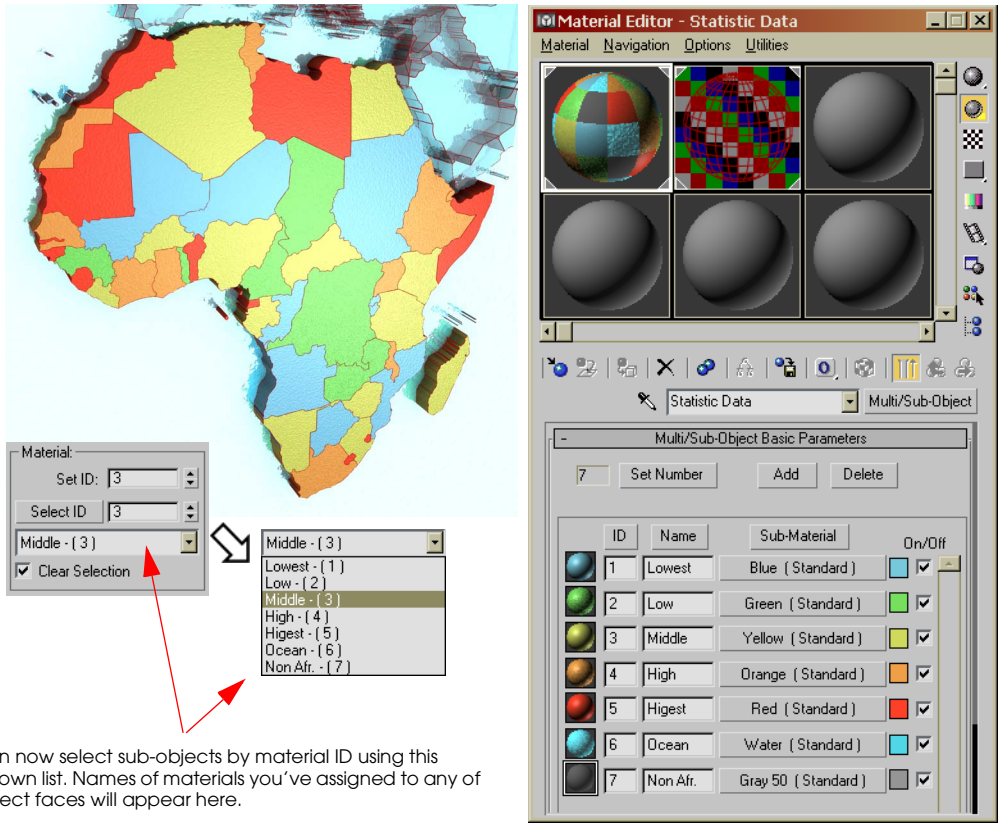

You can now select sub-objects by material ID using this drop-down list. Names of materials you've assigned to any of the object faces will appear here.

Note that if more objects share a single Edit modifier, or if the object does not have a Multi/Sub-Object Material assigned, the list will be unavailable.

New Materials Utilities

A couple of utilities have been included to help manage the complexity of sub-materials. They are accessible from the Utilities menu in the Material Editor, or from the Utility panel, in the list opened when you click the More button.

Clean MultiMaterial

Clean MultiMaterial finds all the materials in the scene that have unused sub-materials, presents them in a list, and lets you chose which ones need to be cleaned. Of the selected materials, the unused sub-materials will be removed. Materials IDs remain unchanged, and the used sub-materials are unaffected.

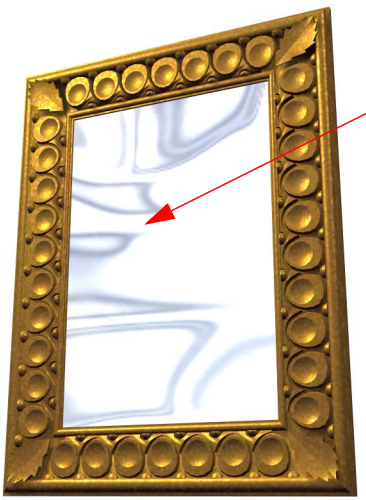

An object using only selected sub-materials from a Multi-Sub-Object material can be quickly and automatically cleaned, removing the unused sub-materials from the list.

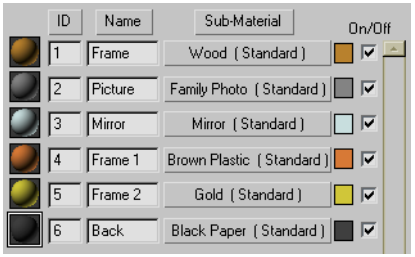

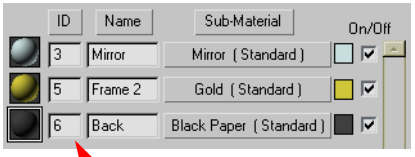

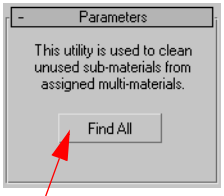

Click Find All to start the search for unused sub-materials.

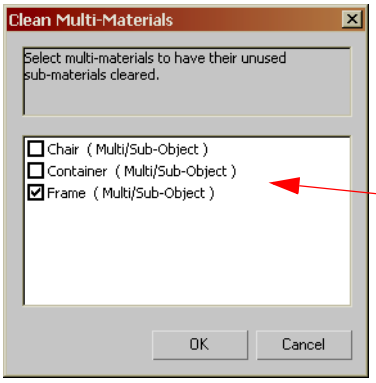

The utility lists all the materials in the scene that have unused sub-materials. Select the ones you want to be cleaned, and click OK. The remaining sub-materials will maintain their original IDs.

Instance Duplicate Maps

Instance Duplicate Maps searches all the materials for references to the same texture file. The duplicated textures are then listed; for each, the material and map where they are used is displayed. You can then choose to instance all of them, or select and instance only a subset.

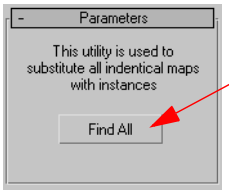

Click Find All to start the search for duplicated maps in all the materials used in the scene.

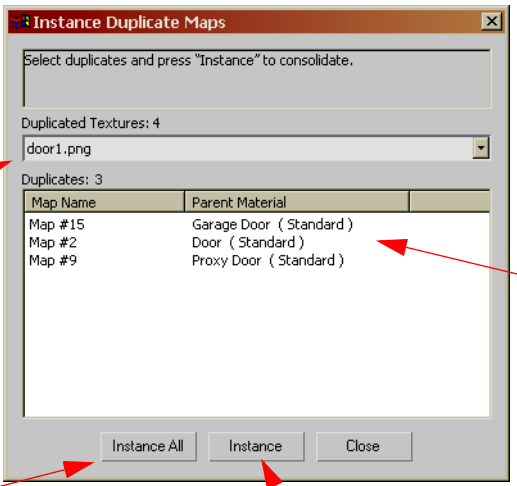

The list of duplicated maps appears as a drop-down list. Choose a map and the materials using it will be listed.

From the list of materials using the selected duplicated map, you can select the ones you want to instance.

Click Instance All to instance all the duplicated maps in the scene. The dialog will close after that, since no more duplicates exist at that point.

Click Instance to instance the selected map for the selected materials. The dialog stays open so you can select other materials and instance the map for them as well, or switch to a different map from the duplicated textures drop-down list.

reactor 2

Version 2 of **reactor** adds functionality and improves workflow, as well as speed and quality of the dynamic solutions. Here is a brief overview of the new functionality and the major changes from the previous version.

For details of the other minor technical improvements and additions, see the topic *What's New in reactor 2* in the online reference.

User Interface

As you saw in the *User Interface* topic earlier in this guide, **reactor** now has its own toolbar docked to the left side of the **3ds max** window, and a menu with all the objects, modifiers, and utilities.

Here is a look at the toolbar. It's a single column in the default docked location. You can undock it and make it a multi-line toolbar if you prefer.

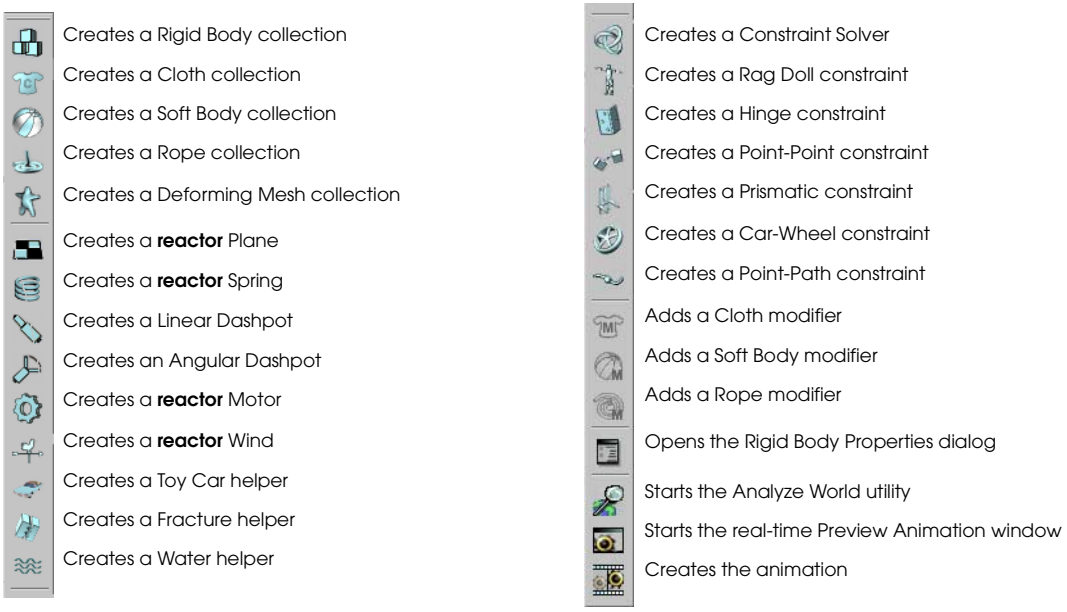

While access to the Utility panel is still necessary to set parameters and options, the **reactor** toolbar and menu provide a way to work faster once those values have been set and the simulation just needs to be tweaked.

The menu has the same options as the toolbar, plus some utilities like View Stored Collisions, and key reduction or deletion.

Simple and Cooperative Constraints

Constraints have been divided into two types: Simple and Cooperative. The first type uses constraints that are added to the scene and automatically solved. The second type works with a Constraint solver, where the constraints are listed and need to work together.

Cooperative constraints require more time to solve, but provide more robust and realistic results, due to their interaction.

Hinge Constraint

This new constraint works like a hinge between two objects, allowing them to rotate around one axis. A single object can also use Hinge, with the World space as its parent.

The large arrow indicates the hinge axis.

Note: If you enter sub-object mode, you can rotate the parent and child spaces. A set of red arrows indicates the child space, and will be visible when the two are oriented differently. This behavior is common to most constraints.

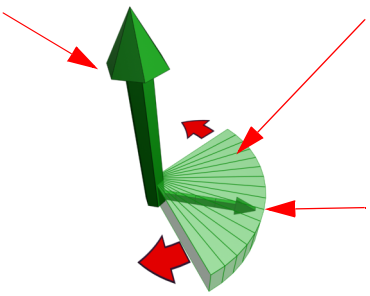

You can set an optional min and max range of the swing, visually indicated by this angle gizmo.

The small arrow indicates the origin (zero angle) for the limits.

Rag Doll Constraint

The Rag Doll constraint restricts the movement between two objects (or one object and the World) in a way similar to human body joints.

This is the main constraint plane, the one that defines the orientations of the other limits.

These two symmetrical cones define the plane limits.

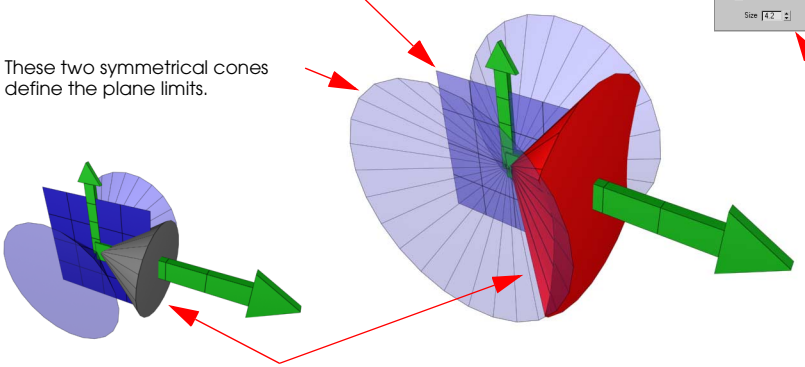

This gray cone defines the range of motion of the joint. The range is reduced by the plane limits, making the effective range the area highlighted in red.

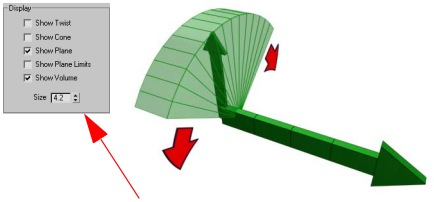

To avoid confusion, you can turn off parts of the gizmo or make it bigger. To set the Twist, for example, you should disable the rest and enable just the twist gizmo as shown above.

Arrows always indicate the constraint axis and limits origin.

Using Rag Doll constraints or a combination of Rag Doll and Hinge constraints, you can simulate the behavior of a lifeless human body (hence the name Rag Doll), and create digital stunts like falling and rolling down stairs.

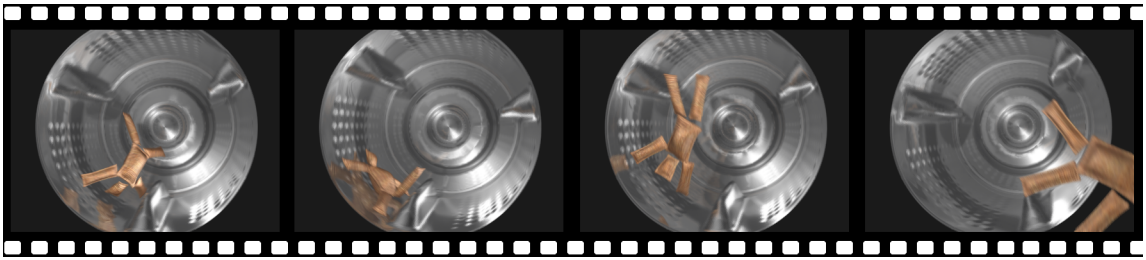

A simple rag doll tumble-drying.

