# Making It Beautiful

*3ds max has always had one of the fastest renderers in the 3d industry. Its native scanline renderer has been the backbone of everything from cinematic special effects to architectural rendering to high-resolution print work. In short, it's a fast workhorse that gets most every rendering job done quickly. (People who want better control and more rendering options can also employ*

# Making It Beautiful

## rendering tips

*some excellent third-party renderers, such as the Brazil Rendering System [http://www.splutterfish.com/], Cebas's finalRender [http://www.finalrender.com] [http://www.cebas.com/], and The Chaos Group's V-Ray [http://www.chaosgroup.com].)*

*However, for those desiring more high-end render control within the core 3ds max package, fear not: With the release of 3ds max 6, the acclaimed mental images' mental ray 3.2 renderer is included. The mental ray 3.2 renderer brings fast raytracing, true sub-pixel displacement, advanced shaders and a complex shader language to adventurous 3ds max users.*

*In this chapter, I'll present several dozen tips on how to get the most out of your native 3ds max scanline renderer. In addition, contributor Aksel Karcher (a freelance designer and lighting technical director, at [http://www.akselkarcher.com/] weighs in with several mental ray tips for those 3ds max users who want to stay on the cutting edge.*

## RAY TRACING: SPEEDING UP THINGS (OBJECTS)

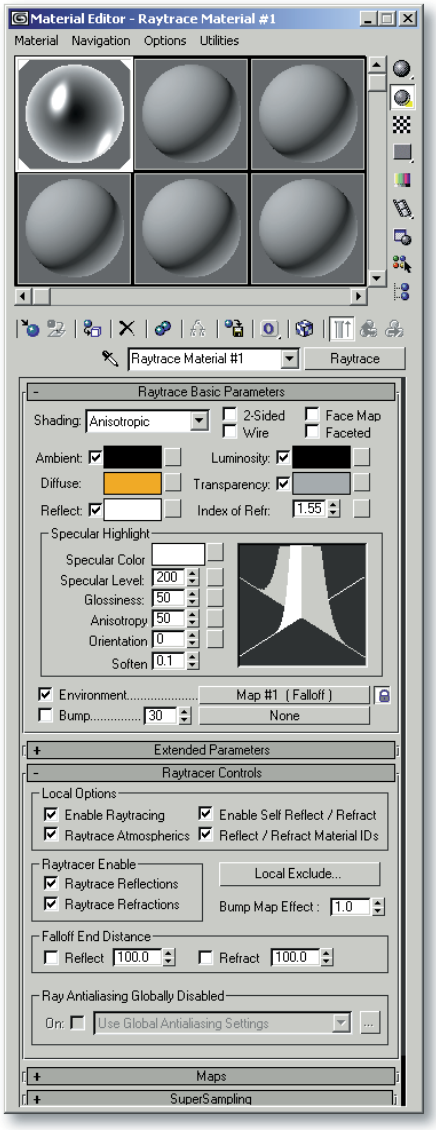


If you love the look of ray tracing in your scene (for glass and metallic surfaces) but you're not a fan of raytraced material rendering times when using the 3ds max scanline renderer, don't worry: You can speed up your renderings by doing a few simple things.

First, make sure antialiasing is unchecked in the Rendering > Raytracer Settings > Raytracer/Global Ray Antialiaser menu when you're doing test renderings, then turn antialiasing back on when you're doing your final production rendering. (Note: You must be using the 3ds max default scanline renderer as your production renderer. If you have mental ray chosen instead, the Raytracer Settings and Raytrace Global Include/Exclude menu items are grayed out.) Second, check your scene object integrity: make sure you've welded the cores of Lathed objects, that objects have unified face normals, and that the objects aren't degenerate. (That is, they should not have missing or coincident faces, overlapping vertices, and so on.) Third, if you don't need to keep the modifier stacks active for some or all of your scene objects, then collapse them to the modifier stack results (preferably Editable Meshes). Fourth, keep your Raytrace material as one-sided instead of two-sided, unless it's absolutely necessary to represent surfaces such as thick glass.

## RAY TRACING: SPEEDING UP THINGS (GLOBALS)

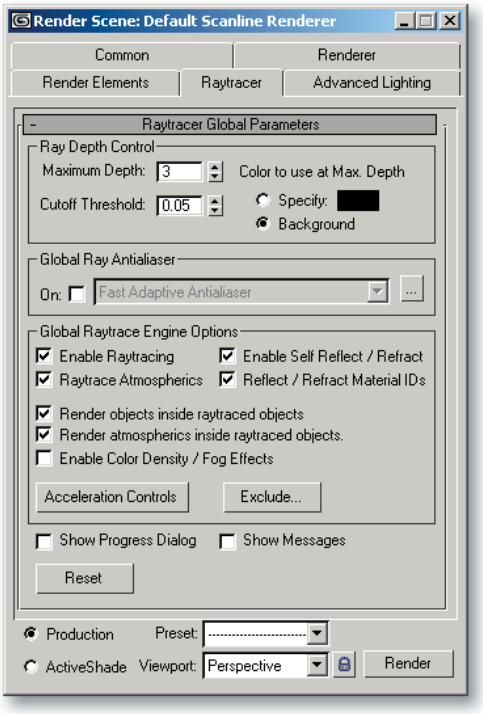


If you go to the Rendering > Raytracer Settings > Raytracer tab, you'll see that the default Ray Depth Control (Maximum Depth) is set to 9. This refers to how many times rays are "bounced" through the scene, and how many reflective objects will reflect in each another. ("Department of Redundancy Department here….") However, this is usually overkill for most basic scenes; unless you really need a "hall of mirrors" effect, try turning this setting down to 2 or 3; it will render much faster.

Note that the Raytracer Global Include/Exclude settings (in the Raytracer menu mentioned earlier) are for use with the 3ds max default scanline renderer only; changing these settings has no effect if you're using the mental ray renderer, which has its own ray tracing controls, under the mental ray rendering menu section Rendering Algorithms.

## RAYTRACING: INDEX OF REFRACTION (IOR)

One of ray tracing's great strengths (besides creating physically realistic reflections for chrome surfaces and the like) is its capability to replicate the look of transparent materials. When light passes through a transparent surface, the light is typically bent or distorted. This distortion is known as refraction, and the amount of refraction is known as the *index of refraction* (IOR). The IOR results from the relative speed of light as it passes through a transparent medium relative to the medium that the viewer is in. Often, the more dense the object, the higher the IOR value will be.

Do you want to render realistic transparent objects using the 3ds max Raytrace material, but you can't find that convenient chart listing common substances and their IORs? Well, copy this down and put it in your pocket so that you can answer accurately the next time some stranger asks you, "Hey, dude, you gotta tell me—what's the IOR of table salt again?"

| Material | IOR |
|---|---|
| Air | 1.0003 |
| Water | 1.33 |
| Ethyl alcohol | 1.36 |
| Glass | 1.5–1.7 |
| Lucite or Plexiglas | 1.51 |
| Crown glass | 1.52 |
| Sodium chloride (salt) | 1.544 |
| Quartz | 1.544 |
| Flint glass | 1.58 |
| Diamond | 2.42 |

To use these values, just place them in the IOR spinner. (In the Raytrace material, the spinner is in the Raytrace Basic Parameters rollout; in the Standard material, it's in the Extended Parameters rollout.)

## RENDERING: FILTERS FOR STILL IMAGES VERSUS VIDEO

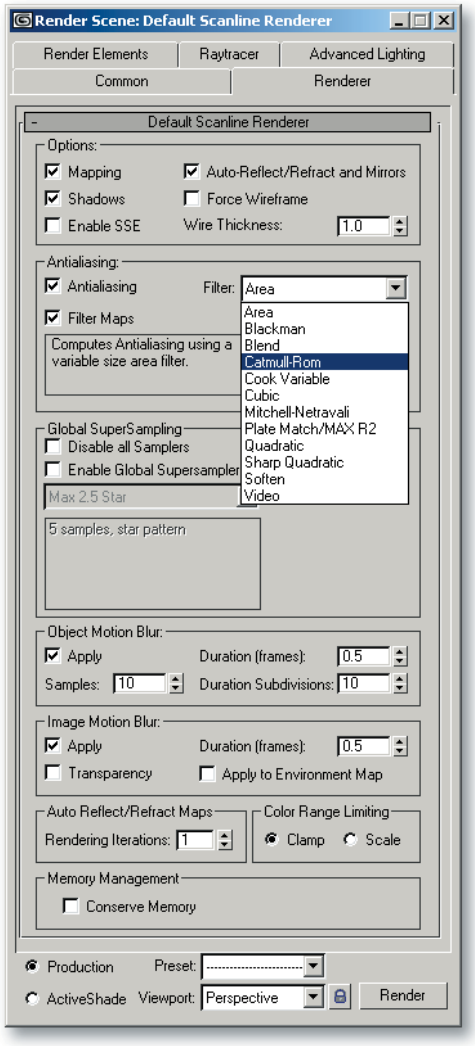


Rendering to video presents a different set of considerations than when you're simply rendering beautiful CG still images for print. If you're rendering to video, you should probably avoid using the 3ds max program's "sharpening" filters (in the default scanline renderer); those filters tend to add edge detail and might create aliasing problems on object edges, or texture scintillation (swimming) of both bitmap and procedural textures. The "sharpening" filters include Blackman, Mitchell-Netravali, and especially Catmull-Rom. They might look great for print, but be careful when using them for video; do extensive (animated) rendering tests first and preview your video work on a standard television before you commit to production rendering. (The scanline renderer video filter is designed specifically to mitigate aliasing on horizontal lines, so you should probably use that instead.)

