THE NEWSLETTER FOR LIGHTWAVE 3D ANIMATORS

LIGHTWAVEPRO

an Avid Media Group, Inc. newsletter

Volume 2 Issue Number 11 • November 1994 • $8.00

ScreamerNet Shakedown

Inside:

LightWave & PhotoCD

Lens Flare Mania!

All About Fog

What's on the LIGHTWAVEPRO Disk? See Page 15

Normal Render
This is the final rendered image from the normal render test. Image courtesy Amblin Imaging.
Copyright 1994 Universal Television and Amblin Entertainment

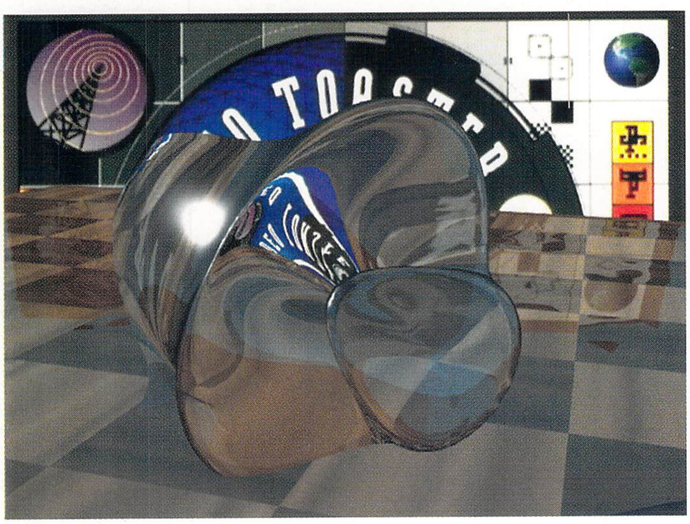

Nasty Render
This image is the final rendering of the nasty render test.
Copyright 1994 Grant Boucher

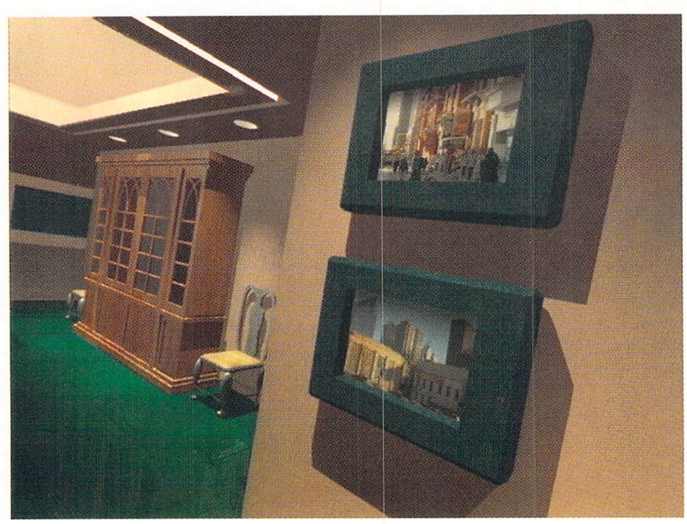

Hallway Closeup
Each 24-bit image in the picture frames only consumes about 200K of memory, yet looks sharp and clear as the camera zooms close to it.
Copyright 1994 Dan Ablan

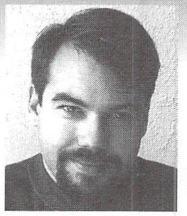

EDITOR'S MESSAGE

by John Gross

I am writing this month's "Editor's Message" from a hotel room in Waikiki. We are here doing the annual LightWave training in Hawaii (it's a tough job, but someone's got to do it). I would rather be lying on the beach at this point, but duty calls, and I'm a dutiful kind of guy. I like to think the information contained in *LWPRO* is timely, even if the newsletter might not always be ... I guess I'll have to take credit for that. It's not as easy as I thought it was going to be to get a bunch of writers to turn in their work on time. And of course, no matter how many writers turn in articles on time (you know who you are—I could count you on Homer Simpson's hand), there are always a few writers that straggle along (you know who you are). To top it all off, I'm usually the one that's the latest (just ask me how many "Dear John" columns I have due).