The way to use a rag doll in your work is to prepare a substitute model of your original character in which the bones have been replaced by a matching Rag Doll rig.
Animate your original character up to the point where the Rag Doll model is needed. Match the position of the character, and then hide your character and let the Rag Doll simulation take over.

You can create your own rag doll by manually setting Hinge or Rag Doll constraints for each body joint. A script is provided to create standard humanoids, and to automatically add the appropriate constraints.

The script is called *rctragdollscript.ms* and is located in the *scripts* directory. Run it, and the Ragdoll dialog appears.

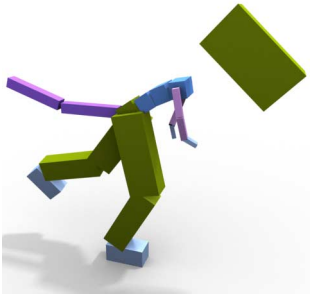

Customize the humanoid to best fit your character by scaling the components. When ready, click Constrain Humanoid.

Other custom body parts can be added if necessary, but they will need to be manually constrained.

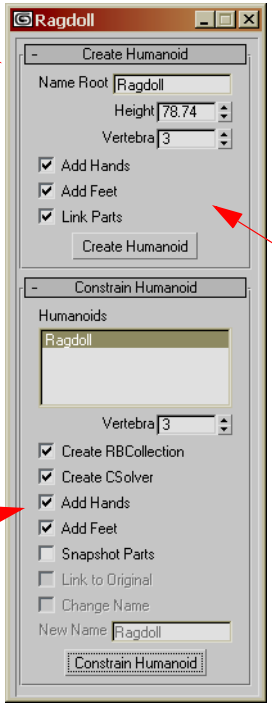

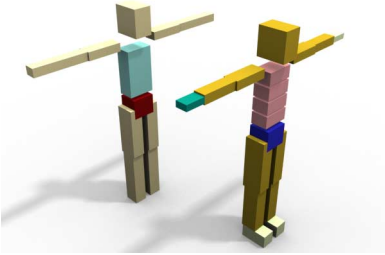

Set name, size, and body parts options, then click Create Humanoid to get a basic reference character.

Add a **reactor** plane and try a simulation. Your character is ready to use. Add some animation if your character has to inherit the original character's movement.

Car-Wheel Constraint

This constraint simulates an individual car wheel. You can use as many as you need for your vehicle and use the same parent for all.

A Car-Wheel constraint has two main axes, one for the spin direction and the other for the suspension.

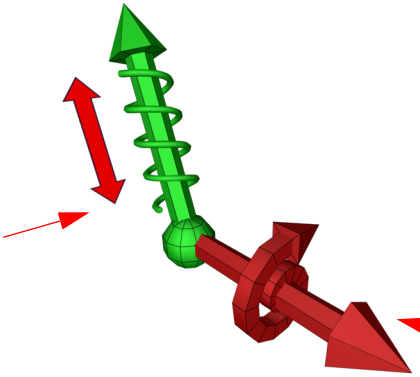

The Suspension axis (in green, parent space) indicates the direction suspensions are working. You can set optional min and max limits to control how much the suspension can slide below and above the constraint attachment.

The Spin axis (in red, child space), should be aligned with the wheel axis.

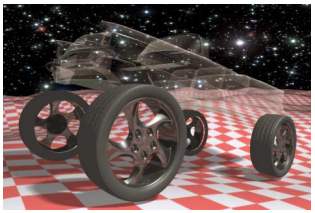

The test above uses four constraints to spin the wheels and push the car on a surface.

Car-Wheel constraints can be used to follow the vehicle's animated, unyielding body and to turn the wheels, or they can apply a force to the parent rigid body and push the vehicle forward.

View Stored Collisions

You can store a record of each collision that **reactor** detects during a simulation. You can then view the list and save it to a text file.

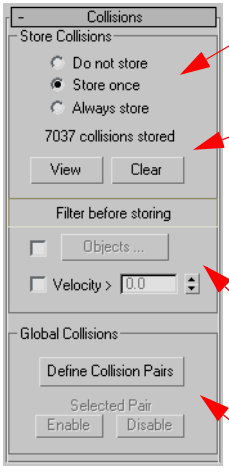

You can store the collisions only once, the next time you create an animation, or every time. By default, no collisions are stored.

Click View to open the dialog shown at the right, listing all the collisions recorded. From this dialog, choose Save to save the collisions as a text file.
Click Clear to delete the records.

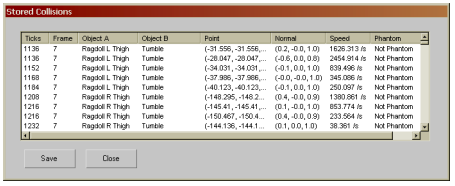

Enable the collision filter and click Objects to open a dialog where you can choose which object collisions to store.
Enable Velocity if you want to store only collisions above a specific speed.

Collisions are now managed globally rather than in a Rigid Body collection. Collision data is now accessible using MAXScript.

Other Improvements

Prismatic Constraint
A Prismatic constraint limits the movement between two objects along a single axis. Only sliding is allowed, not rotation. An example of its use would be a forklift.

Fractures
Fractures has been improved; all fragments now are considered rigid bodies and can be constrained and given an initial velocity. **reactor** can now handle complex fractures with more reliability.

Animatable Wind
Wind parameters like direction and strength can now be animated. Falloff and range controls have been added.

More Information

The workflow and many concepts behind **reactor** functionality have been revised. It's beyond the scope of this guide to reintroduce the basic concepts of **reactor**.

The **reactor** portions of the online reference and tutorials have been rewritten, and we suggest that even more experienced users check them out to make sure all the differences are clear.

Dope Sheet: Curve Editor

Operations in Dope Sheet have been speeded up, optimizing the key computation and using internal caching. Only the keys that are actually visible in the working area are updated.

By carefully limiting the tracks displayed to the ones you really need, you can improve responsiveness and perform tasks much more quickly.

Now, Auto-Expand is by default set to show just the selected objects, and not to expand any tracks or children.

Keyable Tracks

Previously, you could change the keyable attribute of a track only by clicking the keyable icon. Now, it's possible to select multiple tracks, and set them to keyable or not with a single step.

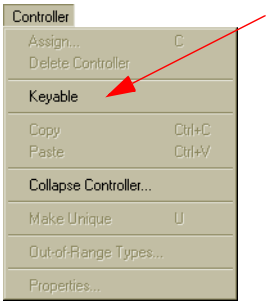

Use the new Keyable menu command to toggle the Keyable flag of the selected controllers.

In case of a multiple selection, the toggle sets or resets the Keyable flag based on the first track's status.
After choosing the menu command, all the selected flags are unified.

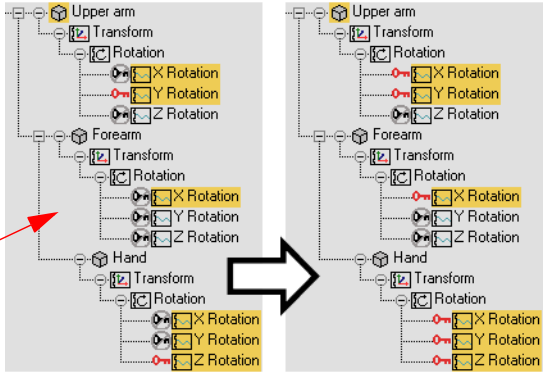

Also, a new filter has been added to the Filters dialog. The Keyable Tracks filter displays only the tracks that use keys. This helps minimize the number of tracks shown, and leaves only the ones that actually have keys or curves to edit.

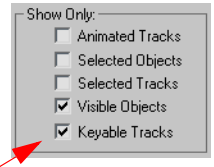

Turn on Keyable Tracks in the Filters dialog to show only tracks that allow keys.

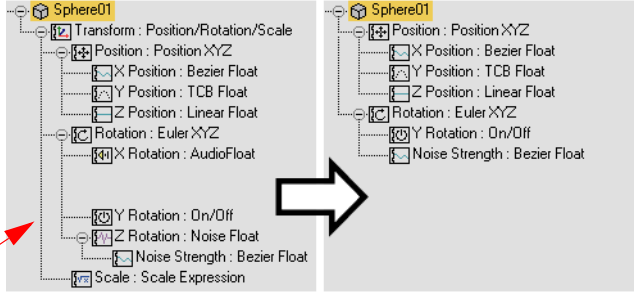

In this example, various controllers have been assigned, and when the Keyable Tracks filter is enabled, tracks with controllers like Expression or AudioFloat are hidden.

mental ray

mental ray is a powerful rendering system fully integrated within **3ds max**. It was previously available as an optional package, but this version makes it available as a standard feature.

The image below shows many mental ray features.

Global Illumination

Caustics light effects

Full support for all **3ds max** materials

Reflections and refractions

Realistic shaders

Fast raytracing

Area lights

Depth of field

Distributed rendering

mental ray uses a scene-description language. Its integration within **3ds max** is based on a translator that translates the scene data to mental ray prior to a rendering.

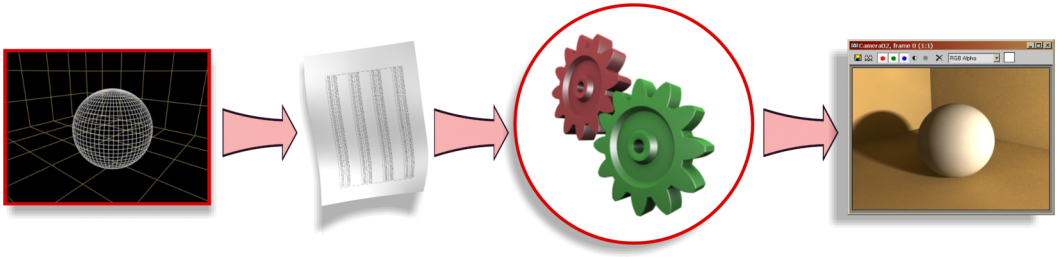

With mental ray selected as the active renderer, render the scene.

In the background, the scene is translated and transferred to mental ray.

The mental ray renderer re-creates the scene and processes it.

The scene is rendered in the frame window or to a file, as usual.

All this processing is done internally, and to the user, a mental ray rendering looks like a standard rendering.

Integration

mental ray integration with the **3ds max** user interface has been made as seamless as possible. The translator uses the standard **3ds max** materials and lighting parameters, and some extra parameters that pertain only to mental ray.

Controls specific to mental ray appear in the following locations:

- A panel in the Object Properties dialog
- A panel in the Preferences dialog
- New material types
- New map types
- New light types (mr Area Omni and mr Area Spot)
- New shadow types
- New properties for lights and materials (when mental ray extensions are enabled)
- A whole new set of panels and rollouts on the Render Scene dialog

Complexity

mental ray uses both scanline and ray-tracing algorithms. The true power of mental ray is in its fast and extensive use of ray tracing. Ray tracing is used to calculate not only reflections, but also diffuse and caustic lighting. This power doesn't come without some complexity and numerous parameters to juggle.

If you are new to mental ray, you might find its complexity a bit intimidating at first. The integration with **3ds max** allows you to use mental ray as the rendering engine, while continuing to work with lights and materials as you have in the past. mental ray renderings will match those of the existing scanline renderer.

In the Render Scene dialog, you can add global illumination and caustics. Using Preferences to enable the mental ray extensions allows you to add specific mental ray features to standard **3ds max** materials. As your understanding grows, the more you can add features to your scenes.

If you are an expert, you will find options to access mental ray features directly, like the mental ray material, contour shaders, and camera shaders. Later in this topic, you will see how to access these options.

This introduction gives you an overview of mental ray and the translator. We suggest you read more in the online reference, the tutorials, and the third-party mental ray documentation.

Shaders

The term *shader* is used quite often when talking about mental ray. Shaders are the basic elements that control a scene rendering. A shader in its most simple form can be a surface shader, like a procedural map. But mental ray shaders also control volumetric effects, photons, lens effects, file formats, and much more.

Combinations of shaders that define all aspects of a material, are called *phenomena*®.

Using mental ray

Let's first take a look at the changes implemented throughout the user interface to integrate mental ray. We will then introduce concepts and features, and then look into specific mental ray user-interface and rendering options.

We start with materials and maps, which play a major role in mental ray renderings. You need to have mental ray set as your active renderer to access the features illustrated. To see how to switch renderers, see the *Rendering Dialogs* topic.

To better handle differences between the scanline and mental ray renderers, the Material/Map Browser displays extra materials (phenomena) and maps (shaders) specific to mental ray using yellow icons. Incompatible materials and maps for the active renderer are grayed out.

