

**ARTICLE: Choosing the Best 3D Rendering Application for Your Needs (New for 2005)...**

{Last Updated: 02/27/2005 — Renderer-specific parts of this thread article are entirely new!}

**CHOOSING THE BEST 3D RENDERING APPLICATION FOR YOUR NEEDS (NEW FOR 2005)**

You hear it all the time: "**Which 3D program should I buy?**" Nearly every artist has explored this question for his or her own needs—and those who haven't, should. What you learn might surprise you, as all too often the better question to ask is "**Which renderer should I buy, and then which 3D program works best with that renderer?**"

Misinformation about 3D apps and their renderers is surprisingly prevalent, and incorrect advice is regularly propagated from one artist to the next with such fervor that biased rumor is soon erroneously viewed as industry fact. This article will help separate fact from fiction—from program tools to renderer functionality and quality.

**FIRST, A FEW FUNDAMENTAL MISCONCEPTIONS EVERY ASPIRING ARTIST MUST UNDERSTAND WELL:**

**Which 3D Software is the Industry's Best?** Almost every 3D artist has heard this one: "Blah-blah program is the best! All the top artists use it." Anyone who suggests that one particular 3D software application is the clear-cut winner and best solution overall is someone whom you should promptly ignore. Dismiss them as religious zealots. It's amazing how artists go to such lengths to defend their beliefs and choices in 3D software, as if to convince themselves as much as others. Moreover, interestingly it's often the least talented artists who are most adamant in preaching about a particular 3D program's superiority. This is politics and religion for the artist, and online debates rarely prove more than who is the most tenacious or talented speaker. The 3D industry is too diverse for any particular program to consistently shine above all others. Now, "Which software is best for specifically doing \_\_\_\_\_" is indeed a question closer to having a clear-cut answer, and luckily this answer is easier to discover with a little online research.

**The Artist or the Tool?** A professional Hollywood stunt driver will maneuver a sloshy Chevy sedan through a tight obstacle course far better than a novice driver does in the seat of a tight Porsche. 3D rendering is no different: a 3D industry professional can pick up a \$99 off-the-shelf rendering package and produce better art than the no-talent hobbyist can with a \$15,000 highend integrated software solution. Of course, just as the Hollywood stunt driver could employ a Porsche to far greater effect than a sloshy Chevy, so too can a talented 3D professional in front of a highend software tool. Therefore, while it's important to realize that natural artistic talent is your most important prerequisite in reaching 3D sainthood, having the wrong software tool can slow you down, limit your expression, and ultimately put you at a competitive disadvantage compared to other well-equipped artists.

**Judging 3D Software by Art You've Seen Done With It.** Never do this because in reality you're judging the artist's talent and not really the software. Worse, there's no way for you to know how long those artists may have fought with and tweaked their renderer's settings nor if all the tweaking was even limited to the 3D software itself (they may have done a lot in Photoshop). Also, some of the most fantastic images come from artists who've honed their skill by working for major studios where particular software was already in place for reasons not necessarily having to do with quality nor the features you may require for your particular work needs. Finally, the most fantastic image art stills you've ever seen in your life say absolutely NOTHING about how good that renderer will be at producing animations where antialiasing flicker and texture or geometric edge artifacts could suddenly 'come to life' leaving that fantastic output entirely unusable for TV or film due to poor quality not discovered until animation time.

**Software Used for Movie FX.** It's important to understand that not everything you see in the way of special effects in movies is necessarily 3D.

For example, in *Star Wars Episode 1: The Phantom Menace* with exception of the high-speed pod race shots and a few other planet-side and high-speed space motion shots, all of the starships you saw in the Ep1 movie were physical models; NOT 3D. Same with Starship Troopers and most of the film versions of Star Trek. The reason is that the physical models still look better, but this is in the process of changing now as we saw starting back with *Star Wars Episode 2: Attack of the Clones*. Also, when a software manufacturer markets its program as "used in \_\_\_\_\_ movie", they often neglect to inform you that their software might have only been one tool of many and not necessarily the primary tool. For example, their program might have been used for only one minor FX shot in a movie full of hundreds of other 3D FX shots. Worse, some companies are simply better at marketing than other companies. Alias (makers of Maya) is quite infamous for its highly-effective ad campaigns. Finally, the biggest and best studios often use particular 3D programs with ability to translate well to a preferred renderer technology and a total solution pipeline, and not necessarily because of the value of that 3D software's toolset itself. In other words, studios usually begin with a rendering pipeline and then port to it with whichever 3D software applications are best-suited to translate.