Anyway, now that I have cleverly made some excuses and hopefully laid a few guilt trips, I can't help but feel I have been slightly remiss in keeping you up to date with changes in *LWPRO*. By now, you've probably noticed the color covers (that's a joke). A few people have asked about the resolutions of these images. They are printed at 2640x2640 at 300 dpi, but only a few have been rendered at that resolution (the LightWave hallway on July's cover, for instance). A few were rendered in Print Resolution (like this month's cover) and one was even rendered in High Resolution (no, it wasn't June's—see if you can figure it out).

In the premiere issue of *LWPRO*, I said that each issue would cover a different LightWave topic. It took me about 9-10 issues to realize that was going to be a difficult order to keep filling. While topics for articles abound, it can be an arduous task to compile seven or eight articles about the same subject. With this in mind, I've gone to a more "open" style for issues where a number of different subjects are covered in one issue. There will still be issues devoted to entire topics (when I can think of them).

see Editor's Message, page 5

TABLE OF CONTENTS

LIGHTWAVEPRO
an Avid Media Group, Inc. newsletter

Editor..........John Gross
Managing Editor..........Jim Plant
Editorial Coordinator..........Douglas Carey
Art Director..........Helga Nahapetian Taylor
Art/Production Coordinator..........Kristin Fladager
Associate Editors..........Joan Burke
..........Corey Cohen
Contributing Writers..........Dan Ablan
..........Grant Boucher
..........Mark Glaser
..........Taylor Kurosaki
..........Glen David Miller
..........Mojo
..........John F.K. Parenteau
Group Publisher..........Michael D. Kornet

Editorial Offices: Avid Media Group, Inc.
273 N. Mathilda Avenue, Sunnyvale, CA 94086
Telephone (408) 774-6770; Fax (408) 774-6783
John Gross can be reached electronically at:
jgross@netcom.com (Internet); 71740,2357 (CompuServe)

Printed in the USA© 1994 Avid Media Group, Inc.

Are you interested in writing for LIGHTWAVEPRO or submitting images? If so, contact us at our offices or electronically.
Avid Media Group, Inc. its employees or freelancers are not responsible for any injury or property damage resulting from the application of any information in LIGHTWAVEPRO.
LIGHTWAVEPRO (Vol. 2, No. 11); (ISSN 1076-7819) is published monthly by Avid Media Group, Inc., 273 N. Mathilda Ave., Sunnyvale, CA 94086-4830. A one-year subscription (12 issues) in the U.S. and its possessions is $72 (U.S.); Canada/Mexico, $84 (U.S.); Overseas, $108 (U.S.). Allow 4 to 6 weeks for first issue to arrive. Second-class postage rate paid at Sunnyvale, CA and additional mailing offices. POSTMASTER: Send address changes to LIGHTWAVEPRO, 273 N. Mathilda Ave., Sunnyvale, CA 94086-4830.

About the cover: This month's cover image was created by LightWave artist Mark Glaser. Each piece of the knight was built with splines and can be patched for any resolution in similar fashion to the workings of outline fonts. Rendered using LightWave 3.5, the chain mail texture utilizes the new bump array procedural texture. The remaining textures rely on 1.61 megabytes of image maps created with DPaint and one digitized wooden fence image that's wrapped onto the barrel. Actually part of an animation, the scene displays several examples of secondary animation where bones are used to bend the feet and make the chain mail kilt undulate.

Replacement Animation

by Glen David Miller

Replacement animation? "Huh, what's that? I've never heard of that. What good is it?" Most people respond the same way when hearing this term for the first time.

Replacement animation is simply the replacing of one object for another or nearly similar object. To be a little more specific, an example of replacement animation would involve a scene showing a dripping area of water. This is a common sight after most rainstorms. First, the water flows to a low spot on the edge of a window frame. As the water collects, a small teardrop shape begins to form in the lowest area. Then, as the teardrop grows in size, its weight overcomes the surface tension. Suddenly, the teardrop separates from the main area of water and drips downward.