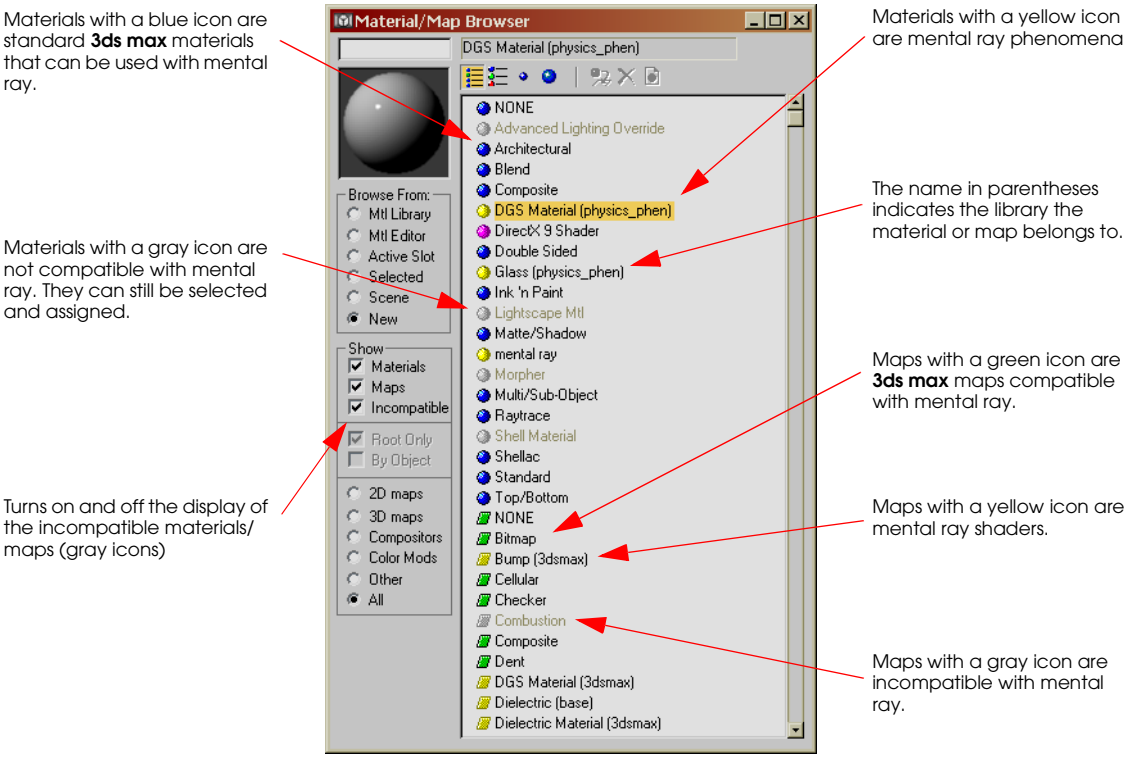

mental ray shaders are very specific and technical, and to a first-time user they might cause some confusion. mental ray is a powerful rendering system that requires a certain degree of technical knowledge to use fully.

The translator can use standard **3ds max** materials and add specific mental ray functionality (like photon shaders) to them. This allows you to use the powerful features of mental ray while continuing to use the materials you are accustomed to.

Colors

Another difference you will notice is the Color Selector.

Whenever a color has to be chosen for a mental ray feature, the Color Selector dialog has values in the range 0.0 to 1.0 instead of the 0 to 255 used for all other **3ds max** colors.

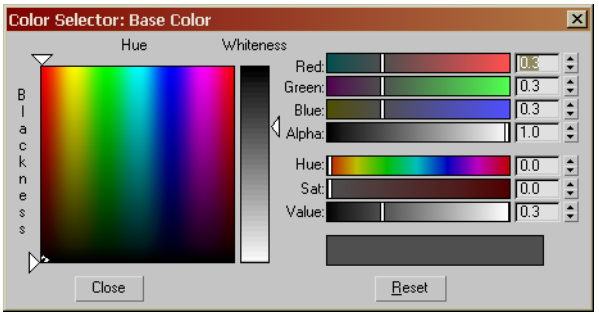

Lights and Shadows

Two new lights have been introduced to match mental ray area lights: mr Area Omni and mr Area Spot.

Standard and photometric lights can still be used and will render correctly.

Skylight and IES Sky are supported when Final Gather is enabled in the Indirect Illumination options.

The best choices for shadows are Ray Traced Shadows, which use the mental ray raytracer, or "mental ray Shadow Map". Advanced Ray Traced and Area Shadows are not supported and will be rendered as raytraced shadows. Ordinary Shadow Map shadows are rendered, but "mental ray Shadow Map" provides better results.

Errors and Warnings

When translating a scene to mental ray, errors and warnings on items that are incorrectly set up or unsupported are displayed by the translator in a window that only opens if any issue is found.

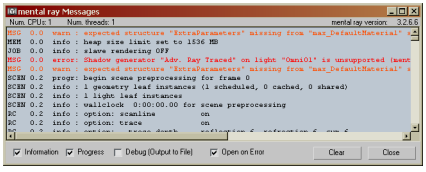

Error and warning messages appear if something requires your attention. Optionally, other status messages can be displayed, too (in black).

Rendering

When it comes to rendering, mental ray presents you with an extensive set of options.

Before we look at them, it's better to introduce functionality and concepts that will help you understand the meaning and purpose of those settings.

Indirect Illumination

Indirect illumination is one of the most powerful features of mental ray, handling both diffuse light and caustic effects.

Later in this topic, we will explain the options and parameters, but first, we want to explain some concepts about how mental ray indirect illumination works.

Photons Basics

mental ray uses *photons*, samples emitted from the light sources, to simulate indirect illumination. Each photon carries energy and interacts with the scene, hitting surfaces, bouncing, and being reflected or refracted. Energy decays over distance, and whenever photons hit a surface.

Photon interaction with the scene is handled by special *photon shaders*. While photon shaders can be added and tweaked manually, it's better for a novice mental ray user to let the translator do it.

For example, a Raytrace material automatically handles photons to create reflected and refracted caustic light effects.

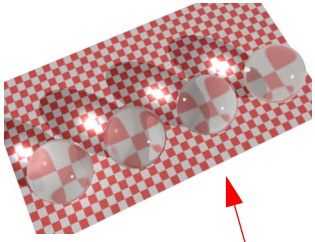

When rendered with mental ray, all materials give the same result. All the standard **3ds max** materials handle diffuse and caustic light.

This scene is using different materials to achieve the same effect on each sphere. From left to right, the materials on each sphere are: standard material with raytraced reflection and refraction maps, raytrace material, architectural material, mental ray glass phenomenon.

When rendered with the scanline renderer, the **3ds max** materials are still working. The mental ray material is not rendered. This makes standard materials a more flexible choice in many cases.

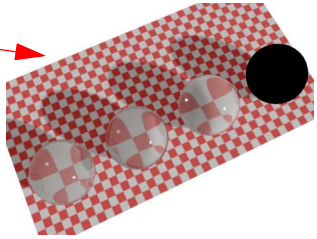

One suggestion is to use standard **3ds max** materials whenever possible. Your scenes will have less problems if you need to render with the scanline renderer or other third-party renderers.

There are two kind of photons: Global Illumination and Caustic. The first simulates diffuse light; the second, reflected and refracted light.

It's important to have your system units set to realistic units for the scene. Light energy from photons depends on the scene size.

Global Illumination Photons

Global illumination (GI) simulates lights bouncing off surfaces in a scene, creating diffuse or ambient light. **3ds max** already has two systems to simulate it: radiosity and the Light Tracer. mental ray adds a third one, based on a different technology.

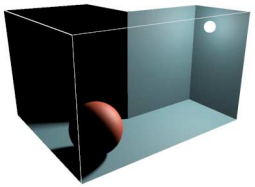

This is the scene with only direct light, no global illumination.

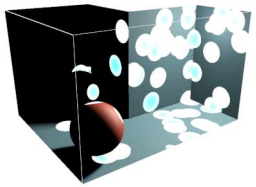

With only a few photons emitted, each photon carries a lot of the light energy. By default, the size of the photons is calculated based on the scene extents.

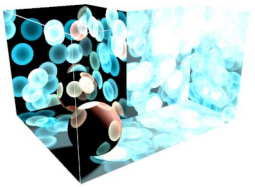

As more photons are emitted, they blend and bounce, assuming the color of the last surface they hit.

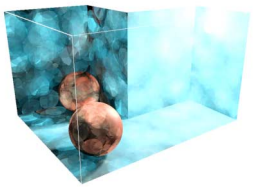

A higher number of photons shows that the energy that each photon carries is now smaller. The blending of the photons starts creating the GI effect.

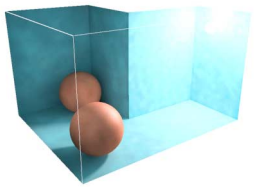

With an even higher number of photons, the GI effect is more clear. Still, some artifacts due to photon granularity are visible.

The correct number of photons depends on the scene complexity. Large, flat surfaces will always require high amounts of photons. It's not unusual to use millions of photons.

When photons alone can't give you the desired result, you can use Final Gather to refine the image.

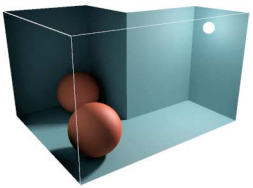

Final Gather alone, with no photons, already provides a bit of light in the shadow areas, but that light is the result of sampling only the direct light in the scene.

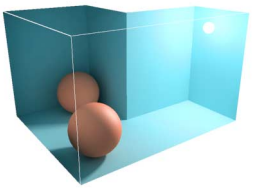

When used after photons have been emitted, Final Gather can smooth the result without the need to use extremely high numbers of photons.

Final Gather for mental ray works the same as Regather in radiosity, by sampling the scene for each rendered pixel, or averaging between more sparse samples.

The best results are often achieved with a medium-to-high number of photons, and Final Gather enabled for the final rendering. Final Gather might affect the lightness of the scene. Make sure you do some test renderings using Final Gather in Preview mode to get quick feedback on the overall scene lighting.

Caustic Photons

Caustics simulate refracted and reflected light. An example of caustics is the beam focused by a magnifier lens, or the play of lights at the bottom of a pool. Caustics is a new effect mental ray provides that was not previously available in **3ds max**.

With no caustic photons, the water surface is just a transparent object lit by the direct light.

A few caustic photons show that the light is now calculated using samples, and each sample goes through the surface (refraction) and bounces off it (reflection).

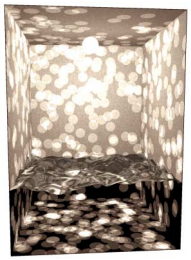

More photons gradually show the caustic effect on the overall scene.

At this stage, photons starts combining into contiguous caustic effects. This stage can be good for test images, since the details let you preview the final look.

For final renderings, increase the photon count even more, and achieve the details that make the effect realistic.

The working principle of caustic photons is the same as GI photons, but these are specialized for reflections and refraction. The photons are smaller, and create more detailed effects.

Final Gather does not improve caustic effects. Use a higher number of caustic photons and adjust the samples to increase the details.

Other mental ray Features

Contour Shaders

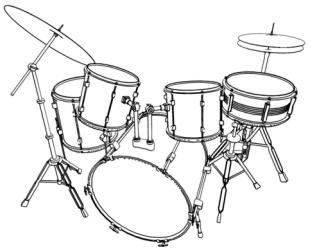

Contour shaders provide a very fast way to get the effect of ink lines, as in a cartoon. Compared to the Ink 'n Paint material, mental ray contours are extremely fast and accurate. They do not, however, completely replace the Ink 'n Paint features.

Contour shaders are applied post rendering, based on information gathered during the rendering. A special, *Contour Store* shader takes care of recording this information.

To use contour shaders, you have to assign a contour shader to the material. To do this, you have to enable the mental ray extensions. Also, you have to enable the contour scene shaders on the Render Scene dialog. You are shown how to do make these changes in the following pages.

Camera Shaders

Camera shaders affect the rendering as a whole. You can add a global mist effect, or lens effects like Barrel, Pincushion, or WrapAround.

Example of mist effect.

Example of WrapAround output. WrapAround creates a 360-degree view of the scene, suitable for spherical mapping.

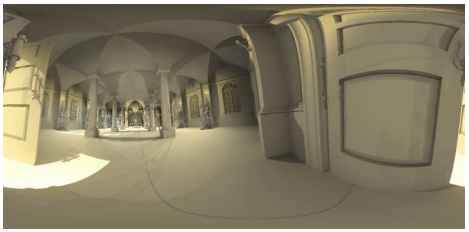

Depth of field

Depth of field simulates the blurring of objects in front of and behind the camera focus plane. The amount of blurring and the focus plane are set for the camera using an f-stop value and the target distance.

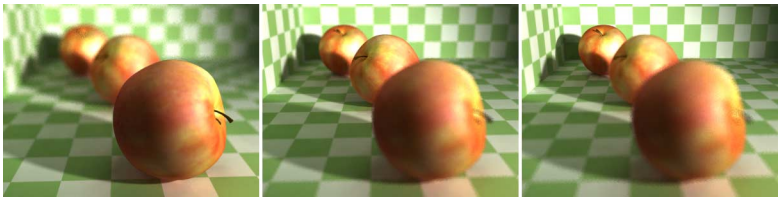

Focus plane set at different distances, from the closest to the farthest apple.

Enable the "Depth of Field (mental ray)" multi-pass effect for the camera, not the standard **3ds max** Depth Of Field effect.

Displacement

mental ray displacement is very efficient and detailed. There are differences between the standard material displacement maps and mental ray displacement. Parameters set to work for one renderer will not produce the correct result when rendering with the other.

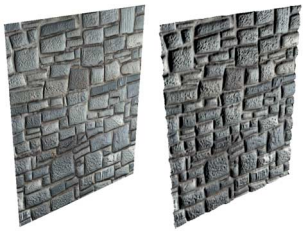

Even a good bump map is no match for real displacement.

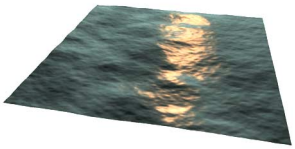

Displacement using a procedural map; in this case, the Ocean shader from the Lume library.

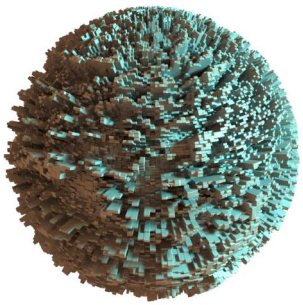

Displace can create some crazy stuff.

3ds max displacement works on patches, poly, mesh, and NURBS objects, but all other objects (like parametric objects) require a Displacement Approximation modifier. mental ray displacement works on all objects.

Motion Blur

Motion blur in mental ray is based on the Object Motion Blur technique. Object Motion Blur has to be enabled on the Object Properties dialog for all objects you want blurred.

The amount sampling is then done by mental ray and has to be specified on the Renderer panel of the Render Scene dialog.

Unlike the scanline renderer, mental ray can blur shadows as well as objects.