Also, long-standing industry artists may prefer one program over another because they probably 'grew up' on SoftImage and the predecessors of Maya because that's all that was available back then for professional use. Why would these artists learn a whole new package when they don't need to? Also, large studios customize their 3d programs in ways you are not able to (thanks to entire teams of programmers) in order to produce the toolsets that actually get used in movies—meaning you're judging their custom code rather than the off-the-shelf program tools.

**You Get What You Pay For with 3D Software.** Yes and no. Highend software was over \$50,000 per workstation seat only a few years ago. With today's Intel-based PC architecture now outperforming all other 1-, 2-, and 4-way CPU workstation hardware solutions, software pricing gaps have closed tightly as dictated by market demand. So is Discreet's 3ds max software a better performer because it costs nearly double what Alias' Maya software does? Don't be too quick to judge in favor of either; what you learn might surprise you. Why would Alias lower Maya's pricing from over \$7,000 to just under \$2,000 in just one year alone? There is no clear-cut answer, but software must be priced to sell—perhaps aggressively as a marketing strategy—and it must also prove itself as a good total value buy. Therefore, you must consider pricing-to-worth by ballpark figures only. In this case, anything between \$1,500 and \$5,000 is roughly comparable, though some applications are clear-cut winners in specific fields and for specific duties. **Renderer technology, however, happens to be at this time rather close to a 'price equals quality and value' ratio.**

**Computer Hardware:** What's Fast and What's Not. When purchasing a workstation and determining which platform to learn 3D on, look to the present (and to some extent the future); not the past. Learn to separate what was fact only a few years ago from what holds true today. Not long ago, SGI hardware was tops in the 3D field; nothing could touch it. Today it's the slowest solution imaginable and a dead-end career choice. Even huge FX houses such as Industrial Light & Magic, Sony Pictures Imageworks, and Pixar who have untold millions of dollars invested in SGI hardware pipelines are aggressively switching to Intel-based PC machines. Most of those aforementioned studios have almost entirely completed their move to PC technology. Today, Intel/AMD technology is a clear-cut winner in sheer speed-to-cost ratio. The fastest scalable server 'workstations' are still Sun Microsystems and SGI supercomputers, but these have price-tags so high that 3D software is no longer practical on them. Macs are awesome machines and excel in many areas—but not 3D rendering speed. They come in a distant second, and even Steve Jobs (CEO of Pixar as well as Apple Computer) knows that and has directed his Pixar Animation Studios to employ Linux on Intel-based rendering machines instead of Macs for Pixar film renders—though in the front office for all non-3D needs (HR, Marketing, etc.) they exclusively use Macs, of course. Also, remember that in hardware your CPU and memory translate to your workstation's rendering speed whereas your video card (and to a minor extent your CPU) handle your onscreen viewport performance. That is, with few exceptions a fast video card does nothing for final rendering speed. Understand the difference.

### **RENDERING: YOU NEED MORE THAN JUST 3D SOFTWARE; YOU NEED THE RIGHT RENDERER!**

Most 3D artists produce rendered results: the final bitmap images used in web art, print, television shows and commercials, and in motion picture film. Not all 3D careers require this, however, since a video game modeler might strictly produce geometry for a game engine and thus have no need to render it. For everyone else, the following is the best-kept secret in the 3D industry:

#### **There is an enormous difference in the rendering quality you require for rendering a still image versus rendering an animation.**

There is also a further enormous difference in the rendering quality you require for producing high-contrast high-detail photorealistic results versus simple smooth shapes, soft shaders, and toon animations.