While it may be easy to imagine or even draw these four steps, in LightWave, it's not quite as easy. And without the replacement animation technique, it would be nearly impossible to accomplish.

While there are many more uses of replacement animation, our technique involves using four objects. The first is a flat plane with numerous faces that have been tripled. The second object is a form of the first object-plane. With the third object we will continue to finesse the drop shape until it reaches the point just prior to the break. The final two objects will complete the break. The first will resemble the first plane with a slight bump in the middle of the plane. We can reuse one of the above shapes for the plane. The second will be the teardrop shape that falls below us.

If at this point, the above explanation seems a bit much or even over your head, stick with me. I feel that one of the key rules every animator must follow is to study nature. While some might say, "The water drips," you and I both know that there are numerous specific steps that occur when a drop of water falls. An animator must know how to identify those steps and mentally or physically make a list of the useful steps he or she will need in order to complete the project. When, for example, Disney animators begin a sequence that involves wildlife, they study the specific creatures they will be drawing. Preplanning is a step too often overlooked by beginning animators. Get specific! You might even want to redraw the water drip sequence to become a little more familiar with what we'll be modeling, or even go make your own drip. Whatever the case, study, study, study.

On the technical side, in this article, we'll be covering spline modeling, metaforming, transforming polygons and points, and finally, we'll quickly cover some ideas behind our use of envelopes. Let's get started.

The Objects

We'll begin in Modeler. In the lower left-hand window, the front view, we need to create five points. As you can see in Figure 1, these points will make up the shape of our water plane and water drop. After entering the points, deselect all of them. Click, in order, the first and second points (again, see Figure 1) and then press Ctrl-p to make a spline. Re-select the second, third, fourth and fifth points and again enter Ctrl-p to make a spline. There is a specific difference in the two splines we have created (Figure 2). While the first is a straight line, the second is a curve. Remember, when spline modeling, two points that are connected become a straight spline-line. Three or more points form a spline-curve.

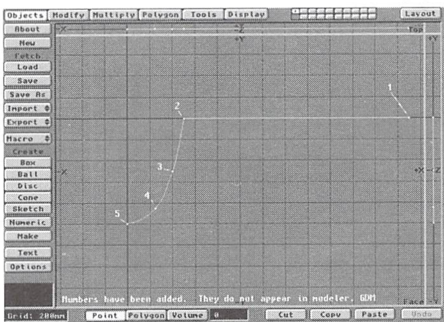

Figure 1

Make sure when creating two-point splines that you do not accidentally create two-point polygons. To illustrate the difference, go to a new layer and enter four points, then de-select them. Select two of the points and press Ctrl-p. Now, select the other two points and press (p).

They look pretty much the same, don't they? Under **Polygon** select mode, select both lines. The first difference to be noted is that the spline has a small "head" on one end. Also, when looking in the polygon stats window (w), you will see that there are two polygons, one face and one curve. Figure 3 shows all of the final possibilities you might run into. The first shape is our object with a spline-curve; the second shows multi two-point segments that have been smoothed (Ctrl-s) to form curved splines; the third shows all two-point splines; and finally, the last object shows two-point polygons or faces. The overall warning is this: a face won't patch and it won't smooth. A spline won't render. Remember to keep your p's and Ctrl-p's straight.

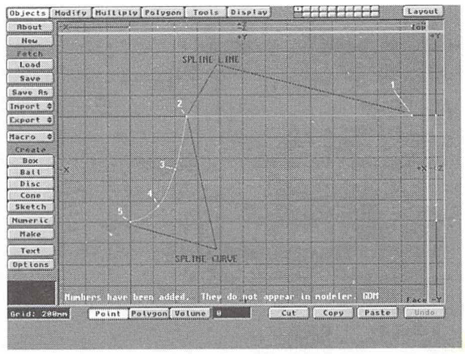

Figure 2

Go back to the original layer. Now that we have the basic cross-section of our object, go to the **Multiply** menu and select **Lathe**. Align your Lathe axis with the fifth point of our object running along Modeler's Y axis.