## RENDERING VIDEO: RENDER FRAMES, NOT FIELDS

Although many video production old-timers like rendering to fields, for a more realistic look, it's almost always preferable to render CG animation to individual frames and use motion blur (object, scene, or image) to smooth out fast movement in the frame. Frames are easier to deal with in paint and compositing programs (which many professionals use to tweak their results before committing to tape), and the end result looks much more like film than video. In addition, the filmic look helps to take off the inherent "CG look" of your renderings, whereas CG rendered to fields has a sharp "video" quality that many find undesirable. (However, if you want to show a CG representation of a "video" point of view in your 3D scenes, you might render those scenes in fields, and the rest in frames; the difference will be noticeable upon playback.)

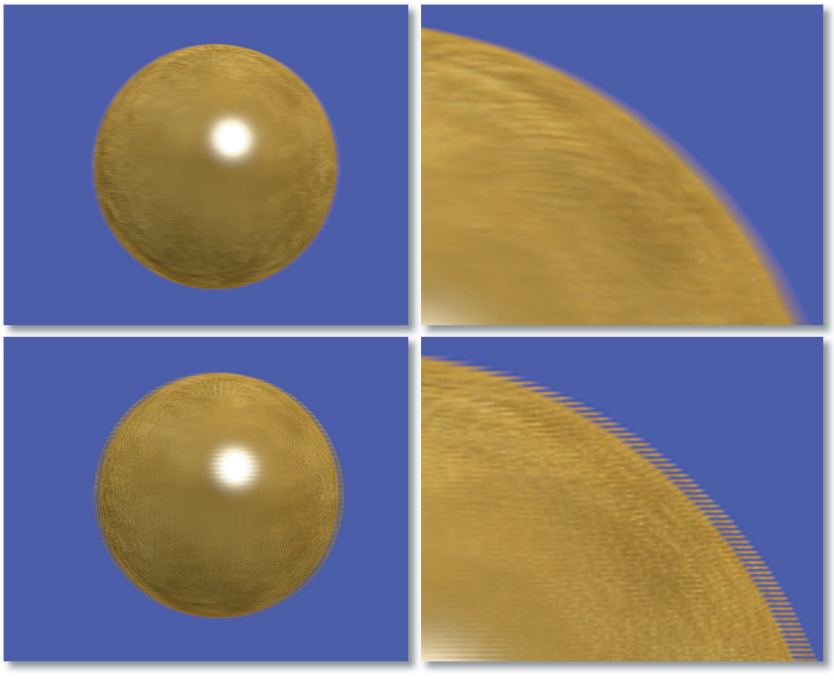

## RENDERING VIDEO: MAKE SURE VIDEO COLOR CHECK IS ON

When you're rendering for video playback on a standard TV set (either NTSC or PAL), you need to make sure your renderings are color safe for video. Colors that are too bright or saturated smear or burn when broadcast on the average TV, making your work less impressive.

To address this problem, go to your Customize > Preferences > Rendering menu. Under Video Color Check, there are three options: Flag with Black (so you can do test renderings, see exactly where the offending colors are, and change them), Scale Luma (so you can bring down the brightness of the offending colors), and Scale Saturation (so you can bring down the offending Color value to acceptable limits). For peace of mind, I suggest selecting Scale Saturation by default; when you go to the Rendering > Render > Common menu, make sure that the Video Color Check box is checked under the Options area. Then, render away, safe in the knowledge that you're not going to offend your TV screen—or your viewers' eyesight.

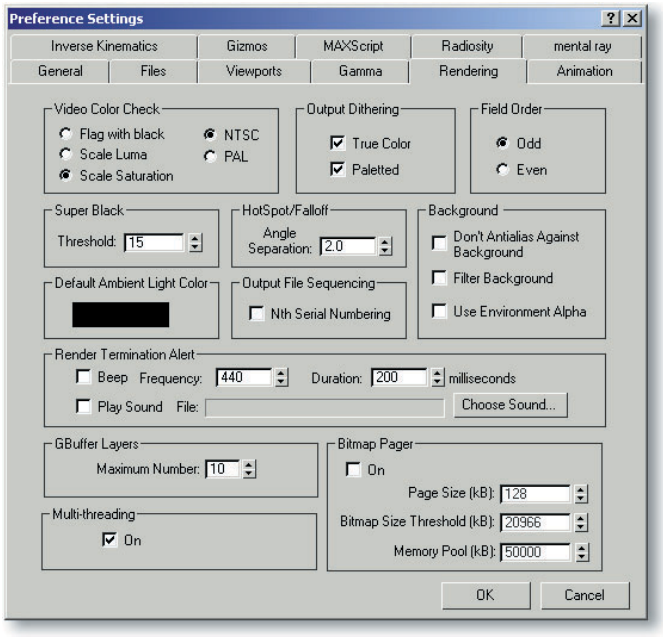


## RENDERING VIDEO: CHECK YOUR GAMMA!

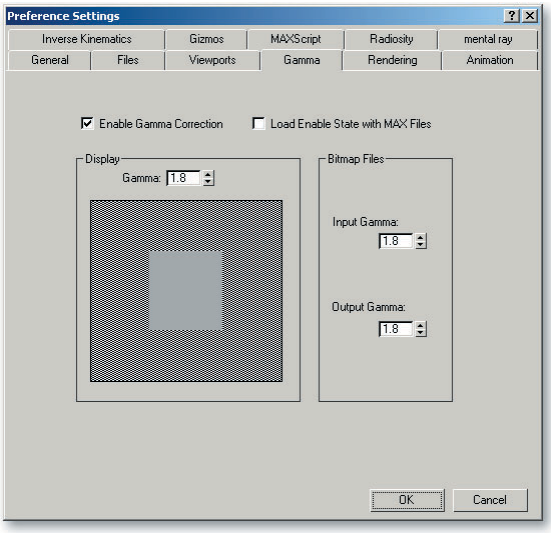


Changing your gamma settings before you render won't turn you into the Incredible Hulk, but it can improve your video renderings. By default, 3ds max has its gamma correction turned off, but if you're rendering to video, you should probably turn it on, or your renderings will appear extremely dark on broadcast video. To set this, go to Customize > Preferences > Gamma, check Enable Gamma Correction, and then adjust your display settings. For video, I usually leave my Display Gamma at 1.8, and I change my Bitmap Files' Input and Output Gamma to 1.8 as well. Note that if you have your Material Editor menu open when you change your gamma settings, clicking OK in the Preference Setting menu shows the change in luminance immediately on the sample spheres and their backgrounds in the Material Editor.

A disclaimer: Gamma correction changes your input, display, and rendered output results across the board, not only for video images. You should not make this decision lightly without understanding the algorithm and nature of the changes to your renderings. Certain bitmap formats do not store gamma values, and some 2D applications do not gamma-correct; therefore, tracking your bitmap assets when mapping or compositing in a multiuser production environment becomes a bit trickier than when gamma is off. If you're in such a production environment, instead of setting Display Gamma to 1.8, you should adjust this setting to match each artist's monitor. With gamma set correctly across all devices, your rendered output should look the same, regardless of your final medium.

## RENDERING GREAT BIG, GIANT, HONKING, ENORMOUS IMAGES SUCCESSFULLY

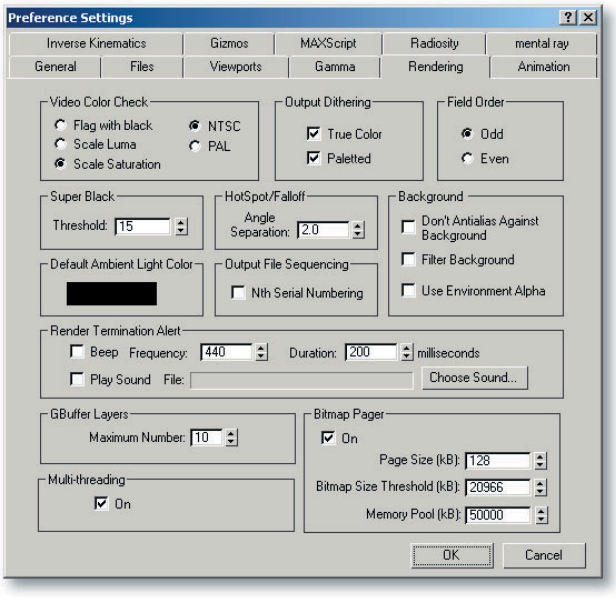


Sometimes you might be called upon to render an enormous, multi-thousand-pixel resolution image, such as for a poster, a billboard advertisement, or other large sign. Even though 3ds max can render images of up to 10,000 × 10,000 pixels, just processing an image like this can drain gigabytes of RAM right out of the fastest, best-equipped workstation. In these cases, if you simply can't get a giant image to render without your machine running out of RAM and dying a painful death, then you need some hard drive cache help. Just go to Customize > Preferences > Rendering, check the Bitmap Pager "On" box, and adjust the memory size parameters as necessary. (You'll have to close and restart 3ds max for this setting to take effect.) 3ds max pages the additional memory required to render the scene to your hard drive, and you get your "Burma Shave 3000" billboard rendered.

Also, in this same dialog: If you're running a dual-CPU system (or one of the new Intel Hyperthreaded CPUs), make sure that the Multi-Threading box is checked. You want both your CPUs to be working hard on your rendering!

## DON'T RENDER MOVIES—RENDER FRAMES!

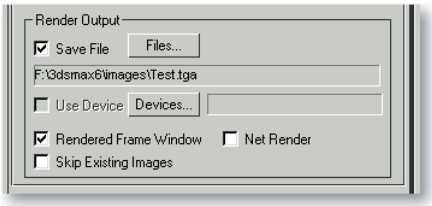


Unless you're rendering a small-format, quick-and-dirty test animation, get into the habit of always rendering image sequences (.TGA, .TIF, .PNG image formats; avoid using .JPG unless you simply don't have the hard drive space) rather than movie files (.AVI, .MOV, and so on). Image sequences have many advantages over straight movie formats.