Accessing mental ray Features

So far, we have seen what mental ray can do and how it is integrated into **3ds max**. The concepts we have introduced will now become more clear as we start looking into the options of the rendering panels on the Render Scene dialog.

Make sure you have mental ray set as your active renderer. For more information, see the *Rendering Dialogs* topic in this guide.

Renderer Panel

The Renderer panel is the first of three panels in the Render Scene dialog to have mental-ray-specific settings. The first two rollouts are explained on this page, the remaining two on the next.

Edge antialiasing (as well as the quality of practically any other effect) is controlled by the sub-pixel sampling set here. Depending on the contrast changes in the image, the sampling goes from the minimum (1/4 means a single sample every four pixels), to the maximum when contrast changes are more marked (4 means four samples per pixel). You can use a maximum setting of 1 to speed up test renderings, and then increase it to 16 or more for really high-quality renderings.

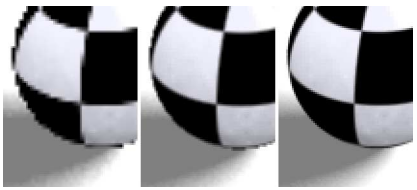

Example of Maximum set to 1/4, 1, and 16

This group specifies the contrast threshold for the subdivision of samples.

Lock Samples forces the same pseudo-random sample pattern to be used for each frame. If turned off, the samples are different at every rendering. Jitter adds a variation to the sampling to help reduce aliasing effects.

mental ray can use different rendering algorithms. Depending on the requirements of the scene, one or all are used. Disabling one will disable rendering effects that depend on it.

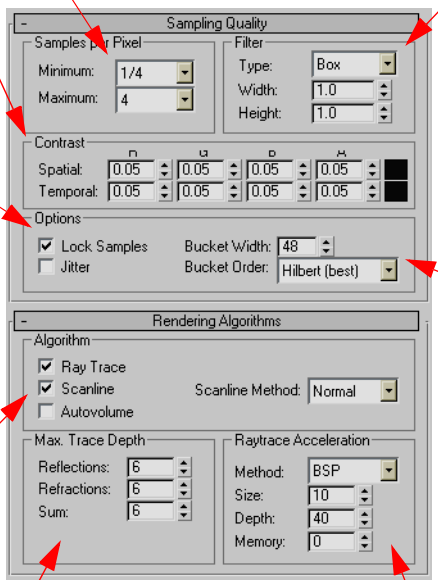

Filter combines multiple samples into one pixel. There are various filter types, and for each type you can specify the area to evaluate.

A *bucket* is a small portion of the image used as a rendering unit. It will become obvious as soon as you look at a rendering in progress what buckets are and how they tile up in the image. Bucket Width sets the size of the buckets. Bigger buckets take longer to render, but there are fewer of them in the image. The Bucket Order settings specify the sequence in which they are rendered.

These settings limit the number of ray reflections and refractions. The lower the values, the faster, but less accurate, the rendering. The Sum value is a limit for both reflections and refractions added together.

Different algorithms for raytracing acceleration are available. Settings depend on the scene size and the effects used. Default settings will be satisfactory for most scenes. Refer to the online reference for more information.

Renderer Panel (continued)

This group enables depth of field when you render a Perspective viewport. If you render a Camera view, these settings are set on the camera's Parameters rollout, and this group is ignored.

Focus Plane is the distance at which the objects are rendered sharply. The depth of field can be set using f-Stop values, as for a camera lens, or by Near and Far distances if you choose the In Focus Limits option.

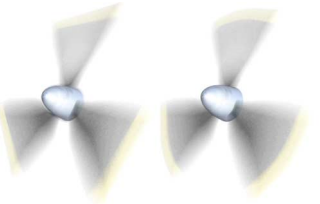

After you have assigned contour shaders to a material (see *Enabling mental ray Extensions* later in this topic), you need to enable the global settings here. The shaders in this section use contrast to determine where edges are. The Store shader records the information during the rendering pass. The Contour Output shader combines the rendering with the contours calculated post-rendering.

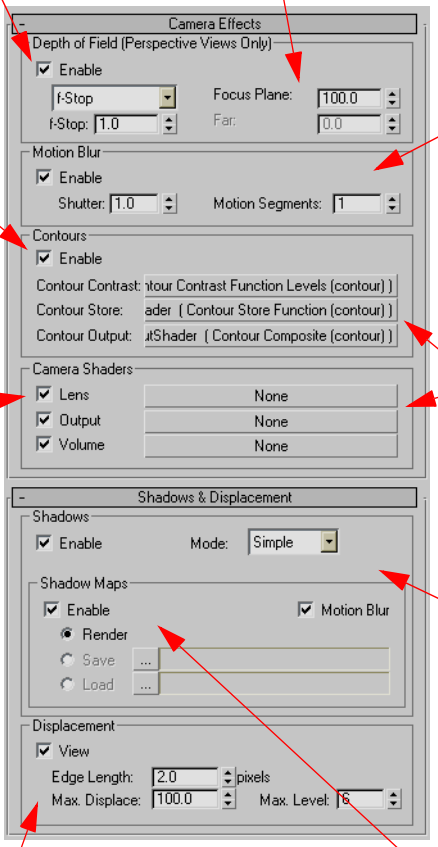

Here, you can enable motion blur. You can also set the shutter speed (which affects the amount of blurring), and the number of segments, if you need to correct straight blurring as illustrated above.

Camera shaders are assigned here. Lens shaders control distortion and other corrections during rendering. Output shaders instead apply changes to the rendered image. Volume shaders add volumetric effects to the image, like glows, mist, and underwater effects.

Parameters for Contours and Camera shaders assigned on the Render Scene dialog are accessible by dragging them onto a Material Editor sample slot.

Shadow settings let you enable shadows globally, and also select the order of shadow processing. For example, when using volumetric effects, the Segments mode is more appropriate,

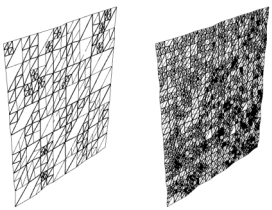

Shadow maps can be set to support motion blur; they can also be saved and loaded from a file instead of reprocessed at each use. This option is useful for complex scenes or animations where shadows do not change.

Here you can fine tune displacement mapping. View enables view-dependent displacement. Edge Length sets the minimum allowed subdivision size as illustrated above (in pixels if View is enabled, in scene units otherwise). Max Displace sets the global maximum displacement amount. Max Level controls the maximum mesh subdivision.

Indirect Illumination Panel

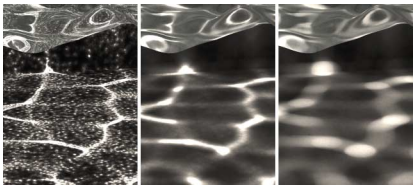

Caustics effects are enabled here. The Samples setting, as illustrated above, determines the level to which photons are blended together. Fewer samples give crispier edges, more samples give softer results.

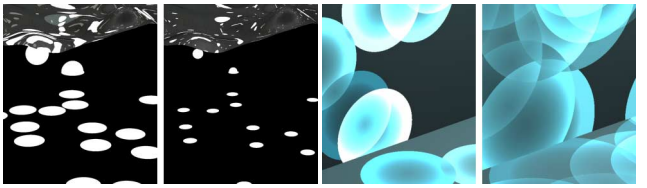

The radius of caustic photons is set based on the scene size (by default, 1/100). To override it, turn on Radius and specify a new value.

As with caustics, the default GI photon size is based on the scene extents (1/10), and can be manually adjusted.

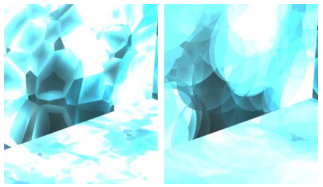

Enable Global Illumination here. The Photons value is similar to Samples for caustics; higher values increase the blending between photons, as illustrated above.

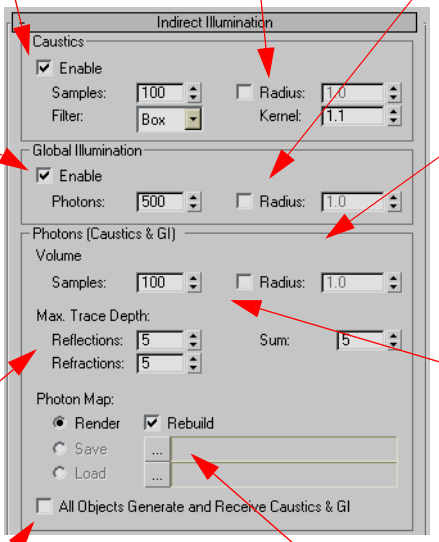

The settings in this group are common to both photon types.

Max Trace Depth sets the limits of reflection and refractions specifically for photons.

Volume sets the samples and size of photons for volumetric photon shaders assigned to materials in the scene.

Here, you can set all the objects to generate and receive caustics, instead of enabling this for individual objects on the Object Properties dialog.
By default, new objects are set *not* to generate caustics.

Photon maps can be saved and reloaded to save time on complex scenes.
Specify a file name to enable Save or Load.

Indirect Illumination Panel (continued)

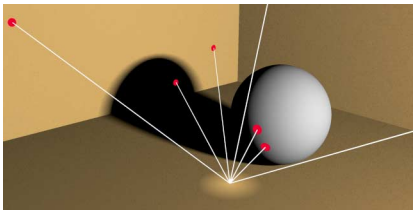

When enabled, Final Gather samples the scene from specific computed points of the image. These points are more or less dense, depending on the details in that area.
For each point, rays are cast and the lighting in the scene is examined and added to the point in question.
See *Global Illumination Photons*, earlier in this topic.

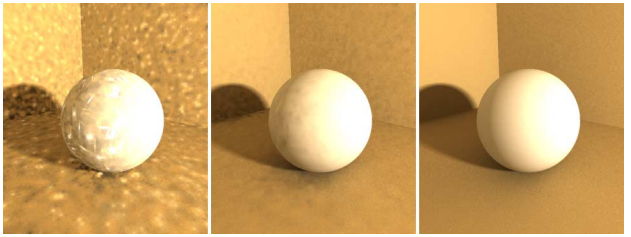

Samples specifies how many rays are cast into the scene to sample lighting. More rays produce more smooth and accurate results, but increase rendering time.

These settings control the quality of the sampling points.
Use Preview for quick test renderings to have an idea of how Final Gather affects the scene.

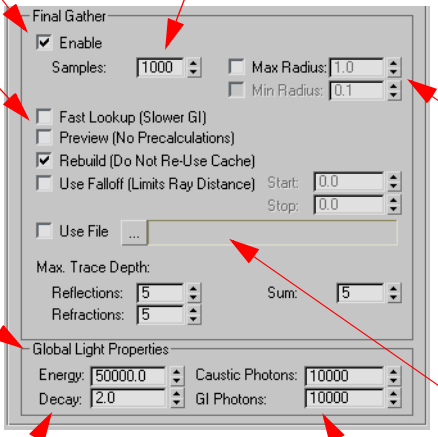

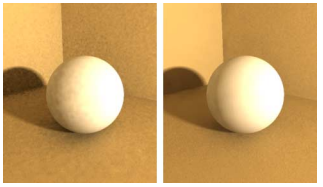

Max Radius and Min Radius set the distance for the blending of the Final Gather results for each computed point. The larger the size, the smoother but less precise the result.
You have to find the best combination of size and number of samples for your scene.

All lights in the scene are, by default, set to use these global properties.
This lets you lower the number of photons globally to speed up test renderings, and then raise it again for the final image.
See the following pages for how to override these values by enabling the mental ray extensions.

You can save the sampling information to a file and reuse it to speed up rendering if the scene does not change.
Turn off Rebuild to load the file; turn it on to save the file.

Energy sets the intensity of the photons, not of the light itself, as illustrated on the right. A Decay of 2.0 means that the energy fades with the square of the distance.

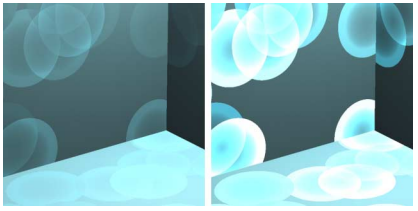

Here, you can specify the number of Caustic and Global Illumination photons each light will emit. For more information and sample illustrations, see the photon descriptions earlier in this topic.

Processing Panel

These options affect the translator and its interaction with mental ray. They are not options specific to mental ray.

Conserve Memory and Memory Limit instruct the translator to discard geometry and textures and keep the memory usage under the amount specified.

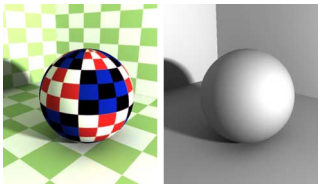

Placeholder objects save memory by delaying the sending of the geometry to the renderer until it's necessary. This is useful when most of the scene is outside of the field of view.

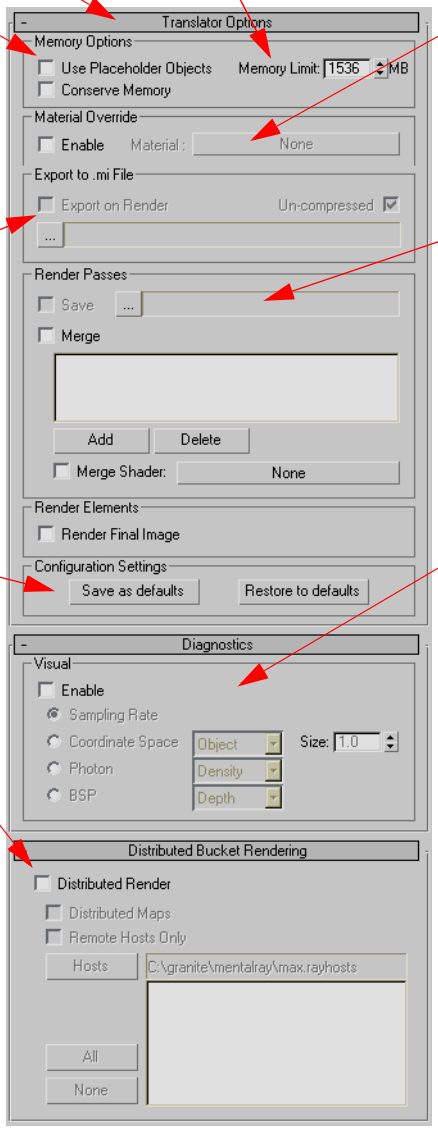

Material Override forces all the geometry to use the material specified. This is a useful option, for example, to check the scene's lighting by forcing a light gray material with no textures.