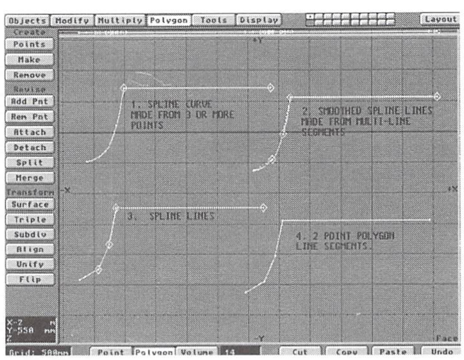

Figure 3

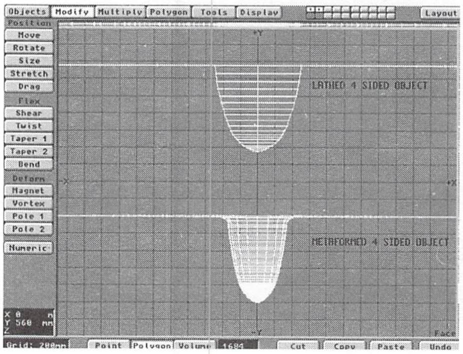

Figure 4

Select Numeric and enter 4 into the number of sections. Lathe the object. Our raindrop doesn't look too smooth, does it? Since the object is made up of quads, it is the perfect subject for Modeler's Metaform tool. Go to the **Polygon** Menu and select **Subdiv**. Under the subdivide menu that appears, select **Metaform**. After Metaform finishes working its magic, you'll see that the object is smoother. Metaform it again. See Figure 4 for the desired shape.

Save this object as "Drop 2." Copy the object and paste it into another layer. Place this copy in the background and re-select the original layer as the current foreground layer. Use the **Volume** select mode to choose only the polygons that make up the lower drop segment (Figure 5).

Your next step is to select the **Taper2** tool. Using

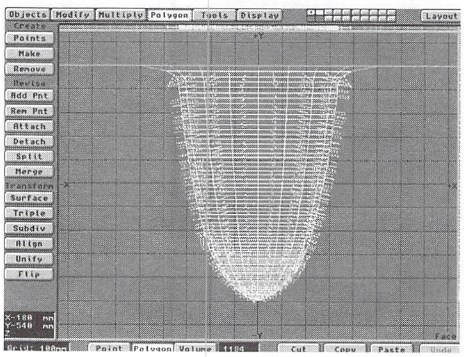

Figure 5

Numeric, select **+ Sense**, **Ease In**, **Y Axis**, and then finally, **Keep**. While holding the left mouse button, move your mouse (centered on the object in the Top view) and taper both sides of the drop of water. Save this object as "Drop 3."

Go to a new layer and place the drip layer (the layer you were just working in) in the background. In the front view, add points around the right-hand edge of the drip shape in the front view in the form of the side of a water drip. Make sure that your first and last points fall on the Y axis line of the original drip shape. Turn these points into a spline by pressing Ctrl-p. Next, lathe the shape around the Y. If you want, count the number of edges and match this second drip object to the first drip, but it isn't important. Save this object as "Water Drip."

Now go back to the second layer (the one we copied the first time) and choose Stretch. Place your pointer on the same level as the flat plane and stretch your object until the drip shape becomes almost a flat section of the plane. Save this object as "Drop 1."

The Animation

Enter Layout. Load all of your objects in the following order: Drop1, Drop2, Drop3, and finally, Water Drip. Here is where the fun comes in. If at any time you feel light(wave)-headed, slow down and reread this section.

Under the Objects panel, select your **Current Object** as "Drop1." Under **Metamorph Target**, select "Drop2." Next, change your Current Object to "Drop2." Place "Drop3" as the Metamorph Target of "Drop2." Also, set the **Object Dissolve** of "Drop2" to 100 percent. Next, change your Current Object to "Drop3." Here, the only step you must perform is to set "Drop3's" Object Dissolve setting to 100 percent.