First, unless you originally used a lossless codec for your sequence, you'll experience an ugly drop in image quality when playing back your animation. Second, if you render an enormous (that is, memory-hogging) animation, you have to load that entire file into RAM to play it, and if you have a slow graphics card, it will run like a turtle dipped in caramel. Third, it's tougher to do any kind of post-production manipulation, such as compositing, on a movie file—especially an .AVI that uses a lossy codec—than an image sequence. (You can't save an alpha channel for compositing in an .AVI file.) Fourth, if you encounter an error during the rendering, or 3ds max or your computer crashes, you'll lose only the unrendered frames, rather than your entire animation (which might get corrupted during the crash). Fifth, your mom said no, and she doesn't want to have to tell you again!

The bottom line is, don't render movies; render image sequences, and load them (or resize them as necessary, upon loading) into the RAM Player. (Choose Rendering > RAM Player from the main 3ds max toolbar.)

## SPEAKING OF THE RAM PLAYER…

Now, despite all the previous warnings, what if you're just itching to create a movie file? (Maybe you want to distribute your animation over the web or on CD.) You can do this in several ways in 3ds max. You can create Windows Media Player .AVI or QuickTime .MOV files from an image sequence by loading the sequence in the RAM Player and then clicking on one of the Save Channel buttons. From there, you can choose to save to any other file type.

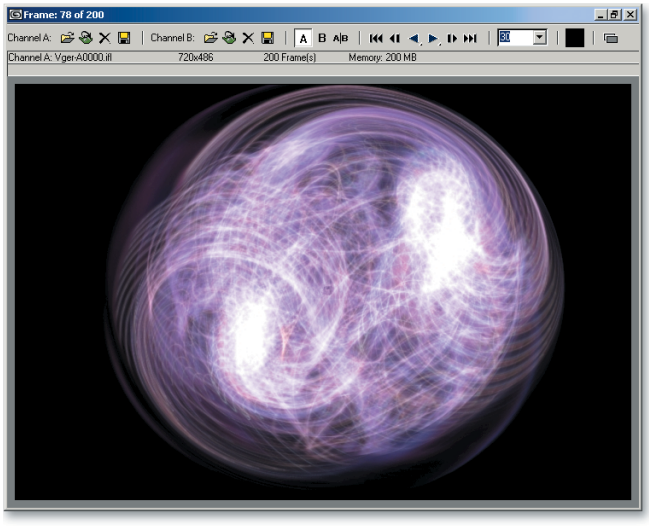


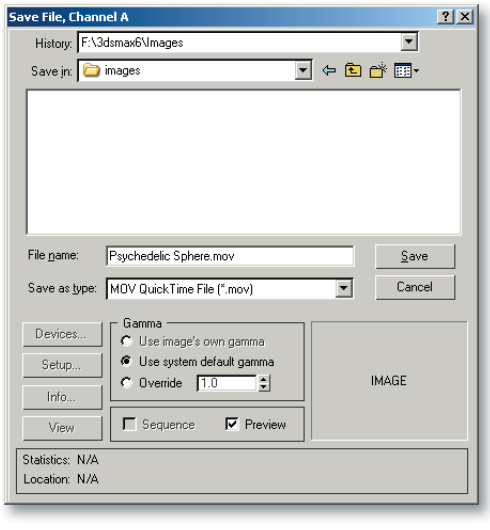


## RERENDERING FROM THE ENVIRONMENT BACKGROUND

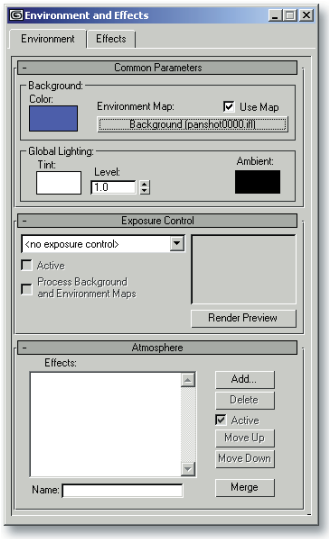


If using the RAM Player seems too easy, and you want to make your life more complicated (well, only a tiny bit), you can load an existing image sequence into the background of an empty 3ds max scene file (using Screen Coordinates, not Spherical Environment; otherwise, it will look wacky) and then render it out again from any viewport. Just choose a different format, and you're ready to go.

## OKAY, SO YOU REALLY WANT TO RERENDER YOUR ANIMATION…

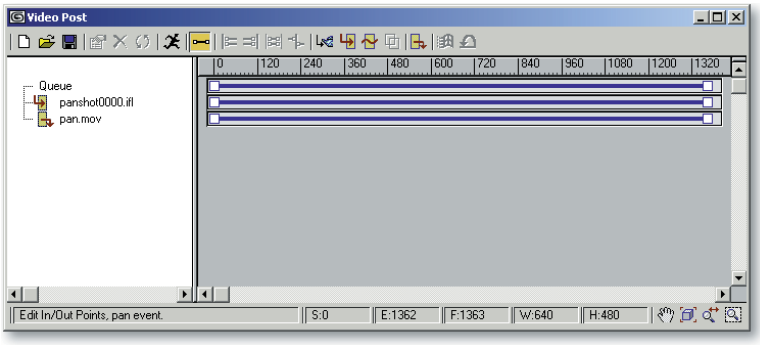


All right, continuing the previous discussions: Say that you've rendered an NTSC D-1 (video) video resolution (720 × 486) image sequence or animation, and you want to resize it for web playback (half-size, like 360 × 243).

In 3ds max, go to Render > Video Post, load the image sequence as an image input event, choose your new output size, and then add an image output event. Save the file in the format of your choosing.

Or, if you have Discreet's Combustion compositing program (or Adobe AfterEffects, or another compositing program), you can simply load the image sequence and rerender it (as done previously) to the size and format of your choosing.

## IFL = IMAGE FILE LIST

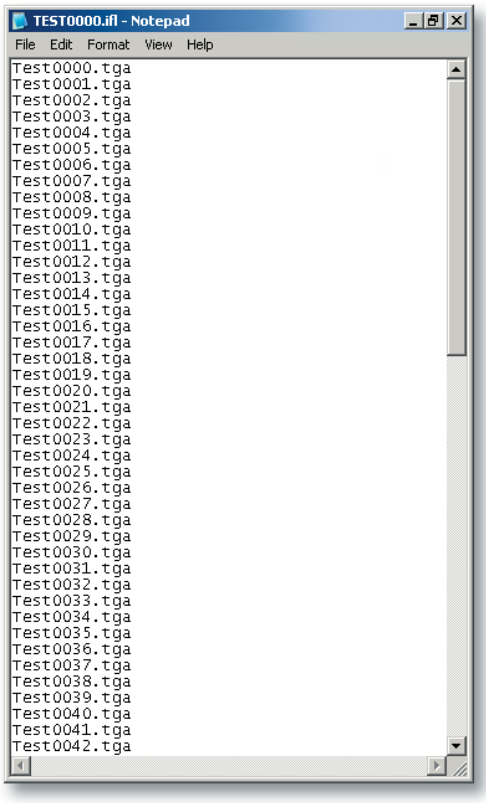


When you pick any frame of a sequentially numbered image sequence in the 3ds max program's File > View Image File, Video Post, an Environment background, the Material Editor, or elsewhere in the program, 3ds max automatically creates an image file list, or .IFL file. (You have to have the Sequence box checked in the relevant dialog; otherwise, you'll load only the individual frame you pick.) An .IFL is simply a text file that lists each frame of the sequence.

For example, if you had previously rendered a 100-frame sequence, starting on frame 0, called Test.tga, and you picked the first image to load as a sequence, then a file called something like TEST0000.ifl would be created in that same directory. If you opened this file in a text editor, it would simply be a listing like this:

Test0000.tga

Test0001.tga

Test0002.tga

Test0003.tga

… (and so on, until you get to the last frame, which follows)

Test0099.tga

(Note: You don't have to choose the first frame in the sequence; any frame in the sequence will do, as long as the Sequence check box is checked.)

## YOU CAN'T CREATE .IFL FILES ON READ-ONLY MEDIA!

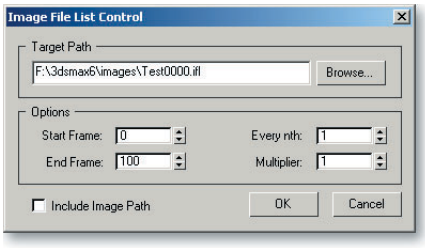


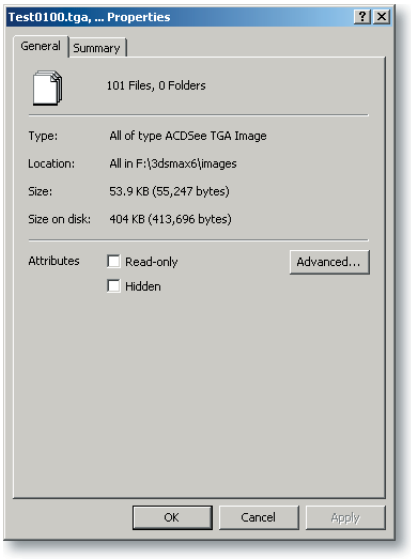


By default, the 3ds max program attempts to create .IFL files in the same directory in which you choose a frame of the image sequence. Consequently, you get an error message saying that the program can't create an .IFL file if you try to pick a frame from a CD-ROM folder, for example. (The program can't write the .IFL file to the CD.)