Set a file name and enable Export To .mi File to save a copy of the mental ray scene-description file created by the translator.
Expert users might modify this file and submit it manually to the mental ray renderer.

Render Passes is an advanced feature. Files containing special rendering information are saved using a proprietary file format, and can be merged with other renderings. For more information on this feature, refer to the online reference and the mental ray documentation.

Save the Translator options as defaults, or restore them to the last default settings that were saved.

Diagnostic options overlay information to the rendered image. You can check information about coordinates, sampling, photons, and ray-tracing optimization.
Below, examples of local coordinate space and sampling.

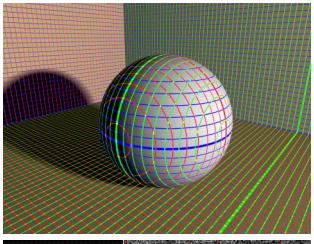

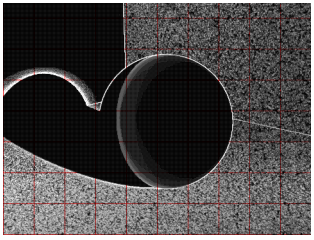

if you have multiple mental ray licenses, you can distribute your renderings.
Buckets are assigned for rendering to different servers. Here is where the bucket size set on the Renderer panel becomes important.
For more information on how to set up distributed rendering and the mental ray installation requirements for the rendering servers, refer to the online reference.

Enabling mental ray Extensions

Enabling the mental ray extensions allows you greater control over light and material options. Extra rollouts appear on the Modify panel and the Material Editor.

To enable the extensions, open the mental ray panel in the Preferences dialog.

Lights

On the Renderer panel, we have seen the global settings for lights. By default, all lights use these settings. Enabling mental ray extensions gives you the ability to override the global settings and fine-tune each light in a relative or absolute way. The rollout is available on the Modify panel when a light is selected.

By default, lights are created with Use Global Settings turned on. This way, they will get their energy, decay, and photon count from the Global Light Properties settings on the Indirect Illumination panel.

Leaving Use Global Settings on, you can change the light settings in a relative way.
Turning it off lets you set the light properties as absolute values.

Light shaders are not provided in this version of **3ds max**. If you own third-party shaders, you can use them here to modify the light's behavior.

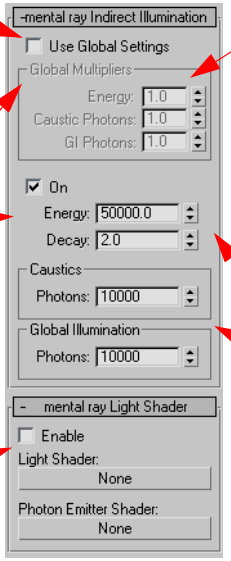

Using multipliers, you can change the light settings in relation to the global settings. This is useful, for example, in order to maintain the ability to change the overall light energy in a scene (or to speed up rendering by lowering the photons emitted for all lights), and still have the ability to adjust individual light settings.

Setting absolute values lets you fine-tune each light and its influence in the global illumination of the scene.
Some lights might be more important to caustic effects in a scene, while others might be marginal. Here, you can lower or increase the appropriate photon count, energy, and decay.

Materials

The translator has preprogrammed ways to translate materials and maps into mental ray shaders.

Usually, these defaults are all you need to create the materials you want and to maintain compatibility with the scanline renderer, and at the same time, have access to all the mental ray features.

When you want to override these defaults, you can enable the mental ray extensions, and an extra rollout on the Material Editor appears. This rollout allows you to substitute or add shaders.

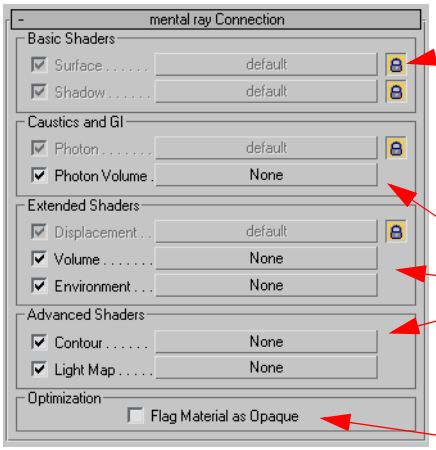

The mental ray Connection rollout shows you the shaders that can be assigned to the mental ray material.
The locked ones are those the translator assigns from the standard material maps.
Unlock them to assign a different shader.

Like any other **3ds max** map, shaders can be enabled or disabled.

Other shaders are specific to mental ray and are not used by the translator. These shaders enable special features like Contours and Volume effects.

Flagging a material as opaque, optimizes the rendering by letting mental ray know that the material does not have any transparency.

The mental ray Material

For expert users who have a deeper knowledge of mental ray, a special material is available that allows direct shader assignments to the material properties.

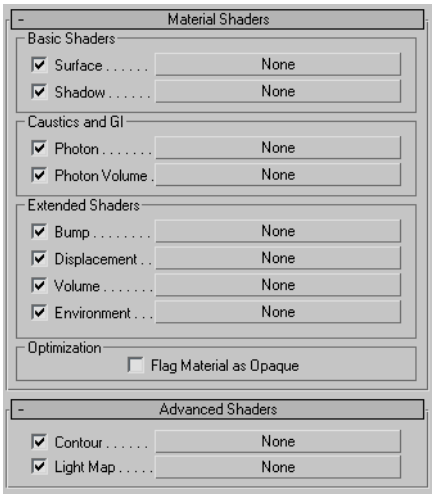

This material gives you full access to mental ray material features. As you can see, the shaders are similar to the ones listed on the mental ray Connection rollout shown above.
While the mental ray Connection rollout overrides the translator settings that translates a **3ds max** material into a mental ray material, this material gives you a way to build a material from scratch.

More Information on Shaders

While we have covered how mental ray is integrated into **3ds max**, and where to find the general settings, most of the power of mental ray relies on the shaders and how they can be combined together to create the desired effects.

A number of shaders are included with **3ds max**, and many more have been developed by third parties. Each shader has specific settings and requires its own documentation.

An explanation of the shaders and their capabilities is beyond the scope of this guide. For more details on these shaders, refer to the online reference and the mental ray documentation.

Rendering Dialogs

All the rendering-related dialogs and settings have been reorganized and grouped into tabbed dialogs. The separate dialogs that you could open from different places are now accessible in a more organized way. This, plus the fact that the dialog has been made smaller, makes the display area less cluttered when working on refining the rendering settings.

The Two New Dialogs

All the rendering settings are now on two dialogs that use tabbed panels for each rendering feature. The two dialogs are Render Scene and Environment And Effects.

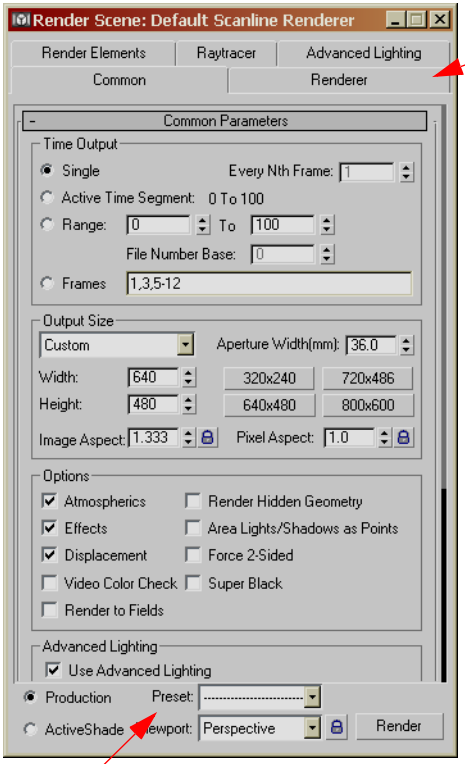

The tabs on the Render Scene dialog update depending on the renderer selection, so that only the relevant options are available. See later in this topic how to switch the renderer.

The Draft settings have been replaced by the more flexible and powerful Render Presets. See later in this topic for how to use them.

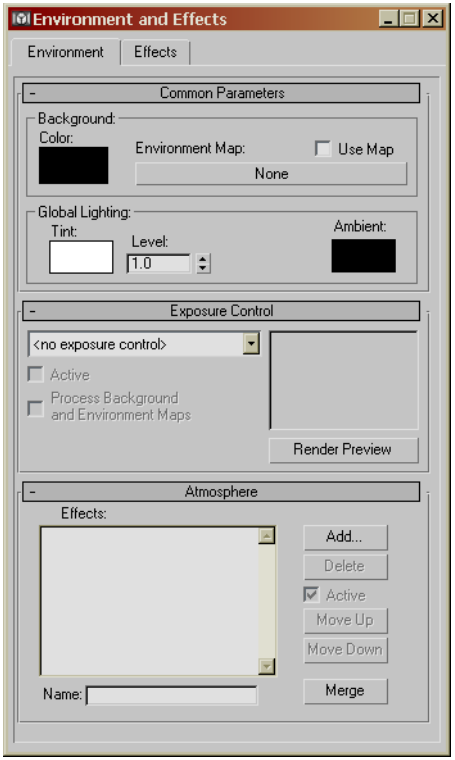

The legacy menu entries that used to access the features now in the panels (like Rendering Effects, Environment, Raytracer, Advanced Lighting), are still present, but now all open the Render Scene or the Environment And Effects dialogs to the appropriate panel.

Switching the Renderer

On the Render Scene dialog, at the bottom of the Common panel, you will find the Assign Renderer rollout. Here, you can set the renderer to use Production renderings, Material Editor sample slots, and ActiveShade windows.

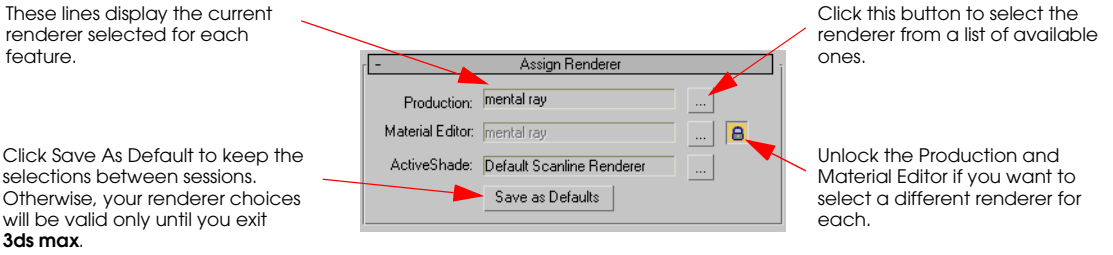

Based on your current rendering selection for Production, the panels will change. Only the panels that are relevant to the renderer will be displayed.

Scanline Renderer

mental ray Renderer

Changing renderers affects other parts of the application, like the material types available in the Material Editor. The switch is immediate, and there is no need to exit and restart **3ds max**.

Rendering Presets

Panels play an important role in the Rendering Preset feature. Rendering Presets are files that store settings from each panel into a file you specify. You can then load these files again, and the saved settings are restored into each panel.

Rendering Presets replace and expand the Production and Draft options from previous versions. Just save your production and draft settings, and they will be available from the list in the Render Scene dialog. Of course, now you are not limited to just two settings, but you can create as many as you want, for various rendering qualities and types, and for individual or common projects.

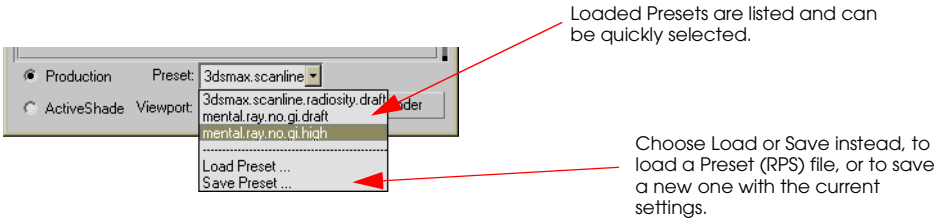

When choosing Save or Load Preset, you will be presented with a list of panels names. Here, you can select or deselect the panels you want to load or save.

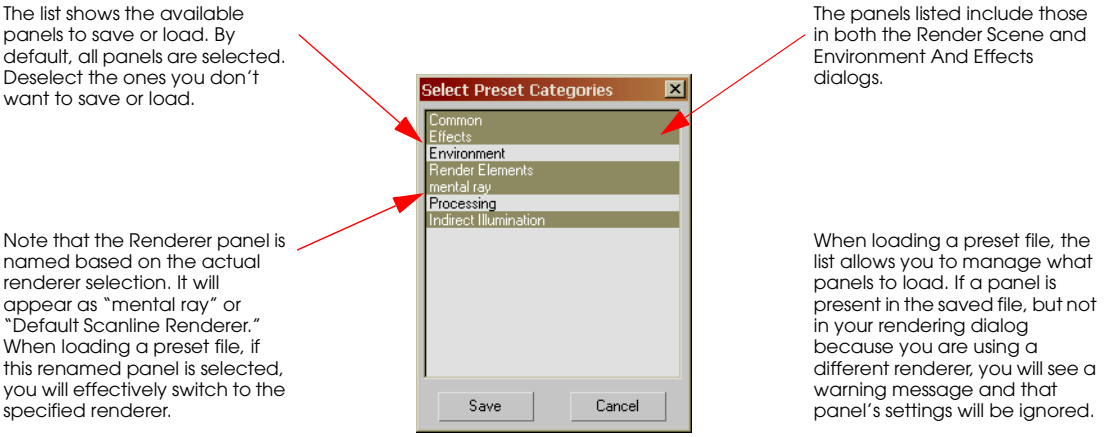

Rendering Presets allow you to share rendering settings between users working on the same project, and to create as many rendering settings as are required for different rendering qualities. In a collaborative project, it might be worth planning the needed settings in advance, and finding a common repository for the files.