Don't jump to conclusions about how great or poor a renderer is that you're testing until you've properly learned how to use it:

For example, slapping a displacement map on a test object in 3ds max and rendering it blindly with Mental Ray may yield atrocious results—notably in rendering speed. This isn't a flaw in Mental Ray, it's that you're now exploring a much higher-end renderer than you might otherwise not be used to using and not yet have a proper grasp on the controls of that renderer. In this example, Mental Ray's displacement architecture uses a powerful and innovative approach far different than that used in most native built-in renderers shipping with Lightwave, 3ds max, Maya, etc., and one must understand Mental Ray's association between the renderer's global Maximum Displacement value, the global Edge Length, whether it's View (pixel) or World-Unit based, and the strength of extrusion in the map settings for your particular object (Maya calls these MR-connection settings by different names). Having an object-independent extrusion strength value far smaller than the renderer's global Maximum Displace value will up render times to ridiculous levels. If you didn't know that, you might jump to the erroneous conclusion that Mental Ray is slow at displacement.

In Photorealistic RenderMan, you might jump to the false conclusion that it's no faster at rendering heavy geometry scenes than most other renderers simply because you might have built and sent to it scenes full of complex polygonal geometry with high polygon face counts designed to achieve smoothing—not knowing that you could have had even smoother results at a tiny fraction of the rendering time if only you had built your scene using true SubDs or NURBS objects because of how fast and smooth PRMan's micropolygonal renderer can process those. In this case, you might not have understood how to take full advantage of PRMan's strengths; how it can render billions of virtual polygonal faces at any resolution in mere seconds providing a perfectly smooth surface that would have otherwise required enormous scene detail with hardwired polygonal geometry that'd take much longer to render and with poorer surface smoothness and inferior quality. So go forth and explore, but do it with an open mind, intelligence, and understanding.

Most 3D programs come with native integrated renderers (usually hybrid renderers). Some come bundled with professional highend renderers, while others require separate purchases entirely. The differences between renderers are tremendous and **choosing the right renderer for your 3D needs is THE MOST important decision you will EVER make in 3D**. This point cannot be over-emphasized enough, especially if you intend to produce TV or film-quality output. In fact, the differences between various highend software programs out there (e.g., Maya, 3ds max, Lightwave, SoftImage XSI, Cinema 4D, Houdini, and the like) are essentially insignificant compared to the importance in choosing the right renderer (and the right connection/bridge/translator between that renderer and your 3D software).

Therefore, this article will treat Rendering as its own distinct subject separate from 3D Programs and the differences in their features.

### **UNDERSTANDING RENDERER IMPORTANCE IN YOUR WORK**

**NOT IMPORTANT** = Non-rendered output, such as video game models. *Focus all of your attention on program toolsets that enable you to model quickly and accurately to meet your modeling needs.*

**SOMEWHAT IMPORTANT** = **Any still image. Nearly any renderer (including free integrated native renderers built into the 3d programs you buy) will produce beautiful still images, usually indistinguishable from the same image produced by any other renderer.** There are of course some minor differences in some abilities, for example the ability to produce sub-surface scattering effects like how light diffuses and scatters inside a wax or marble object, which some renderers support and others do not. Most are capable of producing complex ray-traced and global illumination effects which may be slow to produce but fine in quality for a still image. Generally, this is simply a features-based need. Render times can vary, but the speed difference isn't necessarily critical to workflow. *For this level, focus most of your attention on program toolsets and abilities, with only moderate attention to the renderer you wish to use because all renderers will produce*

*essentially the same result.*

**MODERATELY IMPORTANT** = Simple Animations involving mostly smooth shapes, mostly shader-based imagery, and toon animations—all at television or lower quality. Here enters the need to set sampler levels high enough for quality Anti-Aliasing (to avoid pixel flickering in animations) without driving up rendering times too much. Where the difference of 5 minutes versus 15 minutes in producing a frame isn't a huge factor if you only need one still image, times that by thousands of frames needed for a brief movie shot and quickly it will make or break a studio. *Focus at least half (preferably more) of your attention on renderer **quality and speed**—NOT features lists—with your remaining attention to program toolsets and workflow.*