That's no problem. When you go to create the .IFL file, change the target path in the Image File List Control menu to a local hard drive path, and make sure that you check the Include Image Path box as well.

If you have sufficient hard drive space, you can also just copy the image sequence from the CD-ROM to a hard drive folder, and then create the .IFL file there. (Note: If you want to resave or overwrite the copied images after they're on your hard drive, you need to change their properties so that they're not still set to read-only. Just go into Windows Explorer, select all the files you need to modify, right-click and select Properties, and uncheck the Read-Only box.)

## MANIPULATING IMAGE SEQUENCES USING .IFLS

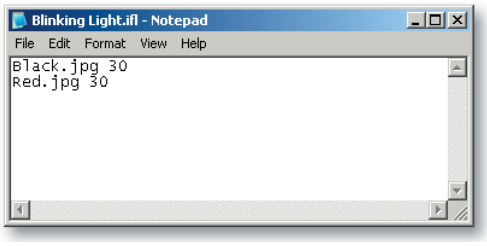


Image sequences and .IFL files are useful for many things in 3ds max: prerendered animated backgrounds, animated bitmap textures, and so on. They're also useful for allowing you to "time" events in a prerendered image sequence to something else happening in your scene.

For example, you could create a "blinking light" material that consists of, say, 30 frames of a black bitmap and 30 frames of a red bitmap, which would cycle on and off during the course of your animation. However, you could also simply use one frame each of black and red bitmaps by having your .IFL file include the duration (in frames) of each bitmap in the sequence, like this:

Black.jpg 30

Red.jpg 30

Placing the number 30 after the name of the bitmaps holds each bitmap for that number of frames in the sequence.

## CREATE NESTED .IFL FILES

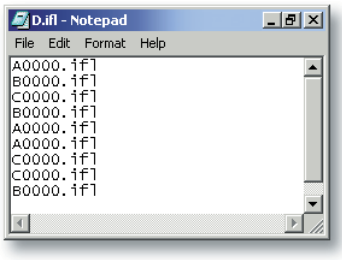


An .IFL file doesn't just have to be lists of bitmaps; it can also contain other .IFL files, for complex effects. For example, if you have multiple .IFL files for .JPG bitmap sequences A0000, B0000, and C0000, you don't have to have one giant .IFL file listing every frame of these sequences. If you've already created the individual A0000.ifl, B0000.ifl and C0000.ifl files, you could "gang" these together by creating a new .IFL (called D.ifl or whatever you want) that reads as follows:

A0000.ifl

B0000.ifl

C0000.ifl

This would then play each 100-frame sequence in turn. You could also randomize the order of nested .IFL files within the larger .IFL file, if necessary for your sequence.

## A CHICKEN AND EGG PROBLEM: HOW DO YOU SET UP AN ANIMATED BACKGROUND FOR A SCENE IF YOU HAVEN'T RENDERED THE BACKGROUND YET?

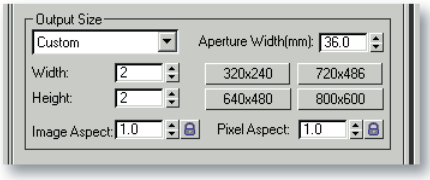


Here's a wacky scenario: Sometimes you might want to do a network rendering of several sequences in order, in which Scene A renders a .TGA bitmap sequence that you want (ultimately) loaded as an Environment background for the following Scene B. However, how do you load (or create) an .IFL list of an image sequence that you haven't rendered yet?

Well, one thing you could do is load an existing .IFL file (that has an equal or greater number of rendered frames as the new sequence you want to render) in a text editor such as Notepad or WordPad. Then just do a search and replace on the words/letters you need to replace with the new words/letters (to correspond with the Scene A rendered frame filenames), and resave the file with a different filename. You'll also have to make sure the frame count equals the frame count of your current sequences.

Here's a fast and dirty way to create the A bitmap sequence and its associated .IFL for the background of Scene B. First, reset 3ds max to an empty scene, set the frame count to be the exact amount as your A animation sequence, set your render resolution at $2 \times 2$ pixels (I'm not joking), and then render a series of blank (black) frames with the A sequence filename. The incredibly tiny frame size makes this black "proxy" sequence render incredibly fast. When you're finished, load Scene B, and load this proxy bitmap sequence as the background (which creates a properly named .IFL). Now, resave the Scene B file, and then load Scene A and render the "good," final A Background bitmap sequence, overwriting the tiny $2 \times 2$ pixel files you just rendered. If you're doing network rendering, you can then load Scene B and queue it up *after* Scene A; when Scene B renders, it renders with the new A Bitmap Environment background.

## COULD YOU MAKE THINGS MORE COMPLICATED, PLEASE? (WHAT ABOUT USING A COMPOSITING PROGRAM?)

Yes, the previous example was pretty verbose and complicated for such a simple trick. And yes, most 3D professionals would say that ultimately you're probably better off rendering all your sequences as separate passes (with alpha channels) and then compositing them with a program such as Discreet's Combustion.

However, not all 3ds max artists can afford a full-blown compositing program (although Chapter 10, "Aftermarket Accessories," offers some inexpensive substitutes). Consequently, tricks such as those listed previously are necessary if you want to break down your scenes into more easily managed layers and use 3ds max as your final compositor. Just remember: You need to make sure you have your sequences and their respective (proxy) backgrounds set up properly in your network rendering before you commit it; otherwise, you'll get a "Missing External Files" warning, and that sequence will fail.

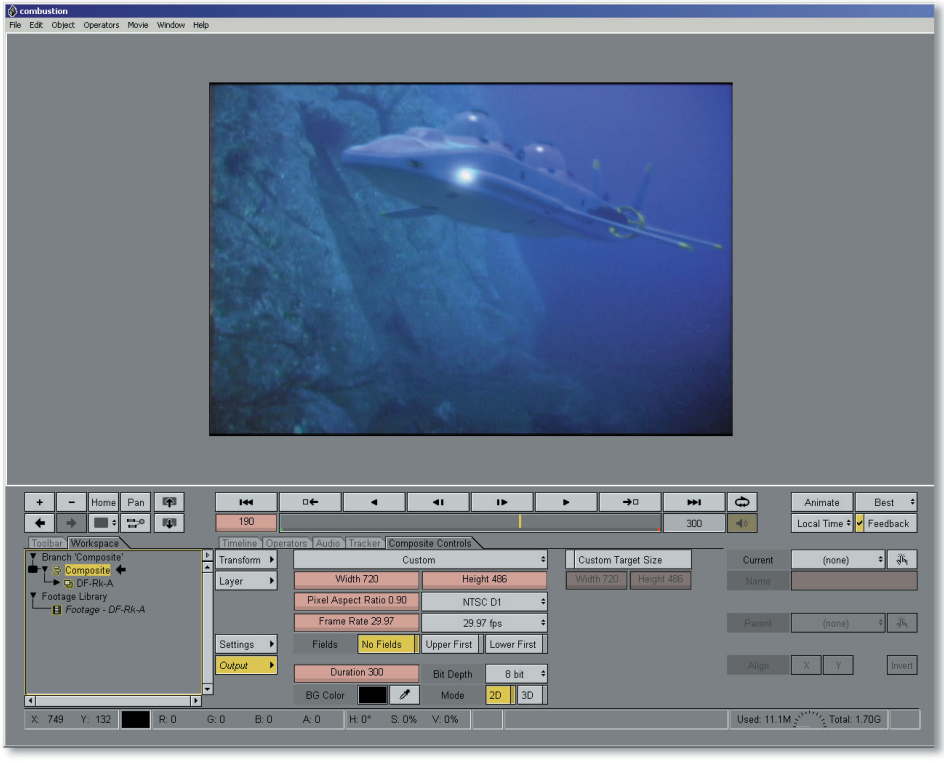


## COMPOSITING USING VIDEO POST

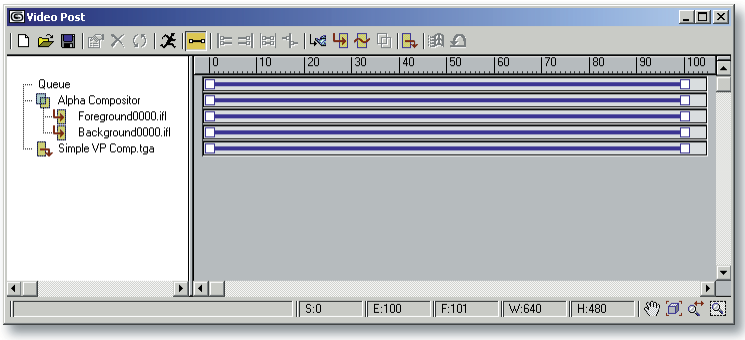


Okay, even if you don't have the cash for a commercial compositing program, you can still do simple alpha composites using the venerable 3ds max Video Post feature. (Discreet occasionally muses about removing this long-neglected feature from the program, yet there are still battle-hardened 3ds max veterans, as well as newbies, who scream "No! Over my dead body!")