Rendering Presets and Command-Line Rendering

Another important role of Rendering Preset files, is that they can be used from the command line to override the settings saved within the MAX file.

If a person in a team is responsible for rendering other users' works, it will be much easier to make consistent renderings by specifying the same rendering presets for all the scenes.

Here is an example:

> **3dsmaxcmd -preset:c:\presets\quickrender.prs renderme.max**
or
> **3dsmaxcmd -prs:\\projectserver\presets\final_with_gi.prs renderme.max**

Global Rendering Options

Two new options globally affect rendering. One option is used to disable area lights and shadows temporarily for quicker test renderings; the other is used to disable or force supersampling.

Area Shadows

Area lights and area shadows can slow down rendering because of the number of calculations required. When their parameters are set and your focus is on something else, it can be useful to disable them and speed up rendering considerably while testing.

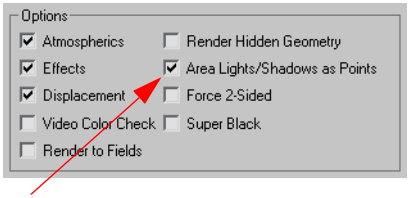

Enable this option to treat all area lights and area shadows temporarily as point lights and shadows.

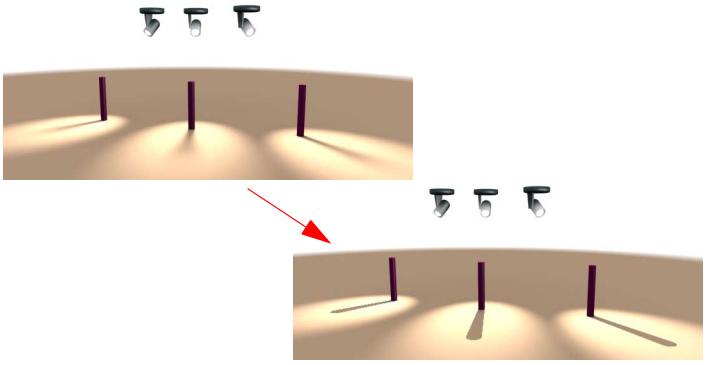

Supersampling

Material supersampling is now set by default to use global settings. On the rendering dialog, on the Scanline Renderer panel, you can specify whether or not you want to use a particular supersampling technique for all materials in the scene.

In the Material settings, Use Global Settings lets you specify a supersampling technique for all materials.

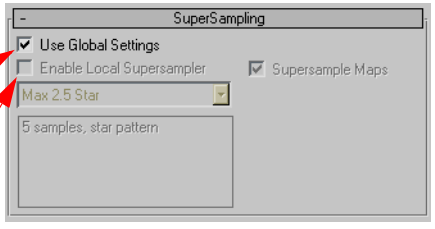

When you turn off Use Global Settings, you can turn on the local option so each material uses its own supersampling technique, as in previous versions.

In the rendering settings, you can disable all supersampling to make quick test renderings. This option disables both global and material-specific samplers.

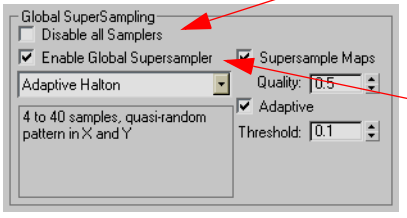

Here, you can set the global supersampler for all materials that use the global setting. Leave it turned off if you don't want any supersampling for the scene's materials.

Supersampling only applies to the scanline renderer. mental ray has its own sampling method, and ignores these settings.

Rendering Lights

Unlike hidden geometry, hidden lights are processed during rendering. This is often desired behavior, so you can hide lights to clean up the viewport display but still render the images with the correct lighting.

But this also makes it harder to hide some lights and test the effect of only a subset of them, especially if the lights, as often happens in architectural modeling, were all instances, and turning one off would turn all of them off.

In this version, lights use the Renderable attribute. If Renderable is off, the light is not processed. The use of this flag has many advantages. First, each instance of a light can be made renderable or not, independently. Second, you can select multiple lights, and use the Object Properties dialog to set the flag for all of them with one click. Third, you can easily turn the flag on or off even from the Layer Manager. Organizing lights by layers becomes a good way to manage them.

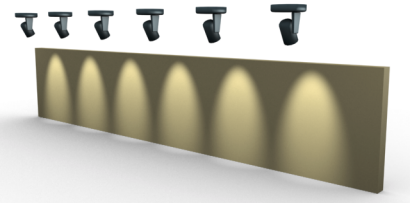

It's common to use instanced lights in architectural modeling. Instancing allows you to control the properties of all the lights of the same kind by editing just one of them.

But instanced lights also are turned on and off all at once. This is useful in some cases, but not in others. When, for example, you want to check the effect of a single light, or test lighting solutions with only a subset of lights turned on.

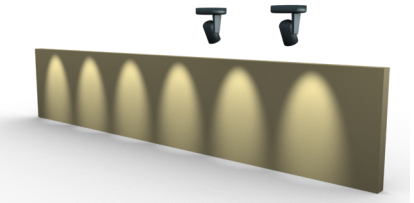

Hiding lights is not the way to achieve the desired result, since lights are processed anyway during the rendering, even if they're not visible in viewports.

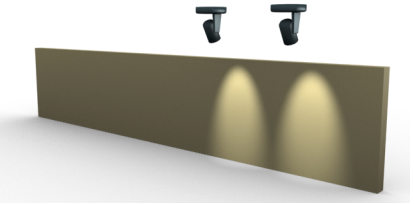

The solution is now easy to achieve by using the Renderable option on the Object Properties dialog. The same option is also available in the Layer Manager.

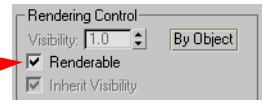

Command-Line Rendering

Rendering can now be performed using a command-line utility called *3dsmaxcmd.exe*, which resides in your **3ds max** installation root directory.

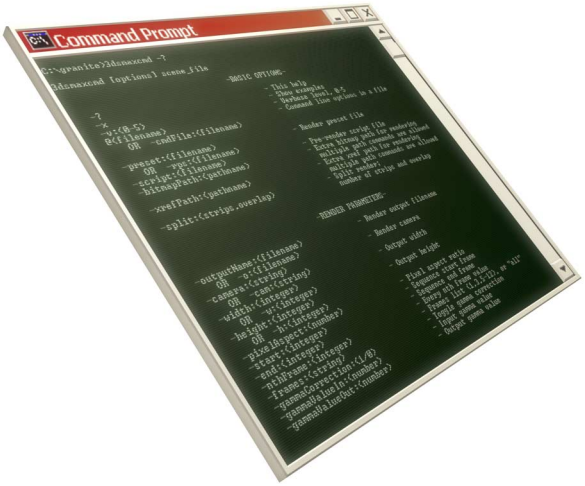

Using *3dsmaxcmd*, you have access to many of the rendering settings available when rendering from within **3ds max**.

The ability to use this utility in a batch file and to have access to rendering presets gives you a powerful tool for fast automating of small rendering jobs on your local computer, as well as for sending multiple jobs to network rendering.

Basic Help

Since the number of command-line parameters is very large, you have access to two quick help pages directly from the command prompt.

> **3dsmaxcmd -?**
This will list all the available parameters. The list is long, so make sure your command prompt window is high enough, or that you can scroll it back.

> **3dsmaxcmd -x**
Shows a series of sample command lines that perform the most common rendering tasks. Find a sample that is similar to what you need, and use it as a template.

Using Command-Line Rendering

The command-line parameters available, are used to override settings that are saved in the MAX file, or to specify other options.

We will not give details on all the options here, but just give you the basics on how this utility operates, and what you can do with it. For a complete list of the options, see the online reference.

In its basic form, you can render a MAX file simply by specifying the file name as the only parameter.

> **3dsmaxcmd myscene.max**

The scene will be rendered using the settings saved within the MAX file. Remember that since the rendering is not interactive, there are no warnings if output files overwrite existing files.

Also, you can use parameters to override the rendering size and the output file name and directory.

> **3dsmaxcmd -width:1024 -height:768 -outputname:C:\pics\final.tga myscene.max**
or
> **3dsmaxcmd -w:1024 -h:768 -o:c:\frames\final.tga myscene.max**

The extension of the output file name is also used to override the output format. Settings for the most common image formats are also available. For example, you can change the JPEG quality.

> **3dsmaxcmd -o:c:\frames\final.jpg -jpeg_quality:75 myscene.max**

Another common override is the range of frames to render, and the paths to bitmaps used for textures.

> **3dsmaxcmd -start:10 -end:50 -bitmappath:c:\maps\myproject myscene.max**

Using Text Files

If you use command-line rendering to render scenes created by other artists, you will have to override quite a few options, and in that case, command lines can become quite long.

To make the parameters easier to maintain and understand, especially when using batch files, you can use text files with the desired parameters.

For example, a text file that overrides output path and format, image size, map paths and sets the rendering to a specific camera might look like this:

-outputname=c:\frames\final.tga
-targa_colordepth=24
-targa_compressed=1
-width=720
-height=486
-pixelaspect=0.9
-bitmappath=c:\maps\myproject
-bitmappath=c:\maps\common
-camera=Camera01

You will then be able to use these settings by referring to the text file. Other parameters can still be used in the command line.

> **3dsmaxcmd -cmdfile:c:\templates\finalrender.txt -start:10 -end:30 myscene.max**
or
> **3dsmaxcmd @c:\templates\finalrender.txt -start:10 -end:30 myscene.max**

Text files can be used to save settings for different kind of quality, like tests or final renderings, or to save common preferred options, so you don't need to open and check each single file for the correct output settings.

Using Rendering Presets

Another way to save rendering settings and reuse them on different scenes, is to use rendering presets files. Rendering Presets are explained in the *Rendering Dialogs* topic.

Rendering presets files can specify more output options, and also can make the command-line rendering task easier to manage.

Just save a rendering preset file with the desired setting, and then use it from the command line:

> **3dsmaxcmd -preset:c:\3dsmax6\renderpresets\mysettings.rps myscene.max**
or
> **3dsmaxcmd -rps:c:\3dsmax6\renderpresets\mysettings.rps myscene.max**

Other command-line options can be used as well in combination with the rendering presets.

Network Rendering

Jobs can be submitted for network rendering, too. You can submit a scene or a single frame, using the new split image feature. See the *Region Net Render* topic for more information.

For example, you can submit a job to a specific rendering farm.

> **3dsmaxcmd -submit:manager01 -jobname:"myscene test01" myscene.max**

Here is an example of how to submit a single frame for rendering in strips on the rendering farm:

> **3dsmaxcmd -submit -split:100,5 myscene.max**

More Information

For your convenience, we include a quick reference for the available command line parameters on the following pages.

For more details on individual parameters, refer to the online reference.

Command-Line Quick Reference

C:> 3dsmaxcmd [options] scene_file

Basic Options

-?	Displays the online reference
-x	Show examples
-v:<0-5>	Verbose level, 0-5
@<filename> *or* **-cmdFile:<filename>**	Command-line options in a file
-preset:<filename> *or* **-rps:<filename>**	Render preset file
-script:<filename>	Pre-render script file
-bitmapPath:<pathname>	Extra bitmap path for rendering. Multiple path commands are allowed.
-xrefPath:<pathname>	Extra XRef path for rendering. Multiple path commands are allowed.
-split:<strips,overlap>	Split render; number of strips and overlap

Render Parameters

-outputName:<filename> *or* **-o:<filename>**	Render output file name
-camera:<string> *or* **-cam:<string>**	Render camera
-width:<integer> *or* **-w:<integer>**	Output width
-height:<integer> *or* **-h:<integer>**	Output height
-pixelAspect:<number>	Pixel aspect ratio
-start:<integer>	Sequence start frame
-end:<integer>	Sequence end frame
-nthFrame:<integer>	Every nth frame value
-frames:<string>	Frames list (1,3,5-12), or "all"
-gammaCorrection:<1/0>	Toggle gamma correction
-gammaValueIn:<number>	Input gamma value
-gammaValueOut:<number>	Output gamma value

Render Flags

-showRFW:<0/1> *or* **-rfw:<0/1>**	Show/hide Render Frame Window during render
-skipRenderedFrames:<1/0>	Toggle skip existing images
-videoColorCheck:<1/0>	Toggle video Color Check
-force2Sided:<1/0>	Toggle force 2-sided
-renderHidden:<1/0>	Toggle render hidden
-atmospherics:<1/0>	Toggle atmospherics
-superBlack:<1/0>	Toggle super black
-renderFields:<1/0>	Toggle render to fields
-fieldOrder:<even/odd>	Field order
-displacements:<1/0>	Toggle displacement mapping
-effects:<1/0>	Toggle render effects
-useAreaLights:<1/0>	Toggle area lights/shadows
-useAdvLight:<1/0>	Toggle use advanced lighting
-computeAdvLight:<1/0>	Toggle compute advanced lighting
-ditherPaletted:<1/0>	Toggle output dithering (paletted)
-ditherTrueColor:<1/0>	Toggle output dithering (true color)

Backburner Job Submission

-submit(:manager_name) *or* **-s(:manager_name)**	Submit to **backburner**
-port:<integer>	Manager port number
-netmask:<string>	Network mask other than 255.255.255.0
-jobName:<string>	Job name
-priority:<integer>	Job priority
-suspended:<1/0>	Toggle initially suspended
-writeJobFile	Write all job settings to XML file
-readJobFile:<filename>	Read all job settings from XML file

Bitmap Parameters

-BMP_TYPE:<2/8>	.BMP type, 2=paletted, 8=true 24
-JPEG_QUALITY:<1-100>	.JPG quality value
-JPEG_SMOOTHING:<1-100>	.JPG smoothing value
-TARGA_COLORDEPTH:<16/24/32>	.TGA color bit depth
-TARGA_COMPRESSED:<1/0>	.TGA compression on/off
-TARGA_ALPHASPLIT:<1/0>	.TGA alpha split on/off
-TARGA_PREMULTALPHA:<1/0>	.TGA premultiplied alpha on/off

Bitmap parameters (continued)

-TIF_TYPE:<0/1/2/3>	.TIF type, mono/color/logl/logluv
-TIF_ALPHA:<1/0>	.TIF alpha on/off
-TIF_COMPRESSION:<1/0>	.TIF compression on/off
-TIF_DPI:<number>	.TIF dots-per-inch value

For each of the following there is a corresponding -RPF_xxx option.