Returning to "Drop1," the following settings need to be entered into the **Metamorph** Envelope (**E** button). Create keyframes as listed with the given values and bias settings:

Keyframe	Value	Bias
0	0%	0.0
30	100%	-1.00
36	100%	1.00
37	21%	-0.48
44	22%	0.0
47	23%	0.0
51	12%	0.0
59	9%	0.0
65	44%	0.0

Now move to "Drop2," and enter the following settings into the **Metamorph** Envelope:

Keyframe	Value	Bias
0	0%	0.0
20	0%	1.00
35	100%	0.0
36	0%	0.0

Since "Drop3" has no envelope, leave it blank.

Move to the "Water Drip" object and enter the following settings into the **Object Dissolve** envelope:

Keyframe	Value	Bias
0	100%	0.0
36	100%	1.00
37	0%	0.0

The only other step is to keyframe the "Water Drip" object at a slightly lowered position than the "drip" at frame 36. The objects should overlap. Then, place the "Water Drip" object much lower out of frame falling down the Y axis and keyframe the object somewhere around frame 60. Now, after saving the scene make a preview. You should see the drop of water form and then break off and fall.

You advanced animators might have noticed that the morph target envelopes overlap at times. This allows a smoother morph from object to object across a range of objects, rather than an A to B movement followed by a B to C movement.

This scene and objects are included on this month's *LWPRO* disk. In a future article, I'll cover morph targets, cross-morphing and envelopes in more detail.

LWP

Glen D. Miller is one of the newest animators to join the Amblin Imaging team. Striving to become a digital creature expert, he animates while listening to Jimmy Buffett, Tori Amos and Patty Loveless. He can be reached on CompuServe at 73223,3535.

EDITOR'S MESSAGE

continued from page 3

Along those lines, it's important for you to remember that this newsletter is for you. We used to say that this was a newsletter for serious (whatever that means) LightWave animators. But the cover is now more accurate. This is a newsletter for LightWave animators. As both the LightWave market and our readership continue to grow, we are realizing that not everyone out there is a master LightWave user. Actually, less-experienced users are the majority.

We have tried to meet the demand for more beginner-type elements by incorporating articles and columns such as "LightWave 101." However, we will always have in-depth articles like this month's "Replacement Animation."

It is a difficult mix to blend each month, and that's where you come in. We can't know how well we are meeting your expectations without you telling us. We want your suggestions, comments, tips and complaints. You can send mail or faxes to our offices listed in the masthead, or e-mail to myself (electronic addresses also listed in masthead) or to Jim Plant on CompuServe at 72242,1623.

I'd also like to talk about the *LWPRO* disks. I hope these have been valuable (again, you have to let us know). Sometimes, there doesn't seem to be enough material to make a disk worthwhile for a particular issue, so I've decided to start including this "low disk" information on the next disk to go out. Therefore, it may be possible in the future to receive a disk that has items from a couple of different issues. You'll notice this happening with this issue. For any articles that mention disk information, you will see it on the next disk.

Before closing, I want to offer a large "Thank You" to all of the people who help create *LWPRO* every month. Many of you are mentioned in the masthead, but there are some of those who aren't, and I just would like to send out a well-deserved thanks to all of you. I certainly couldn't do this alone (remember—I can't even get articles in on time).

In another year, I'll fill you in on everything that happens between now and then... (grin).

LWP

Digital Cinematography

by John F.K. Parenteau

As computer-generated graphics are used at an ever-increasing rate to augment live action, it becomes extremely important for us "CGI people" to understand how our craft relates directly to on-set cinematography. As we have reviewed in past issues, creating a real-world effect in CGI is usually a much more involved task than many realize. Programmers don't consciously try to make our work difficult; it isn't easy creating some effects in the computer without massive computing power. Even packages with longer histories of involvement in the entertainment industry don't always have a direct relationship from their workspace to the tools used on set.