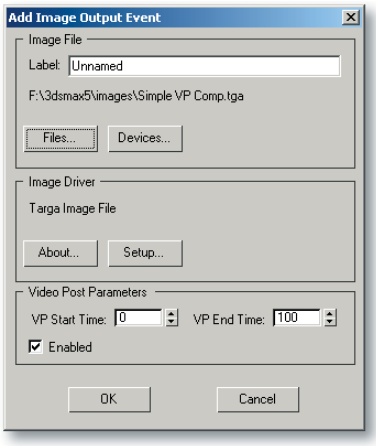


Here's how to do a simple alpha composite in Video Post: First, make sure you have your rendered elements (such as background and foreground elements) rendered with proper alpha channels. (You should have rendered them as .TGA, .TIF, .PNG, .RPF, or .RLA files.) Although the background images don't have to have an alpha channel, foreground elements most definitely *do*. Go to Rendering/Video Post, and click the Add Image Input Event button. Click the Files button, and then browse to the appropriate directory and load your background image sequence first. Then repeat the process with the foreground element. (Sorry. You can do only two elements at once with this technique.) When you have both elements loaded in the Video Post queue, Ctrl-click on each element to select both of them, and then click on the (now active) Add Image Layer Event button. Choose Alpha Compositor from the menu, click the Add Image Output Event button, set your final filename to the appropriate composite name, set your render resolution parameters, and then render. (Note that nothing can be selected in the queue if you want to choose the output event.) If you want to layer more elements on top of this composite, just repeat the process, and make sure you render your composite to an image format with an alpha channel, as above.

## RENDERING WITH SCANLINE MOTION BLUR: MULTI-PASS AND IMAGE

Motion blur on fast-moving 3D objects, as well as depth of field (DOF), help increase the realism of your 3D renderings and give a sense of scale, if used appropriately. Image motion blur simply "smears" a fast-moving object's motion across the frame. Image motion blur is fast, but it works only on linear motion, so it's not effective for things like rotating propellers and fan blades. Object motion blur samples an object's motion through the frame, and then composites multiple "slices" of the object on a per-frame basis.

The 3ds max Camera Multi-Pass motion blur (set in the Modify panel of a selected camera) simulates what a real-world camera does and gives the best results if you're using the 3ds max scanline renderer. In the real world, a camera's shutter is open for a period of time, information at an infinite number of intervals is exposed onto your film, the shutter closes, and another photo is taken. Consequently, everything in your photo gets blurred, including shadows, hair, and other elements that tend to pose complex rendering problems in the 3D world. The disadvantage to this technique when done in the 3D realm is speed; the camera effectively has to do multiple renderings of a single frame, with a moving object's motion "sampled" during the course of the frame.

There are workarounds for this technique, however. If you don't mind compositing your scene elements, then render your scene in layers (as mentioned earlier), and use Multi-Pass motion blur just on the spinning objects in that element. Then use (fast) image motion blur on other object layers.

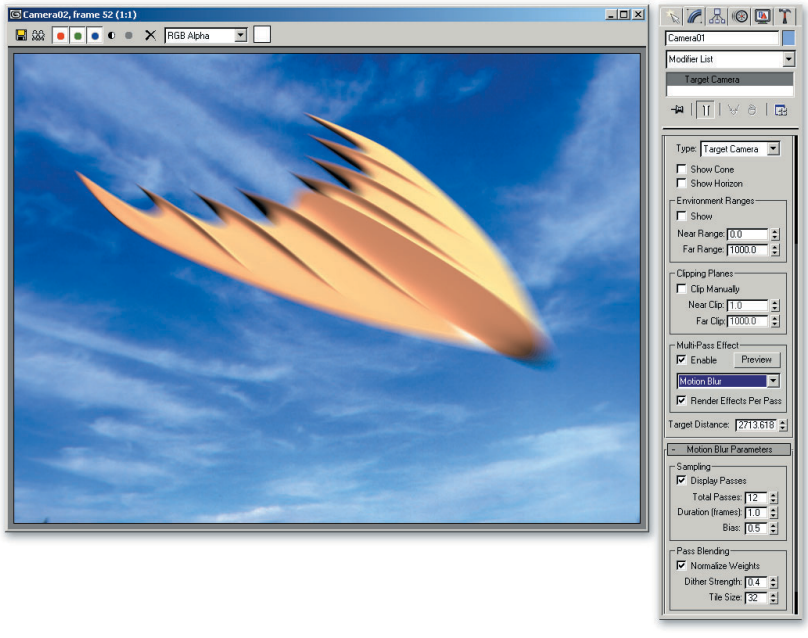


## RENDERING WITH SCANLINE: MIX IMAGE AND MULTI-PASS MOTION BLUR

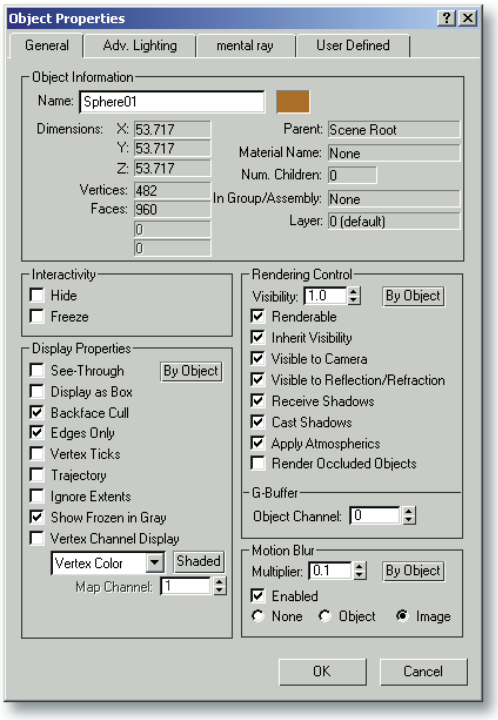


If you really want to improve your 3ds max scanline motion blur effect (and don't mind a hit on rendering time), then look into combining the different types of motion blur. You can combine Image and Camera Multi-Pass motion blur in the same render pass.

For example, if you want the camera shutter to be open for 0.5 (frames), set the Camera Motion Blur Parameters Duration to 0.5, and then set your desired object(s) Properties > Image Motion Blur Multiplier to 0.1. This gives you five sub-frames of accurate camera (or scene) motion blur, and then the Object Image motion blur with blur between those sub-frames.

## RENDERING IMAGES FOR PRINT: TEACH THOSE PRINT FOLKS A LESSON (OR TWO…)

It never fails—I'll have someone in the print industry want me to supply him with a 3D rendered image for a magazine or a poster, and when I ask, "What resolution do you want it?" he responds, "Make it 300 DPI." If this happens to you, first, resist the urge to smack the person. Second, explain that DPI—dots per inch—is simply a way of subdividing (measuring) a fixed, final value (your 3D rendering), *not* the actual measurement itself. (It's like asking, "How far is it from here to the next town?" and the person responds, "55 miles per hour.") Arrghhh!

Instead, tell the client that you need to know both the DPI *and* the final print size. For instance, if that person is running an image that's going to be 8 inches × 9 inches, at 300 DPI, then (in general) you should render your 3D graphic at a figure of 1 pixel equals 1 dot. Multiply 300 DPI by this final print size, and you'll get a rendered image that's 2400 × 2700 pixels. That's what you should render and deliver to the client (along with your invoice asking for immediate payment).

## RENDERING IMAGES FOR PRINT: OH YEAH, ANOTHER THING…

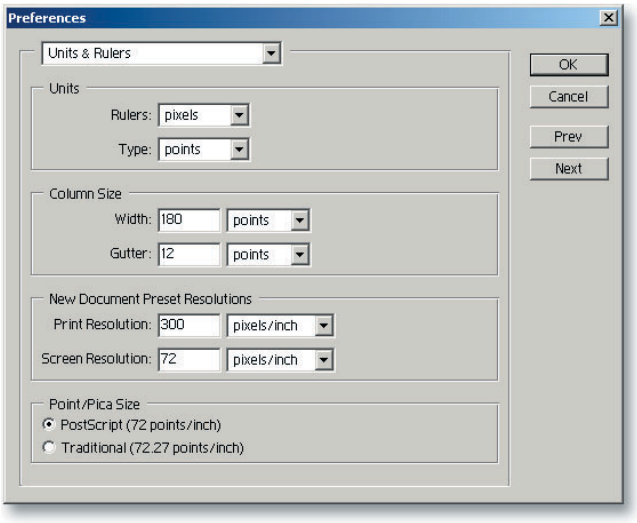


Adobe Photoshop, the *de facto* image-editing standard for all these print folks (and a lot of 3D artists as well) defaults to assigning a figure of 72 DPI (screen resolution) to any image it loads. This causes a lot of print people immediate heart attacks when they load your 3D image and discover it's not "300 DPI," and then they call you up saying, "It's the wrong size!"

Again, as in the previous tip, resist the urge to smack them, take a deep breath, and explain that they can change their screen DPI setting in Photoshop (release 7.x) to whatever their little heart desires by going to Edit > Preferences > Units & Rulers, and making sure that Screen Resolution is the same as Print Resolution. (In Photoshop 7.x, the default new document preset resolutions are 300 pixels/inch for print, and 72 pixels/inch for screen.)

## "THOSE PRINT PEOPLE"—MAKE IT EASIER ON THEM WITH 3DS MAX 6

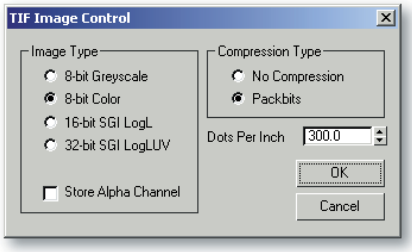


I'm going to beat this dead horse one more time: Most final render image formats that 3D artists use don't have a thing to do with DPI; they're just output to whatever pixel resolution you set them at (such as 720 × 486 for a standard NTSC video rendering).