-RLA_COLORDEPTH:<8/16/32>	.RLA color bit depth
-RLA_ALPHA:<1/0>	.RLA alpha on/off
-RLA_PREMULTALPHA:<1/0>	.RLA premultiplied alpha on/off
-RLA_DESCRIPTION:<string>	.RLA description (in quotes)
-RLA_AUTHOR:<string>	.RLA author name (in quotes)
-RLA_ZCHANNEL:<1/0>	.RLA Z-depth channel on/off
-RLA_MTLIDCHANNEL:<1/0>	.RLA material effects channel on/off
-RLA_NODEIDCHANNEL:<1/0>	.RLA object channel on/off
-RLA_UVCHANNEL:<1/0>	.RLA UV coordinates channel on/off
-RLA_NORMALCHANNEL:<1/0>	.RLA surface normals channel on/off
-RLA_REALPIXCHANNEL:<1/0>	.RLA non-clamped color channel on/off
-RLA_COVERAGECHANNEL:<1/0>	.RLA coverage channel on/off

For each of the RLA options above, there is a corresponding -RPF_xxx option, plus the following which are specific to RPF and unavailable for RLA.

-RPF_NODERENDERIDCHANNEL	.RPF node render ID channel on/off
-RPF_COLORCHANNEL	.RPF color channel on/off
-RPF_TRANSPCHANNEL	.RPF transparency channel on/off
-RPF_VELOCCHANNEL	.RPF velocity channel on/off
-RPF_WEIGHTCHANNEL	.RPF sub-pixel weight channel on/off
-RPF_MASKCHANNEL	.RPF sub-pixel mask channel on/off

Options are not case-sensitive.
MAXScript names for rendering parameters and flags are accepted as aliases, where applicable.

Pre-Render Script

For noninteractive renderings, network and command-line rendering, you can specify a script that will be executed before the actual rendering begins. This script can be used to customize the scene for the rendering.

While rendering parameters can be set through rendering presets or from command-line options, other operations, like hiding and unhiding or changing object properties, can be accomplished using this script.

Network Rendering

In the Network Job Assignment dialog, click the Advanced button. In the Advanced Settings dialog, a new field is available where you specify the script.

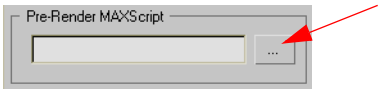

Enter the script file name, or use the Browse button to choose it. Make sure that the path name of the script is accessible to all the computers that will network render.

Command Line

The command-line option **-script:**<filename> specifies the script. If used in combination with the **-submit** option for network rendering, make sure the script is accessible to every rendering server.

For example:

> **3dsmaxcmd -script:c:\myscripts\preparescene.ms renderme.max**

> **3dsmaxcmd -script:\\scriptserver\scriptshare\preparescene.ms renderme.max**

Panorama Exporter Utility

This utility creates six renderings from the point of view of the selected camera, and then seamlessly combines them into a 360-degrees view of the scene. A specific viewer application is launched at the end of the rendering. Files can also be saved for later viewing without rerendering.

You can start the utility from the Rendering menu -> Panorama Exporter, or from the Utilities panel.

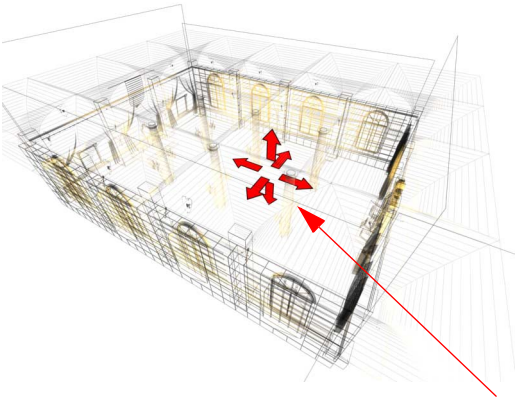

Add a camera to your scene, and orient it toward the direction you want the viewer to display when it opens. Set the desired initial field of view as well.
The utility then renders six images, encompassing the field of view, to create six renderings of the scene around the camera.

The six rendered images are composited into a seamless panoramic image. The viewer lets you look at the image as if you were inside the scene.

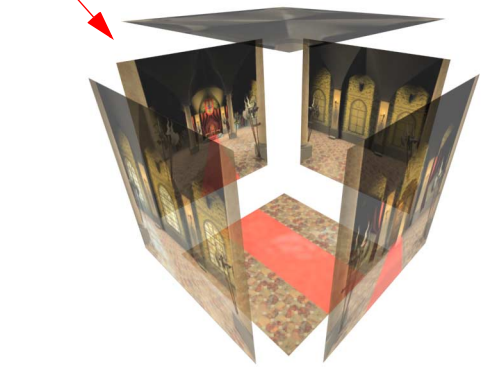

In the Utilities panel, you have access to the two components of this feature: the renderer and the viewer.

Click here to start the renderer. The dialog that appears is a custom version of the rendering dialog (right). The mental ray renderer is supported.

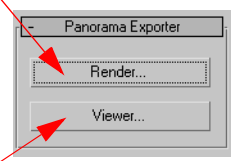

The image size is the size of the composite panoramic image (see next page). Use a small size for testing, and a very big one for the final one, to avoid blurry details.

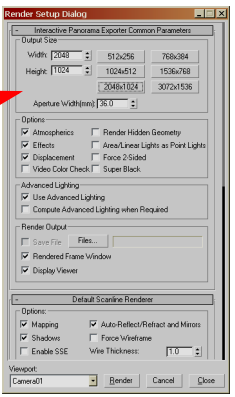

Click here to start the Viewer. After a rendering, the viewer starts automatically, but if you want to load an old image, you can start the viewer independently.

The Viewer

The viewer is a simple window that allows you to browse the rendered scene, starting from the original camera point of view. You can pan the image, zoom in and out, and save the rendered scene, or load another rendering saved previously.

From the File menu, you can load and save the rendered panoramic images. See below for more information.

Click and drag the mouse to pan the image. The position of the pointer in respect to the center of the image sets the direction and the speed of the panning.

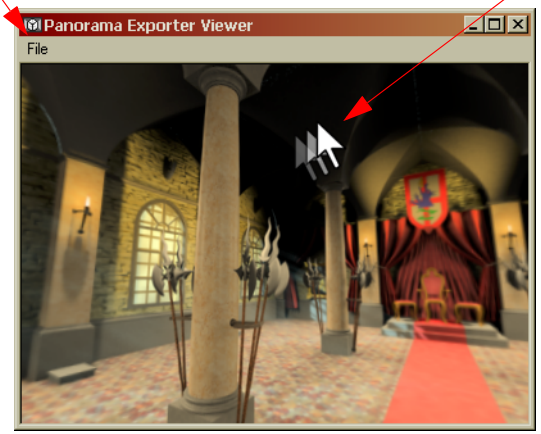

Use the middle mouse button (or press CTRL+ALT and the left mouse button) and drag the mouse up and down, to zoom in and out.

Use the right mouse button (or press ALT and the left mouse button) and drag the mouse, to pan in a different way. The mouse pointer "grabs" the image and you moves that point around in the scene.

Rendered images can be exported and reloaded later for viewing without the need to re-render. The renderings can be exported as cylindrical or spherical projections, and can also be used as maps (for example, as reflection maps) in the same or in another scene.

The viewer can visualize correctly only spherical projections. A pinch is otherwise visible at the top and bottom of the panorama if a cylindrical projection is used.

Spherical projection

Cylindrical projection

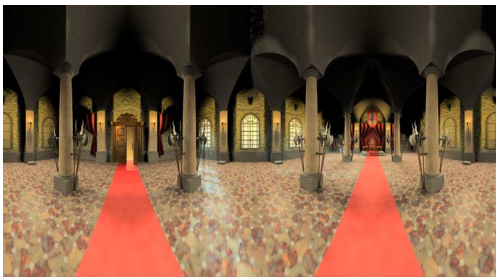

The image size specified for the rendering is, in fact, the size of the spherical projection. For quality results, panoramas need to be rendered quite large, especially if you want to zoom in while viewing.

Print Size Wizard

The Print Size Wizard, available from the Rendering menu, allows you to calculate the size of the rendered image based on print settings rather than pixels.

The final image size in pixels is calculated based on the size on the paper and the DPI setting.

The image can then be rendered directly as a TIF file, or you can open the rendering dialog with the settings just calculated, and set more options there before the final rendering.

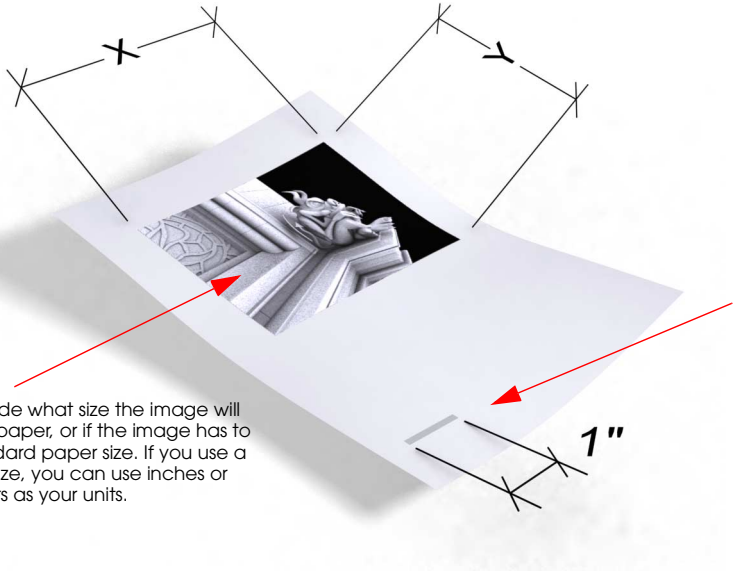

When printing, the image size on paper and the resolution come first. The size in pixels of the image is a consequence of how big and how sharp you want the printed image to be.

The DPI (dots per inch) setting is a standard measurement of the image resolution. Make sure you know the maximum resolution of your printer. There is no need to use settings higher than what the printer can handle.

First, decide what size the image will have on paper, or if the image has to fill a standard paper size. If you use a Custom size, you can use inches or millimeters as your units.

Choose a standard paper size, or a Custom size. If you choose Custom, you will be able to enter the size in (mm or inches) of the image on paper.

Then choose the resolution (DPI) by selecting one of the presets, or specifying a custom value.

Click to open the standard rendering dialog for more options before rendering. Your size settings will be transferred.

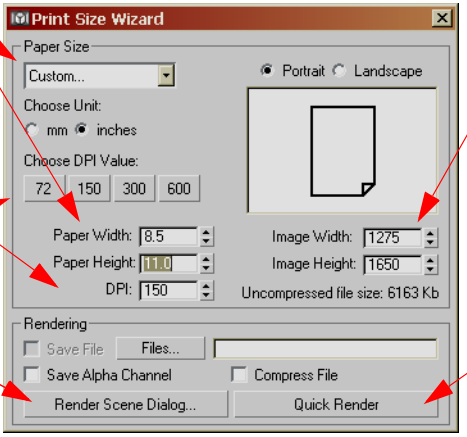

The size in pixels is calculated automatically based on the paper size and DPI setting. You can also change the size in pixels here, and the print size will automatically update.

For a quick rendering, set the optional output file name and use Quick Render. The file will be saved as TIFF, which is a standard format for print files.

Region Net Render

A new option in the Network Job Assignment dialog lets you split a single frame into a set of strips, and assign each strip to a different rendering server. The strips are then combined into a single frame again.

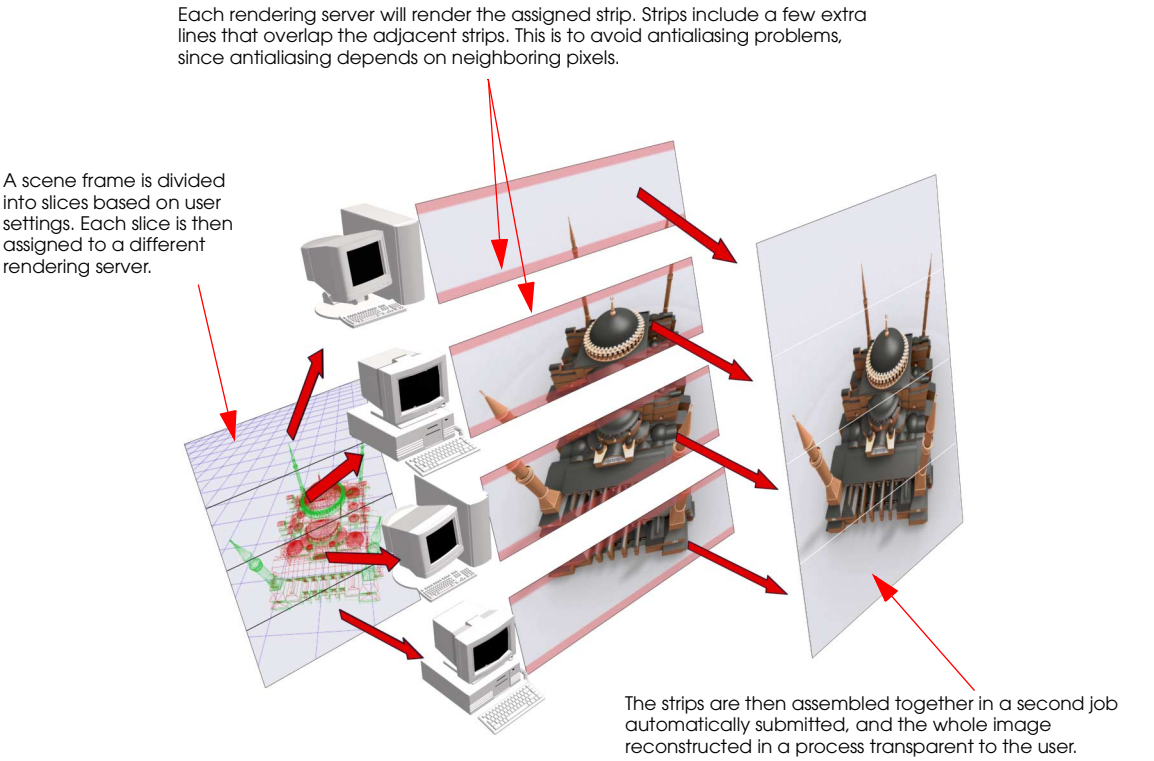

The user interface is pretty straightforward.