**CRITICALLY IMPORTANT** = High-contrast and/or high-detail geometry or displaced surfaces at any resolution, and any animations at film quality resolutions. Here, the speed while maintaining quality issues become even more important while also being joined by the crucial need for efficient computer memory management when dealing with enormous scenes. **As shocking as this may sound, most renderers are absolutely incapable of producing necessary results for this type of rendering output.** *For this type of need, Focus essentially ALL of your attention on choosing the right renderer. For the most part, the only thing you should be concerned about features-wise in any highend 3D program toolset is the extent of its translation link to the renderer you've chosen (how completely the 3d program communicates with the renderer). Last on your list will be the evaluation of program tools and other features goodies, because the difference between 3d program toolsets is trivial in comparison.*

## RENDERERS: WHAT'S OUT THERE AND HOW THEY COMPARE

Only a few highend renderers are discussed below, and that's because those discussed here are the only renderers capable of meeting a CRITICAL renderer need, such as quality film-resolution animations.

For example, depending on who you talk to, Maya's internal built-in renderer is rated anywhere from outstanding to outright crap—and that's because interestingly **those same people are generally speaking from either still-image output experience versus animation output experience**, though rarely do people grasp or realize this difference in perspective. Here, Maya's internal renderer can and regularly does output fantastic still image artwork, sometimes at film resolutions and with great ray-traced quality. Persons using it for this still image need often rate it as a great renderer. On the other hand, anyone producing photorealistic (read: high-contrast; high-detail) television or film output in a studio environment knows that Maya's internal renderer is perhaps the worst and least usable of all program-native renderers sold on the market today. Indeed, attempting to derive the needed quality in it free of motion-generated anti-aliasing flicker artifacts is a no-win scenario bound to induce insanity in any artist forced to use it. The same can definitely be said for 3ds max's internal renderer. Lightwave, the only real exception as far as native renderers go, is discussed below along with other highend renderers.

Remember, when reviewing renderer capabilities bear in mind that these are made from an animation perspective—especially for television and film quality. Most native renderers (3ds max's, Maya's, etc.,) are more than sufficient for producing even the most elaborate highly-detailed quality still image artwork—just so long as it isn't animated into motion.

## A CLOSER LOOK AT EACH INDIVIDUAL HIGHEND RENDERER

**PRMan (RenderMan):** When combined with a 3D program via strong translation, this is one of the best world-class solutions available. PRMan (Photorealistic RenderMan) is the actual rendering software produced by Pixar, whereas 'RenderMan' is just the specification that PRMan adheres to. There are several RenderMan renderers on the market—of which PRMan is the best, most well-known, and the one made by Pixar—so in the strictest sense it remains incorrect to use the term 'RenderMan' to refer to an actual renderer program. PRMan is the renderer of choice by large studios such as ILM (Industrial Light & Magic), Sony Pictures Imageworks, Digital Domain, and many others, and with good reason. The output quality is astonishing and the RenderMan specification provides for countless tweaking controls in order to adjust your rendering speed vs quality trade-offs to the precise levels you require, not to mention its massive extensibility.

**PRMan's STOCHASTIC SAMPLING:** Perhaps Pixar's greatest patent and edge over all other renderers is its exclusive method of implementing Stochastic sampling, a patented form of 'Monte Carlo' sampling. Sampling, of course, is well known in 3d as the method of anti-aliasing/filtering bitmap images both going in (a texture map) and coming out (the final rendered image). **Aliasing in an animation is the annoying pixel flicker indicative of amateur rendering.** Stochastic sampling helps reconstruct a frequency pattern in a texture map with fewer samples than regular sampling methods do but at equal error probability, and thus is faster. For the same render speed, other renderers will experience aliasing (pixel jitter during animations) requiring higher adaptive sampling rates be set by the user at render time and hence much longer render times to overcome incorrect signal reconstruction (aliasing) to achieve the same anti-aliased quality of PRMan. In effect, PRMan behaves exactly like the receptors in our eyes by using a Poisson distribution pattern of point samples to achieve noise instead of aliasing whenever the Nyquist limit (half your signal frequency) is exceeded:



**PRMan's HIGH-ORDER SURFACE RENDERING:** PMan was built from the beginning with core optimizations for NURBS and SubD surfaces. In fact, when it was first developed it was incapable of rendering polygons! **All renderers in the world are polygon renderers, but PRMan is foremost a curve renderer that generates micropolygons (polygons the size of screen pixels or smaller) at render time that follow the NURBS' mathematical surface curvature in screen space and therefore result in all high-order surfaces appearing perfectly smooth no matter how close the camera gets!** And it can generate these high-order surfaces faster than it or other renderers can handle polygonal geometry. **This means when rendering NURBS or SubDs you never have to worry about or specify tessellation settings because nothing is getting tessellated!** This aspect of PRMan's REYES (Render Everything You Ever Saw) architecture is unmatched by any other renderer on the market.

**PRMan's STOCHASTIC 3D MOTION BLUR:** 3d stochastic motion blur is true motion blur which means that objects moving with overlapping Z-depths under different vectors in 3d space all produce accurate motion blur to the camera. PRMan does this fast and at flawless quality justifying the many patents Pixar holds for this technology. Other renderers such as Mental Ray can also produce highly-accurate 3D motion blur, but doing so typically increases render times by at least 500% whereas PRMan's stochastic blur adds very little to render times. **This is where PRMan quickly makes up for any speed losses in areas it's not so adept in (such as ray-tracing).** 2D motion blur, of course, is very fast to render by other renderers but is plagued with quality problems the most notable of which occurs whenever an object passes in front of or behind another moving motion-blurred object.

**PRMan's SUB-PIXEL DISPLACEMENT:** True displacement is another of PRMan's strongest features. In any PRMan surface, because each pixel is one or more micropolygons, you can assign displacement maps that result in extremely fine surface detail **with only a very minor hit to rendering speed** (fast). Other renderers slow to a crawl when subdividing surfaces into greater numbers of triangles in order to accommodate displacement maps because they must use more conventional geometric subdivisions instead of more or less 'just pushing the pixels'. In PRMan, rather than increasing your polygon counts and render times you simply specify the amount of shading (texture sampling) for the displacement image in order to control how sharp and accurate the displacement is. That is, **displacement is handled no differently as far as render times go than having just applied another texture map to your object.**

**PRMan's DISCRETE SAMPLING CONTROLS:** Other renderers provide a single set of controls for sampling textures and final image rendering, usually as an adaptive minimum/maximum number of samples to process dependent on perceived image contrast and complexity. Some even provide control overrides on a per-shader basis. PRMan takes this a step further in a revolutionary approach: **You can control the Shader Rate separately from the Pixel Sampling Rate.** This means if you have basic materials with simple or no textures but very detailed highly-complex geometry then you can set the Pixel rate very high so that very fine geometry doesn't alias without wasting time supersampling the rest of your scene or materials. Furthermore, Pixel sampling (geometric sampling) is extremely fast. On the other hand, if you have very simple geometry but complex texture maps with many fine lines and high-contrast pixels then you can increase the Shader sampling by lowering the Shade Rate while leaving the pixel rate lower. This allows the artist to optimize rendering times in cases of extreme geometry or extreme map detail.

**PRMan's MEMORY MANAGEMENT:** PRMan is capable of handling HUGE texture maps via mipmapping (multiple scaled map sizes as required for various scene elements based on size and distance to the camera so that distant objects call only small maps to conserve memory since larger maps aren't needed until closer to the camera), virtually unlimited complex shaders, and over a billion unique polygon faces in real-world production scenes without fear (not to be confused with render tests involving instanced or copied geometry totaling a similar 1 billion polygon faces which is significantly easier to render).