Take, for example, a soft light effect. The theory to build a real soft box is to bounce light around a white box, then out one end covered with a heavy, yet translucent material. On set, most gaffers can produce a soft box in their sleep, knowing with confidence that this simple design has been used effectively for years. The laws of physics never enter a conscious level of thought.

To program a soft box light type in a program such as LightWave would require massive rendering power and would most likely be too time-intensive. Just the thought of the number of light rays splitting and bouncing around the inside of that box must give most programmers the shakes! Though it would be nice to have a button for everything a cinematographer uses on the set, it isn't really necessary, since most effects are possible with our current tools. As many of my fellow animators are fond of saying, "It'll just take a little messing around to make it work."

In future issues, we will examine how to make the tools at hand work in almost any application. But first, let's take a walk through LightWave and touch on some areas that have a direct relation to live action photography.

Scene Overview

In the film industry, a scene is usually described as a sequence of related shots. In LightWave, we use a scene to describe a single shot. The Scene Overview (Scene panel) feature allows you to examine the layout of your shot at any time. On set, a script supervisor, also known as "continuity," keeps track of each take and the actions that take place within it. If an actor raises a drinking glass with his right hand, the script supervisor is there to ensure that in each subsequent take and any "coverage" (other angles of the same action, such as close-ups), the actor uses the same hand.

Using LightWave's Scene Overview allows you to study each element, its keyframing, and any hierarchies that it is involved in. Suppose you are animating a complex hierarchical design, such as a robot, and the pieces loaded in the scene are not in order. Scene Overview will display the hierarchy in order for you.

Simply start with the shoulder, then highlight the upper arm in the Scene Overview. This selects that object in layout so that you don't have to search the object list for your arm. Compared to the rest of LightWave, Scene Overview is hardly a complex feature, but it is a handy one to have.

Light Types

As I mentioned at the start of this column, most 3D software packages lack the exact duplicate of equipment found on a live action set. This, however, is not a problem to us creative types, since we can create anything with the tools at hand, right? The functions of the three light types LightWave has to offer can simulate almost any form of practical lamp if used correctly.

A **Distant** light should be used in limited instances. Since this light is an overall light from a single source, it is best used to simulate sunlight for day exterior scenes, or space, where the keylight source is usually the sun. Remember, you cannot put falloff on a distant source, nor can an object move into or out of this light unless the light is blocked by another object (and raytraced).

As you may recall, I have an aversion to using ambient light. Though it does serve its purpose in limited situations, I feel it is better to imply ambient light from the source creating it.

For example, the sun emits light rays that strike our atmosphere, and eventually the ground. Though many rays reach the ground, many others are reflected and refracted by particles and debris in the air. This is what produces some of the ambient light visible to your eye. A good practice to help simulate this in 3D when making your sunlight is to apply raytracing settings to create realistic (yet often harsh) shadows. Once you have set this light, clone it, rename the new light as "Ambient" and remove the shadow options for this light. By setting the intensity to one-eighth or so of the sunlight intensity, a sense of ambient light is created from the actual source direction. In addition, radiosity for a day exterior can be simulated with a distant light pointing in the opposite direction of the main light. Keep the intensity low to avoid overfilling your object.

Spot Lights can function as most other forms of light. For example, one of these lights with **Spot Soft Edge Angle** set to zero simulates an open face or lensless lamp on set. Following is a list of practical motion picture lamps and a suggested setting for the corresponding LightWave light. These settings are intentionally vague. It is important to experiment with each to achieve the desired look. Avoid using these as plug and play values. Remember, each situation is different, and it is important to find your own look:

Lamp Type:	Open Face (no lens)
LW Light Type:	Spot Light
Color Temp:	White or slightly warm
Features:	Raytraced, no soft edge angle, low falloff, low cone angle

Lamp Type:	Fresnel lamp
LW Light Type:	Spot Light
Color Temp:	White or slightly warm
Features:	Raytraced, some soft edge angle medium falloff, med. cone angle