However, that's not true of .TIF (Tagged Image File Format, or .TIFF for the Macintosh crowd) files; you can embed the actual DPI information in them, which will delight "those print people" mentioned earlier. In 3ds max 6, when you're rendering images for print, make sure a.) You render them as .TIF files, and b.) When you set the .TIF preferences, set the dots per inch to whatever the print clients want. Doing so embeds the correct DPI information in the rendering and warm the cockles, ventricles, and other parts of their hearts.

## THE PRINT SIZE WIZARD (ENOUGH WITH THE PRINTING STUFF ALREADY!)

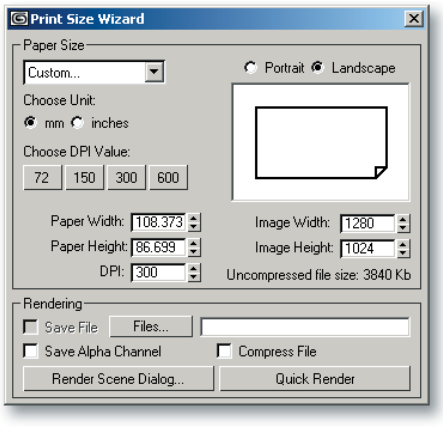


Okay, here's one final thought about outputting your 3D renders to print. 3ds max 6 comes with a dynamite new feature: the Print Size Wizard. Go to Rendering > Print Size Wizard, and a cute little dialog appears. Here, you can dial just about any render-to-print settings you want. Set your output to portrait or landscape, choose the unit measurements (millimeters or inches), choose DPI values, paper width and height, and image size. Simple, yes?

## MENTAL RAY IS IN THE BUILDING!

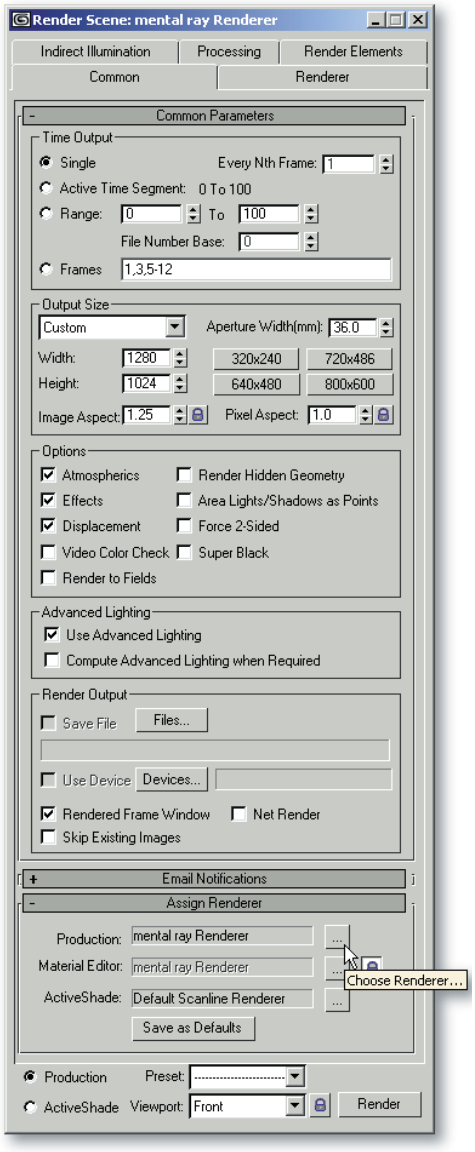


As of the release of 3ds max 6, a node-locked (single PC) version of mental images' acclaimed mental ray 3.2 renderer is included as an alternate to the 3ds max scanline renderer. (Chapter 3, "Waxing the Finish: Materials Tips," mentions this as well, but in case you're skipping around in this book, here you go!) Mental ray is noted for the quality of its global illumination and ray tracing. It's been used by a number of large special effects houses on cinematic visual effects for such films as *The City of Lost Children*, *Fight Club*, and *Star Wars Episode II: Attack of the Clones*.

To choose mental ray as your production renderer, just go to Rendering > Render > Common, and under the Assign Renderer section of the rollout, click the Choose Renderer button and select the mental ray renderer. Note that mental ray works with all core 3ds max materials, but it also includes its own material and shading types, which you can pick from the 3ds max 6 Material Editor.

## MENTAL RAY IS ON THE COUCH!

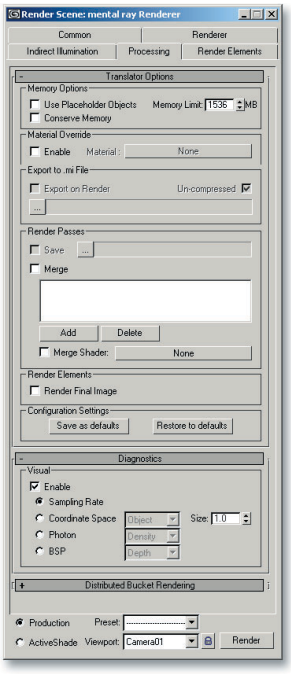


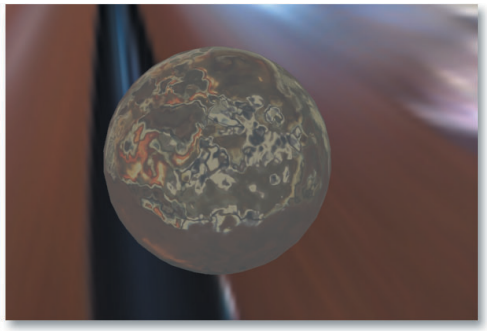

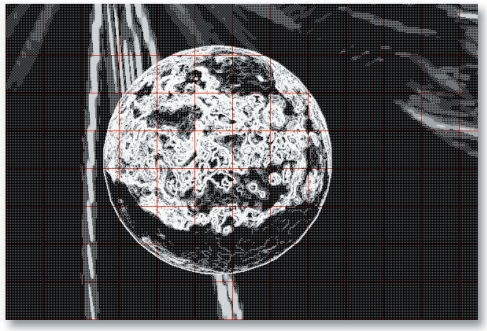

Do you get tired of tweaking your mental ray settings over and over again but not getting exactly the results you're looking for? You can help diagnose your problems by using the aptly named Diagnostics rollout. Just pick mental ray as your production renderer (see the previous tip); then, in the Render menu, click the Processing tab and under Diagnostics: Visual, check the Enable box.

Turning on Diagnostics: Visual gives you an invaluable tool to optimize and tweak your render. You can choose to diagnose your pixel sampling rate (for antialiasing problems), coordinate space, photon emission, or the BSP (Binary Space Partitioning) Tree—all the areas in which your rendering times can be short and sweet or be stretched out until the heat death of the universe.

As an example, to adjust your photon settings, open your scene, make sure you have mental ray as your assigned renderer, and then go to Rendering > Render > Indirect Illumination and make sure you have your photon generators set. (Check the Global Illumination Enable box.) Go to the Processing tab, and under the Diagnostics section, enable Visual and set it to Photon. Now when you render, you'll see a "false color" representation of your scene, showing areas with dense photons as red, and areas that are lacking photons as blue. You can then use this visual guide to tweak your photon count and photon radius until your diagnostic render shows a smoother range of color, with red in the most important areas of your scene. When you're ready to do your final render, turn off Diagnostics: Visual, and you're all set!

## USE MENTAL RAY'S IMAGE SAMPLING WISELY

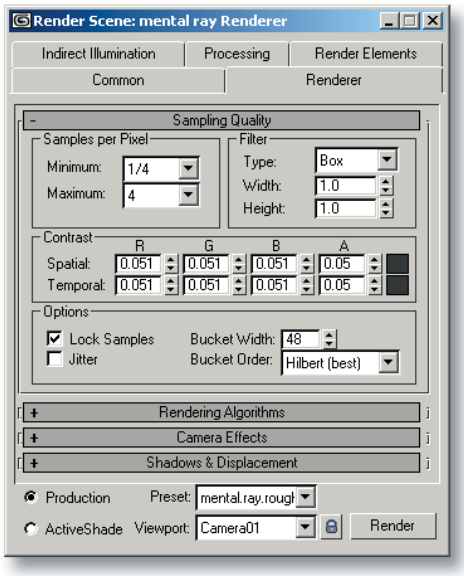


Scene sampling for antialiasing can be a complicated task for any renderer, even a sophisticated engine such as mental ray. If you're just doing preview renderings (to test your lighting, for example), you can speed up your workflow considerably by mental ray sampling only what you really need.

If you're simply working on your overall scene composition and lighting, leaving the mental ray Samples per Pixel Minimum/Maximum values at their defaults costs you unnecessary time. The default values (Minimum 1/4, Maximum 4) mean that mental ray will take 4 samples per 1 pixel in areas with many details, or 1 sample per 4 pixels in areas with fewer details.

If you set both Minimum and Maximum Samples per Pixel to 1/64, mental ray is forced to only take 1 sample for each "bucket" of 64 pixels. This creates a coarse rendering, but it renders at incredible speed and might be just what you need to get your overall look correct before you commit to production-level sampling settings. Just save this setup as Render Preset (on the Rendering > Renderer tab, at the bottom of the dialog box). Call it "Rough Preview" and keep it handy for future use.

## IN MENTAL RAY, CONTRAST CAN SAVE YOUR DAY!

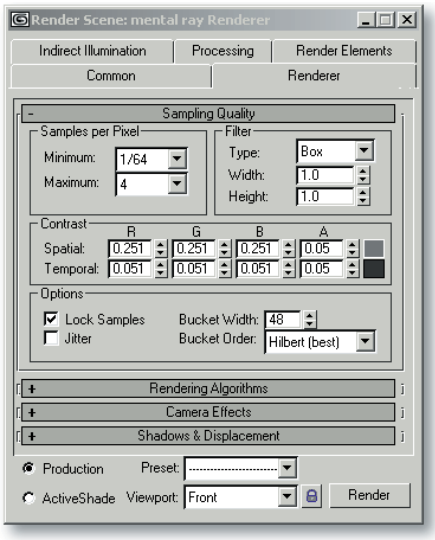


Have you tweaked your mental ray pixel sampling and noticed that you still seem to lose important details, even with a high Maximum Sampling per Pixel setting? You notice that you're losing fine details (such as wispy cables or antennas in an architectural rendering), but at the same time, your render times seem to be going through the roof—and you don't want either of those to happen!