**PRMan's SCALABILITY and EXTENSIBILITY:** PRMan used to possess only an extremely extensible architecture favored by technical directors at large studios because of how well it fit into large studio pipelines by providing for custom shaders, custom object code, and distributed rendering methods. The downside was that PRMan was very involved and difficult for the home hobbyist and small studio user to use. For example, even when paired with Maya via Pixar's own Maya translator (MtoR; RAT), Maya's hypershade nodes were useless because shaders had to be hand-written or assembled graphically in the external application SLIM and then attached to Maya geometry for PRMan. **Today, however, PRMan's scalability now includes the hobbyist and small studio as well via the RenderMan for Maya plugin version which is every bit the quality of the full professional studio version (RenderMan Pro Server) just sans the big studio extensible architecture** (for example, there is no RIB interface allowed for the plugin version). The plugin version even translates Maya's material nodes from hypershade perfectly, supporting Maya better than RAT ever could, all the while providing the same high-quality sampling controls, displacement, fast stochastic motion blur, and more.

**The BAD:** Connectivity with 3d applications other than Maya is a serious problem for any version of PRMan. Via the plugin renderer, the connection to Maya is superb. If you happen to use another 3d program, your options are very limited indeed. Worse, even with those few options you are further limited to RenderMan Pro Server which is significantly more expensive than the plugin version and scene translation is simply not good.

**Mental Ray:** Previously packaged almost exclusively with SoftImage, mental ray (correctly written as a lower case proper noun, though we'll be using upper case for reasons of sentence start clarity) has been around for a long while but its architecture has kept pace with the times. Mental Ray produces outstanding quality in rendered images, but is not the fastest renderer out there—especially at tasks such as producing Global Illumination at quality while fast enough for animation rendering, though it still outperforms PRMan in this area. SoftImage XSI's bundled connection with Mental Ray is the best in the business, offering access and control to nearly every feature this renderer provides. Mental Ray's plug-in translator for 3ds max v5 (and earlier) was previously a serious embarrassment and essentially unusable. However, starting in 3ds max version 6 Mental Ray is bundled directly with 3ds max and has a surprisingly strong translation that nearly rivals SoftImage XSI's connection to this superb rendering technology. Maya, though also bundles Mental Ray, had one of the worst translations (prone to crashes and lacks features) between it and this renderer which was corrected back in Maya v6. Mental Ray's core architecture is even more feature-packed and flexible than PRMan's, and its .mi2 specification is as fully scalable as PRMan's RIB description. MR's photon-based features and overall scene quality capabilities are very mature and extremely powerful and accurate—widely acclaimed as **the most physically accurate ray-tracing and photon lighting** found in any renderer in the world. Displacement subdivision, while not a true sub-pixel 'push', is still handled entirely at render time and is very fast compared to 3D programs where displacement is dealt with at the geometric level prior to translating or enumerating the scene for rendering. Displacement still runs into physical memory limitations when very fine mesh tessellations are called for to match complex sharp-edged map approximations. Motion blur is 3d object-based and very accurate, but is extremely slow. When Motion Blur is needed (which is typical), Mental Ray loses its faster-than-PRMan rendering speeds and can quickly fall behind in an overall production speed. It is typical to see render times increase by more than 500% when motion blur is turned on. Rapid Scanline option attenuates this somewhat, but easily introduces quality issues such as dreaded aliasing. Like PRMan, Mental Ray requires a strong working knowledge of its tweaking controls in order to maximize speed-for-quality trade-offs to the precise levels as required by each particular project, though strong well-integrated translators now found in SoftImage XSI, Maya, and 3ds max make Mental Ray essentially a 'plug-

n-play' renderer with little need to have a Technical Director (TD) on staff. Depending on your production needs, Mental Ray can be every bit as production-useful (perhaps more so in some areas) as PRMan is, but PRMan continues to outshine Mental Ray in many areas such as motion blur speed, sampling quality at speed, fast displacements, and more.