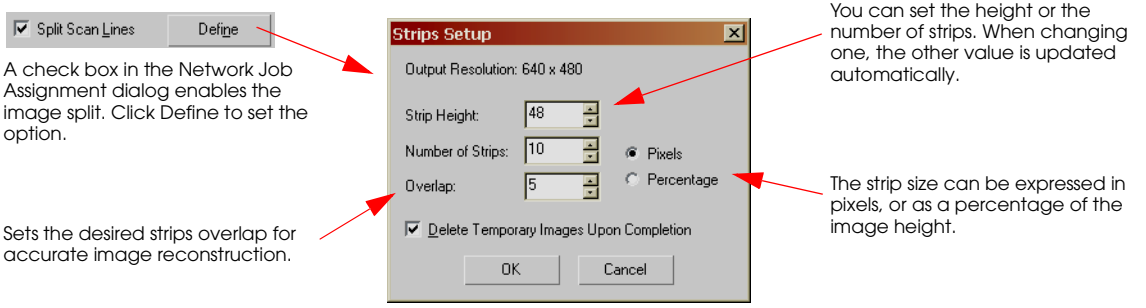

Render Elements is not supported when this feature is used.

New Render Elements

Three new Render Elements have been introduced: flat Diffuse, Lighting, and Matte. They can be used independently, but their main purpose is to allow easy change of diffuse color during compositing.

Diffuse Element

The Diffuse element was in the previous versions, but its output was a combination of the color and the lighting.

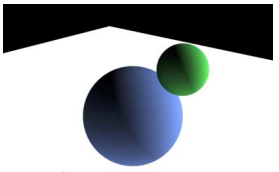

Previously, the diffuse element was a combination of the diffuse color and the lighting.

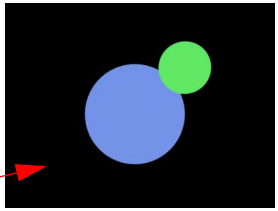

The new Lighting option on the Diffuse rollout allows you to ignore the lighting and output just the flat diffuse color. Notice also that the exclusion of lighting makes the matte material plane disappear.

Lighting Element

The Lighting element gives you a grayscale of the light intensity for one or more lighting components.

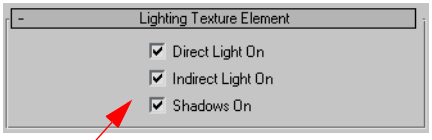

The options on the Lighting rollout let you select what lighting components you want to include.

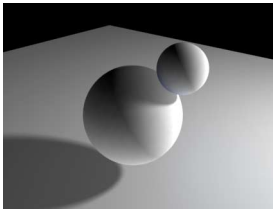

All options enabled

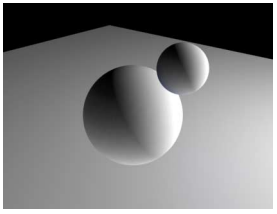

No shadows

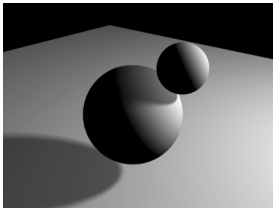

No indirect light

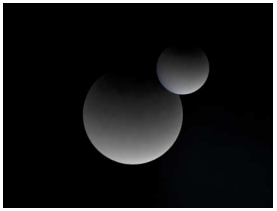

Indirect light only

Matte Element

The Matte element gives you a solid mask that can be used to isolate specific objects in the scene. Transparency is not considered.

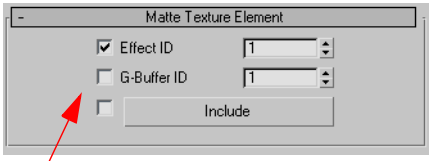

You can use the Effect ID or the G-Buffer ID to select the objects that you want in the mask, and eventually filter the objects even more using the exclude/include list.

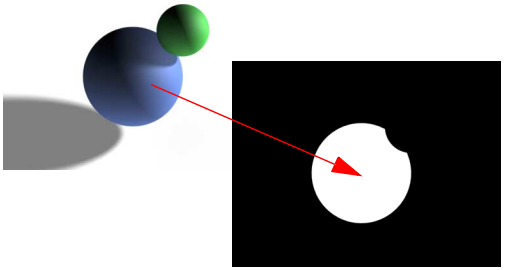

An Example

Let's see how we can use these three new elements to change the diffuse color of an object in a scene.

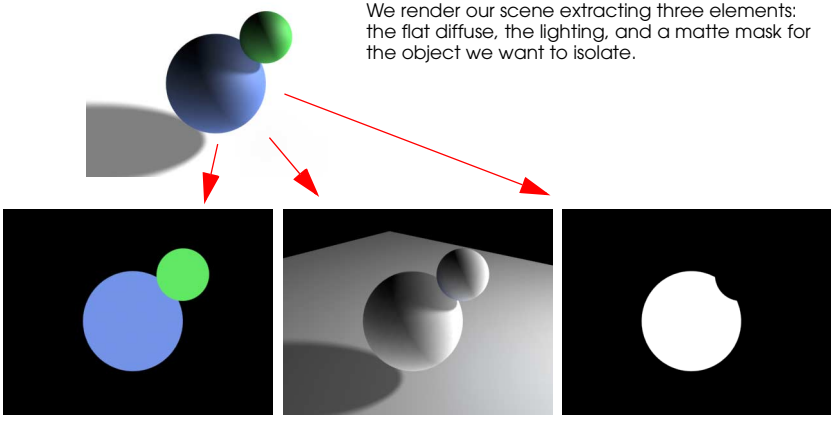

We render our scene extracting three elements: the flat diffuse, the lighting, and a matte mask for the object we want to isolate.

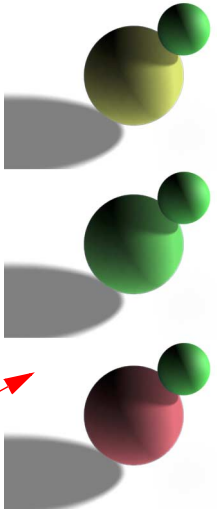

In a compositor program, you can now use the original image, the matte mask to isolate the object, and the flat diffuse combined with the lighting map (multiply mode). With no other changes, the resulting image will match the original one.

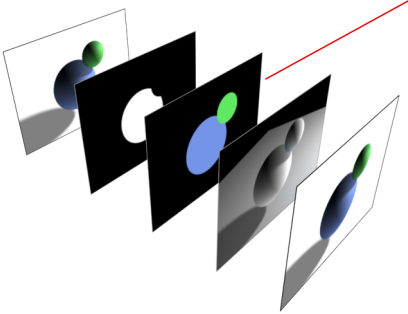

Changing the color of the flat diffuse element (by using a Hue adjustment, for example), now allows us to change the rendered object's color realistically.

Render to Texture

Many changes have been made to Render To Texture. More options and more control over the resulting output make this feature more usable and flexible. Render To Texture can also take advantage of Network Rendering.

User Interface: General Settings

You will notice many changes to the user interface. The way Render To Texture works is fundamentally the same, but you have more control over the output settings for each object.

You can set the mapping channel for the automatic unwrap, instead of using channel 3 by default.

The Objects To Bake group in General Settings has been replaced by a more complete rollout for additional control on each object output.

Skip Existing Files lets you repeat the rendering, keeping what was previously done. Useful when adding output elements.

Rendered Frame Window enables or disables the rendering preview.

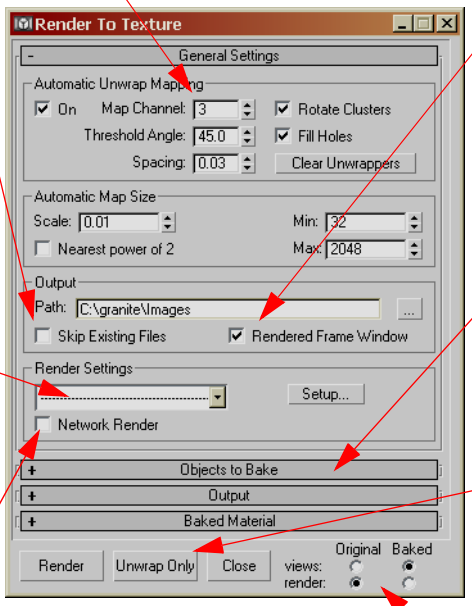

See the following pages for a description of these additional rollouts.

You can take advantage of the Rendering Presets as for regular scene renderings. See *Rendering Presets* in the *Rendering Dialogs* topic.
Note: You cannot use mental ray with Render To Texture. If mental ray is your current renderer, the Render button will be disabled.

Use Unwrap Only to prepare the listed objects and automatically create a mapping channel.

Turn on Network Render to distribute the task to the rest of the rendering farm. Note that the network task is limited to the creation of the textures. Shell materials need to be updated once the network job is completed, by clicking Update Baked Material on the Baked Material rollout.

Here, you can change the Shell Material display option for all the listed objects in the Objects To Bake rollout. This is a quick way to view the rendering result.
You can make individual display settings by choosing only one object from the list, or from the Shell Material parameters in Material Editor.

Objects to Bake

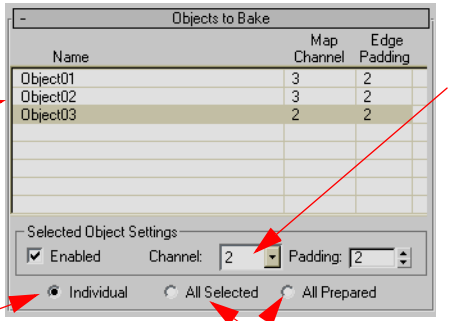

This is the list of objects that will be processed. Make sure everything you need is listed here before starting the rendering.

The map channel used for texture creation can be specified as one of the existing channels available for that object.
Note that this field is enabled only if automatic unwrapping is off on the General Settings rollout.
If automatic unwrapping is on, all objects will use the same automatic channel.

Individual lets you change the selected object Settings for individual objects in the list. Select all the objects in your viewport to see them all listed, then choose from the list above.

When All Selected or All Prepared is chosen, the list is populated with all the objects selected in the viewport, or all objects with mapping channels usable for texture rendering. Selection of individual objects from the list is disabled.

Output

The list of the elements to be rendered is updated based on the choice of Individual or All Selected/ All Prepared.
Only when Individual is chosen will you see the File Name field. You will be able to set a different number of elements for each, plus different target maps and custom file names and sizes.

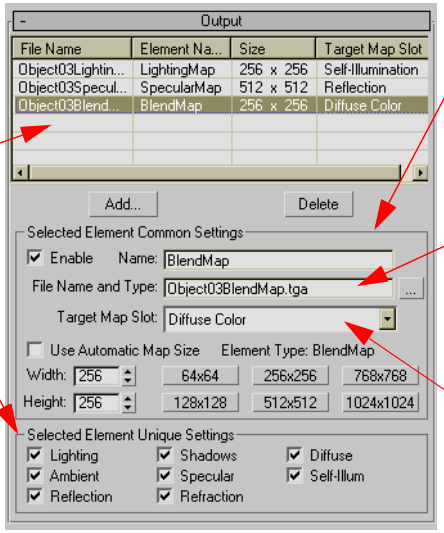

All the settings in this group are set for the objects listed in the Objects To Bake list. Unless Individual is chosen, the changes affect all objects.

You can change the map name only when Individual objects are selected from the Objects To Bake list. It's disabled otherwise.

In this area, different settings appear, depending on the element selected from the list above. In this case, the options are the ones for the Blend map.

Target Map Slot allows you to assign an element to a specific map of the baked material (see the note below).

Each rendered element can be assigned to a specific map of the baked material. Choose the element from the list, and then choose a map from the Target Map Slot.

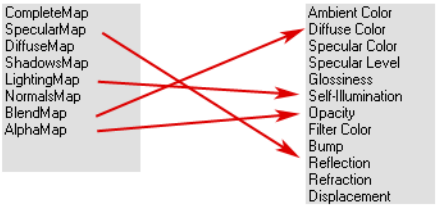

Baked Material

Turn off this option if you just want to create the texture files, without creating the Baked and Shell materials.

Output Into Source will not create the Shell material and save the original one as a sub-material, but *will* replace the maps of the original material directly. This option is useful for the final rendering when no further tweaking is expected. The scene materials will be easier to manage.

Save Source is the default choice, and enables the creation of the Shell material to save the original one.

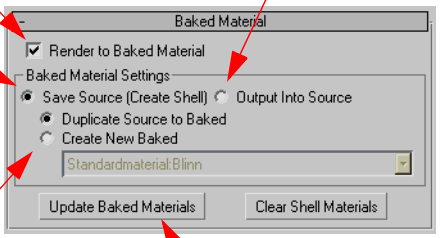

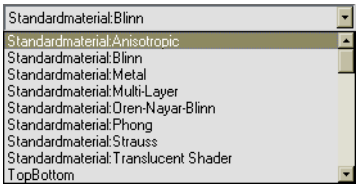

When Save Source is chosen, you have a choice of creating the Shell material as a copy of the original or as a brand new one. In the second case, when you turn on Create New Baked, the drop-down list allows you chose what type the new material will be.

Use Update Baked Materials to refresh the materials and maps association based on the current settings. This is the command to use after Network Rendering has completed texture creation.