You can fix this and render even highly detailed scenes without too much hassle by memorizing two words: Contrast settings. (Actually, you shouldn't just memorize them, you should use them.) The Contrast settings are located on the mental ray Renderer > Renderer tab, right below the Samples per Pixel values, and their defaults are set 0.05. Contrast settings determine how to "weight" the Samples per Pixel values—toward the Minimum or the Maximum values you have set. If your fine scene details get chewed away in the rendering, it's usually not because your Maximum sampling rate is too low, but because your Contrast settings are too low.

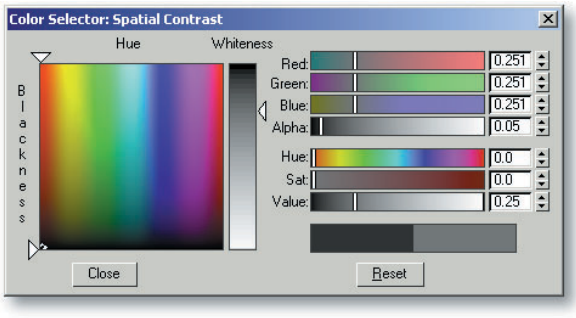


To fix this, keep your existing Samples per Pixel Min/Max values within reasonable limits. (There's no need to go above Minimum 1/64, Maximum 4 for very detailed scenes; for most scenes, even Minimum 1/16, Maximum 1 should work fine with finals.) However, if your scene details show rendering artifacts with these settings, increase your Contrast values slowly, working toward a lighter gray. This triggers the Maximum Sampling values earlier and should fill out those fine details. (Note: The Contrast Spatial settings are primarily for still renderings; adjust the Temporal settings if you need to tweak fine details during an animation.)

## HIDDEN LINE RENDERING: RENDER TO VECTORS IN MENTAL RAY

So, your client wants her 3D model rendered as a hidden-line rendering (in Vector format, no less) for print reproduction, or for the web? (And she wants it today!) No problem. Just go to mental ray's Rendering > Render > Renderer menu, and under Camera Effects, enable the Contours switch. (Note that you'll have to reassign your scene materials using one of mental ray's Contour shaders, but that shouldn't be a big deal. Just adjust your line width and color settings until you get the effect you want.)

Then, in the Camera Effects rollout, click on the Shader button in the Contour Output slot and choose Contour PS as the new shader. Click-drag this slot into an unused sample slot in your Material Editor, and make it an instance. In the Material Editor, you will now be able to adjust your output settings, and then you can just type in your final render path for the vector output, with the Adobe Illustrator .EPS format as the suffix (for example, C:\3dsmax6\images\test.eps). Click on Render again, and you'll produce an .EPS vector file that you can open in Adobe Illustrator or Macromedia Flash without problems. (Note: You might not see rendered results in your 3ds max Rendered Frame window, but you can load the final .EPS file into a program such as Adobe Illustrator or Photoshop, and adjust it there.)

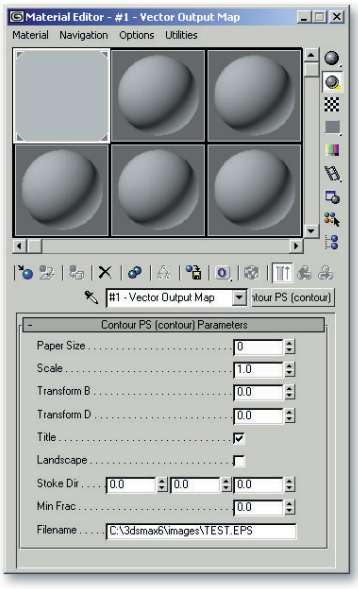


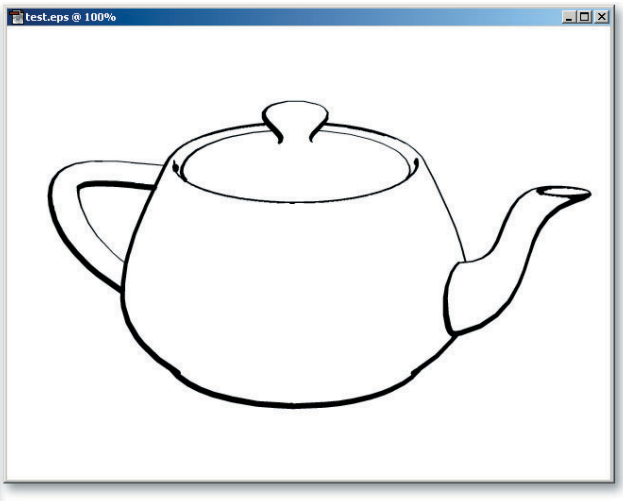


## TRIM YOUR (BSP) TREE IN MENTAL RAY

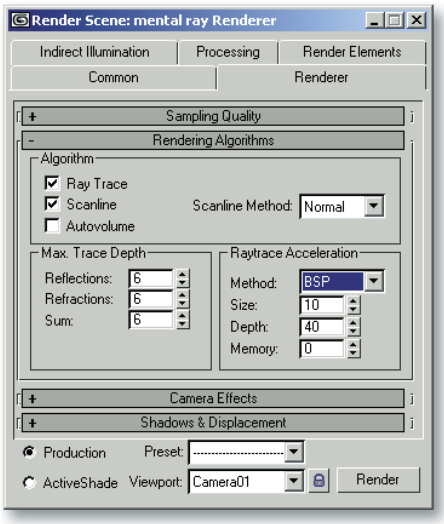


In mental ray's Rendering > Render > Renderer tab, look at the Rendering Algorithms section. Under Raytrace Acceleration: Method, there's a drop-down menu where the default is set to BSP. A BSP (binary space partition) "tree" basically is a method of categorizing your whole scene with the help of an imaginary sub-divided bounding box encompassing your scene. This is necessary because it helps mental ray cast rays much faster in this structure than without.

The subdivision cells of this bounding box are called voxels, and the creation of those is what you control with the two settings Size and Depth. Size is the number of cells along one side, so the default would divide your whole scene into a maximum of $10 \times 10 \times 10$ cells. Depth is the parameter that tells mental ray how many subdivisions of one cell should be allowed. The default of 40 allows a maximum of 40 subdivisions in voxels, with a lot of faces.

That's the theory. In practice, you're looking for a good balance between the creation of your BSP tree (the "idle" time before the actual render starts, when the grid is built and subdivided) and the render time of your image. Increased BSP values can cause a long scene render preparation time but result in a lightning-fast render; conversely, decreased settings result in the preparation/render time equation being reversed. The time difference might not be seconds; it might be much longer.

Depending on the complexity of your scene, pay attention to your BSP settings and experiment with different Size and Depth subdivision values. Due to the nature of the default "bounding box" effect, you might be better off tweaking these settings from scene to scene until you get the results you want (or at least, the results that you can live with).

## MENTAL RAY PREFERENCES: PLEASE LEAVE ME A MESSAGE

You can both display and save your 3ds max 6 mental ray rendering information (in an exhaustively verbose format) by going to Customize > Preferences, clicking on the Mental Ray tab, and then checking pretty much every box in that menu. (Save your .LOG file to your 3ds max 6 \scenes folder and call it something like mr-renders.log.) Then go to Rendering and click on the mental ray Message Window item. You'll get a window that displays a whole heap of rendering information (useful for diagnosing potential problems) when you then render with mental ray.