**Brazil R/S:** A newcomer, Brazil is a special plugin renderer currently available only to 3ds max software users, but with future plans as a Maya tie-in. Brazil is a very high-quality renderer with outstanding (perhaps best in the business) global illumination abilities insofar as quality and speed go, and produces great caustics and reflection/refraction ray-tracing. Brazil's Global Illumination (skylight scene lighting in this case) is so fast while simultaneously operating at GI sampling levels sufficient for artifact-free results that **this GI can actually be used for production work**. Mental Ray's GI via Final Gathering is several orders of magnitude slower, though is more accurate. So while neither Mental Ray nor PRMan really produces GI fast enough for real-world use in production work, Brazil somehow manages to lead the way here, and with flying colors. And despite a lack of physical accuracy in Brazil's GI compared to Mental Ray's GI, the viewer would be very hard-pressed to notice while the 3D artist would likely find artifact-free GI too hard to pass up. Overall, Brazil can probably be most closely compared to Mental Ray given similar features sets. Brazil does lack any real motion blur, leaving the animator with basic 2D blur that is sometimes of little or no use, which for many rules out this renderer as a serious highend studio tool at this time. Similarly, Brazil's lack of integrated displacement shaders is a severe setback to the many artists who rely on that superior-to-bump-mapping technology. Finally, Brazil's pricing is aggressive and makes it quite attractive when compared to Mental Ray and especially PRMan. **VRay**, a renderer with many similarities and qualities to Brazil R/S, is another option though we here at ZAON have not tested it. VRay will soon be available for Maya as well, and many swear by it as fanatically Splutterfish fans do for Brazil R/S.

**Lightwave:** Lightwave's native built-in renderer is the only native renderer actually suitable for production film work, but its core architecture is definitely now showing its age. Lightwave rendering is slow compared to all of the other options available, and while it can produce wonderful results it has more shortcomings than the other highend rendering options—not necessarily surprising, given its age and native development. Lightwave's worst shortcomings are foremost its speed, closely followed by texture filtering, motion blur, and displacement quality. Still, the pricing (free with Lightwave and with **unlimited** network farming licenses) and tight integration with the core package are great strengths.

## CLOSING THOUGHTS

Whether discussing the advantages of SoftImage's non-linear animation workflow techniques to Maya's particle systems to 3ds max's reactor dynamics, the overall functional differences between these industry leaders is pretty minor if you step back far enough to see the whole picture. These minor differences give each program its particular flavor and style, and what's ambrosia to one artist is meaningless to another. Not until you compare one of the six industry leading apps to a mid-level solution like TrueSpace or a low-end one like Bryce 3D do these feature sets really distinguish themselves and define the term 'highend' which at this time applies only to these six leaders: 3ds max, Maya, Houdini, SoftImage XSI, Cinema 4D, and Lightwave.

Aspiring artists are best advised to explore their talent and output needs first, and then select a 3D application from there:

**Choose ANY of the 6 leading apps, based on your own personal preference, IF you're only going to produce still images for web or print art.** All of the 6 leading programs can produce amazing quality images even with their native renderers.

**Choose a Professional Highend Renderer first, and then choose the 3D app with good or great translation/connections with that Renderer IF you plan to produce photorealistic television or film-resolution FX.** Very few renderers can handle this type of duty, so choosing a 3D app before a renderer will only lead to later regret. For example, choosing 3ds max gives you many highend options such as free integrated Mental Ray, or plugin Brazil R/S, finalRender, VRay, and more to choose from. Going with Maya gives you free integrated Mental Ray, with options for the PRMan plugin renderer, full Pro Server PRMan, and soon others such as finalRender and VRay. Lightwave also has many new solutions to choose from, and Cinema4D appears to be hitting the market with a film-quality renderer from the start.

**If you have a particular studio career path in mind, choose the 3D app they use.** If your heart is set on working with Eden FX for producing visual effects for the next Star Trek TV series, then you had better learn Lightwave. If Industrial Light & Magic is the place you want to work, then learn Maya or SoftImage XSI. If painting exquisite digital mattes for movies at Matte World Digital is more your style, then stick to 3ds max. For game companies, find out which gaming platform works best with which 3D app, and then go that route. The list goes on and on, but if you research your specific career path carefully it's easy to learn which FX houses use which software programs.

**If you're going to do a "little of this and a little of that", including animation work, all mainly from a hobbyist's standpoint, then any of the 6 leading apps will work well for you.** Choose based on pricing and other factors such as ease of use, toolsets, and prevalence that you personally identify best with.

Good Luck! and, Good Times!