Lamp Type:	HMI (daylight lamp)
LW Light Type:	Spot Light
Color Temp:	Slightly blue
Features:	Raytraced, some soft edge angle, low falloff, medium cone angle

Lamp Type:	Soft Light
LW Light Type:	Spot Light
Color Temp:	Slightly warm
Features:	Raytraced, high soft edge angle, high falloff, wide cone angle

Lamp Type:	Arc lamp
LW Light Type:	Spot Light
Color Temp:	White
Features:	Raytraced, no soft edge angle, low falloff, wide cone angle, high intensity

What do all those lamp types mean? We've talked about soft lights before, but the others are as follows:

Open face—any motion picture lamp with no lens to focus the light being emitted, creating a harsher, non-softened light.

Fresnel lamp—any motion picture lamp using a fresnel lens to focus light emitted. A fresnel lens is a piece of glass with ridges that help direct and focus light (Figure 1). As the bulb behind the lens is moved forward and back, the lightbeam becomes tighter or wider.

HMI—daylight (5200 degrees Kelvin, nominal) balanced light.

Arc Lamp—daylight balanced light created by an electrical arc between two carbon rods.

As discussed above, many lights have sharp shadows while others have more diffuse shadows. Choosing between raytraced shadows and shadow mapping can help simulate both types. Though shadow mapping will create a fairly realistic soft-edged shadow, it requires a higher setting the closer you are to the shadow, and thus more memory.

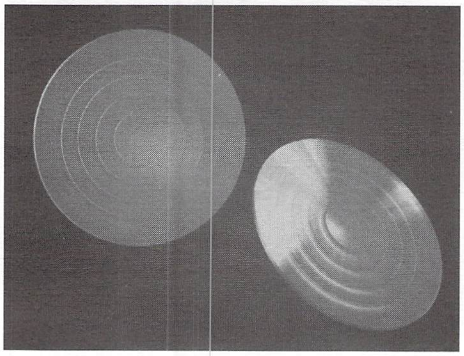

Figure 1: Fresnel lenses.

Camera Panel

The camera panel contains the greatest collection of settings to help simulate live action photography. For those of you who are content producing full CGI shots, the following discussion may not apply (though it may be interesting).

From my standpoint, the truest form of an effect is one that is invisible to the viewer. Granted, many effects are readily noticeable, largely since many of them are monsters, spaceships or aliens. These stand out not because they are unbelievable, but rather because our subconscious mind, and often our conscious mind, knows that these things cannot exist. The best an effect of this type can hope for is to provide a momentary suspension of disbelief. *Jurassic Park*, for example, had CGI dinosaurs. Though these were perhaps the most realistic creatures ever created in the computer, they still stood out because we know that dinosaurs don't exist. This doesn't make them any less thrilling to watch, but the effect is visible. In a movie like *Forrest Gump*, however, many of the effects went unnoticed. I actually had several effects-savvy friends of mine ask if Gary Sinese was truly handicapped (and you tell me you didn't think Tom Hanks was actually playing ping pong). This is the purest form of special visual effect.

To create these effects, it is important to understand how many of the features of the camera panel apply to live action photography.

Basic Resolution settings are fairly straightforward. The lowest resolution usable in the professional world is **Medium Resolution**. With **Low** Antialiasing, the image generated is the equivalent of D1 video resolution, the highest-quality video available today. (It is important to remember when outputting to D1 to set the **Pixel Aspect Ratio** to D1 for proper image sizing.)

Film resolution varies depending on the output medium (e.g., 35mm, 65mm, VistaVision, etc.). Many facilities generate 1K files and size up to 2K for film recording. Though 1K is a fairly non-specific figure considering the number of image sizes this might suggest, the industry standard resolution when referring to this number is 1024x768. 2K usually implies 2048x1536. While scaling from 1K to 2K for scanning to film can be an acceptable solution to high rendering times, it is recommended that you render the highest resolution possible approaching your ultimate target resolution. 2K is suggested to maintain sufficient pixel density on the negative. Remember, use square pixels when rendering frames for film output.