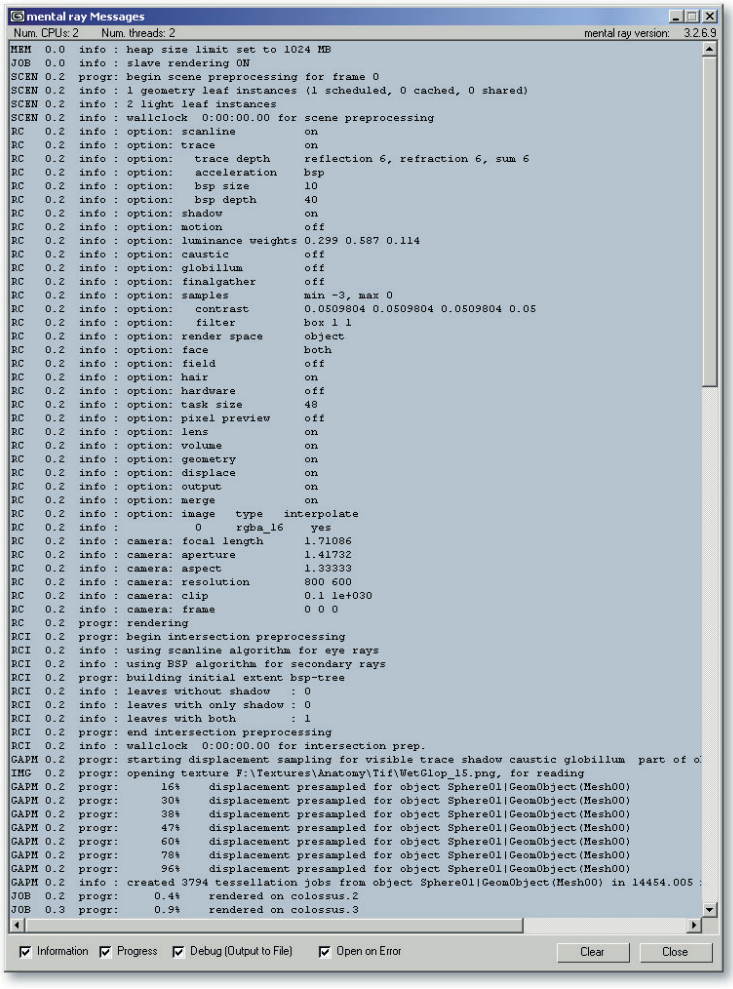


## RENDERING AESTHETICS: OUTER SPACE SCENES

Although multiple combinations of lights (key light, fill light, and "kickers," or edge lights) might work in most 3D scenes, sometimes harsh, unidirectional light (which 3D programs excel at creating) is actually the aesthetic you want. Moody, film noir-ish lighting is one, but the most common unidirectional lighting scheme for 3D artists/science fiction fans (let's face it, there's a lot of overlap there) is when you're rendering outer space scenes. Generally, you'll light a spacecraft model with a single spot or directional light, and little or no fill light, unless it's passing by a planet, an illuminated space station, or a nebula. For increased realism, you could use a colorful nebula background in your scene as the rationale for tinted fill light on your spacecraft models. Other details, such as self-illuminated windows and "practical" lights on the model (making it look as if it has beacons illuminating its surface, like the refitted Enterprise from *Star Trek: The Motion Picture*), increase the realism and scale of your scene.

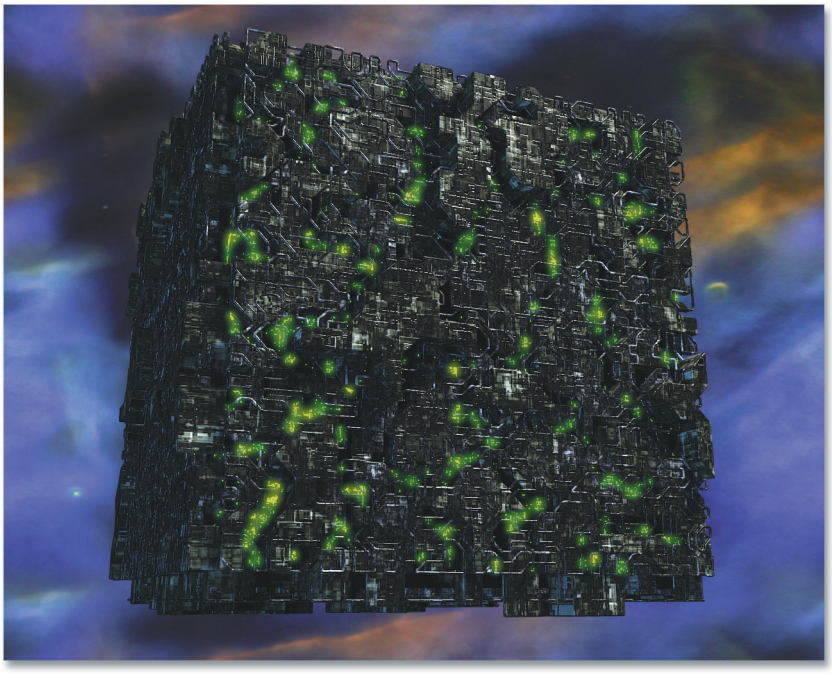

## RENDERING AESTHETICS: UNDERWATER SCENES

To render realistic underwater scenes, you should duplicate the murkiness and light qualities that are inherent in the environment. To mimic these atmospheric conditions, go to Rendering > Environment > Atmosphere, and add fog, volume fog, or volume light to your scene. You should determine the scale of your scene (are you in the ocean, or just in a swimming pool?), adjust your camera's Environment Ranges settings, and then your Fog settings to match the amount of murkiness you need. To create the illusion of depth drop-off in the ocean, apply a Gradient Environment Color Map (such as light cyan to darker blue to navy or black) to tint the fog as well.


## RENDERING AESTHETICS: DISTANT LANDSCAPES

For more realistic outdoor scenes, especially if you're seeing a distant horizon, you should always add a slight amount of atmospheric haze (Rendering > Environment > Atmosphere > Add > Fog). If you look at a distant mountain range (or if you don't have one right outside your window, just grab a travel magazine or pretend), you'll notice how colors become muted and washed out with distance. You can use just a slight amount of atmospheric fog (it depends on the scale of your scene), and the colors will determine the clarity or quality of your "atmosphere." For clear outdoor settings, using a slight bit of white fog is desirable; for sunset or urban settings (where the air might be more polluted), a slight yellowish or reddish cast makes your horizons look better.

Finally, experiment with adding a Gradient Environment Color Map to your fog so that your skies are more realistic. Make your skies darker at the top and lighter on the horizon, as if the sun is rising or setting.

## RENDERING AESTHETICS: STILL LIFE AND MACROPHOTOGRAPHY

The smaller the scale of your scene, the more you should use a shallow depth of field (DOF). In real-world photography, a tiny camera tends to have a correspondingly small ability to capture "distant" objects in a scene—and if you're simulating microphotography, "distant" might be only a fraction of an inch away. The more objects in your scene go in and out of focus relative to their nearness/distance to the camera, the smaller the scene will look. (If you use shallow depth of field in a room interior, it will resemble a dollhouse shot with a snorkel camera.)

Both the 3ds max scanline renderer (through the Camera Depth of Field Parameters rollout) and the 3ds max 6 mental ray renderer offer DOF options; use them carefully when you're rendering your scene, especially because this is a computationally intensive effect. (Also, be forewarned: If you combine depth of field effects with atmospheric effects, raytracing, or global illumination and caustics in mental ray, your render times will go through the roof.)

Note that if you do have a compositing program, such as Discreet's Combustion, you're probably better off rendering your scene to different render layers, and then adding animated Blurring effects (to simulate DOF) as a post-process. (That's how the cinematic special effects experts do it, so you should, too!)

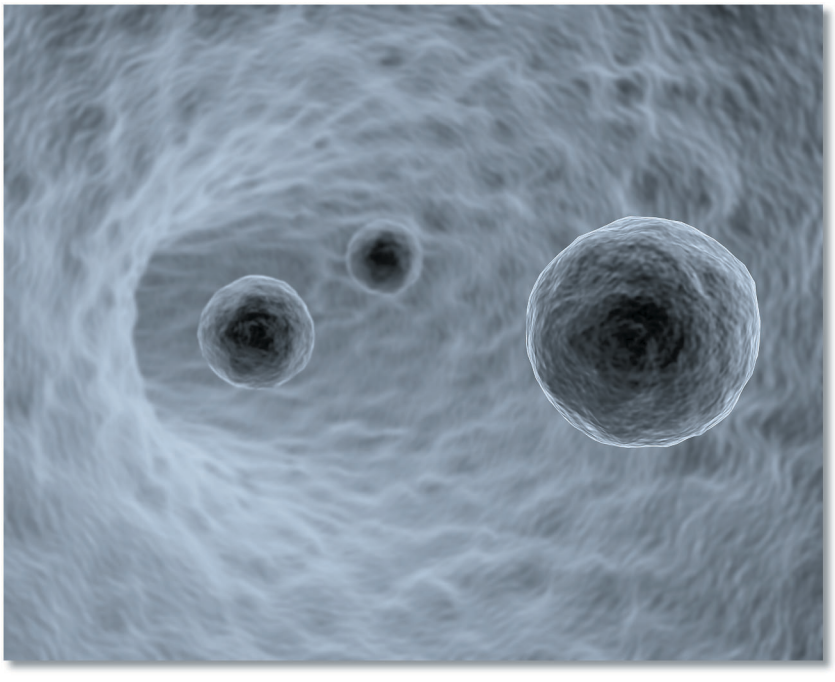

## "HELLO… YOU'VE GOT RENDER!"

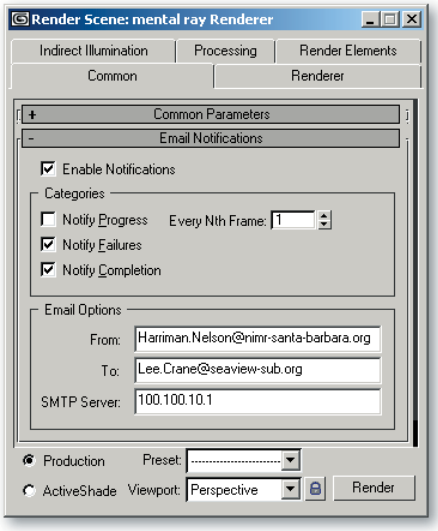


Okay, so you've set up several 3D sequences to render, and you're going to walk away from your computer (or render farm) for a while. If you can't resist the urge to check on your renders (even while you're away from home), and you have an always-on Internet connection, you can have 3ds max actually contact you to let you know what's up. Just go to the Rendering > Render Scene > Common panel, and open the Email Notifications rollout. There, you can enable Notifications and set up various categories of email alerts detailing general frame progress, failures, and/or completion of the job. All you have to do is put in the correct From and To email addresses and the address of your SMTP server, and you're good to go. If you have a cell phone or a personal digital assistant (PDA) such as a Palm Pilot or Blackberry that receives email, you can get the status of your renders while you're on the road, in a restaurant, or even in the middle of a romantic interlude.

"Honey, you have the most beautiful…."

BEEP, BEEP!

"Whoops, gotta go—my rendering of the Enterprise attacking the Death Star just finished!"