When matching live action footage, it is important to match the CGI lens with the actual camera lens. Though LightWave uses a zoom factor setting, an Equivalent Lens display is provided. Below are a list of common 35mm motion picture film lens sizes and their equivalent zoom factor settings:

35mm film Lens	LightWave Zoom Factor
9mm	1.2
15mm	2
20mm	2.6
25mm	3.3
35mm	4.6
40mm	5.3
50mm	6.6
85mm	10.6
100mm	13.3
150mm	20

Accurate framing of the CGI camera is crucial for matching any live action. Though a Safe Areas function is provided under the Options panel, it is sometimes unclear to some as to the accuracy of these markings.

For video resolution, the entire layout screen is rendered as the visible image. Remember, however, that not all televisions can display the entire viewable area, called "TV Raster." Though it is fun to hope that all your animations will be viewed on the best televisions, you must take into account the lowest common denominator, including the 9-inch black and white television. Always frame your work slightly inside the frame of the layout screen to avoid cutting off any important information.

For motion picture work, there is no safe area function for the various aspect ratios. Turning on the "safe areas" display shows an example of television, or action, safe (the outer lines) and title safe (the inner lines). Television safe is the area of picture that is visible on most television sets. Actually, larger monitors will display an area greater than suggested by the larger box. The smaller box, title safe, is the suggested widest framing for any title information, such as a logo or credit. Note that you can only see the safe areas when viewing your scene through the **Camera** view.

By activating the **Letterbox** function (Camera panel), a set of dotted lines will appear on the layout screen to approximate 1:85 aspect for standard 35mm theatrical film (the LightWave letterbox function is actually a 2-to-1 ratio). This wide-screen aspect ratio is the standard projected size for 35mm film. The number actually means a unit of 1 in height by 1.85 units in width. By using a special lens on the production camera, the photographed imaged can be squeezed to a predetermined size, which, when projected through a similar lens that unsqueezes the image, creates the much larger 2:35 aspect ratio, describing a frame 1 unit in height by 2.35 units wide. This is what is commonly referred to as anamorphic or Cinemascope. Films that are shot on 35mm film but intended for release in 70mm are usually shot in anamorphic. The Letterbox function simulates the 1:85 aspect ratio, but remember to shut this function off before rendering your shot since this setting will draw black bars on the top and bottom of the frame.

When scanning to film, the entire frame of film is exposed. The wide-screen look is produced by sliding a matte behind the lens of the theatre projector, not by limiting the area exposed on the film frame. If this matte was removed, you would see additional information (probably the top of the set, a boom microphone, a crew member hanging out, etc.). It is important to provide the film scanning service with a full frame of information.

Motion picture film projected at 24 frames per second (fps) displays a single frame for 1/24 of a second, relying on persistence of vision to carry viewers to the next frame. Video, however, works in a slightly different way. Played at 30 fps, each video frame is displayed not once, but twice. A frame is comprised of two "alternate line" fields of information. For smoother video animations, you should use **Field Rendering**, with which you effectively get 60 images per second as opposed to 30. This will definitely give you a smoother, yet still video-looking, effect. If you are trying to avoid a video look, stay away from field rendering.

Depth of Field has been discussed many times in this newsletter, and I will save my analysis for a future column. It is important to note, however, that LightWave

see Digital Cinematography, page 9

LightWave & PhotoCD

by Dan Ablan

Just about every serious LightWave animator I know uses image maps in his 3D world. Many, like myself, get images from real-world textures, usually via camcorder, videotape or prepackaged disk sets. Some also use textures from CD-ROMs. But another area that many animators overlook is PhotoCD technology.

Kodak introduced the PhotoCD a number of years ago, and those who have attended any of the yearly Consumer Electronic Shows have seen the demonstrations. It was a great idea: take any film, from any camera, and transfer it to a compact disc, and have your images preserved forever. You can even take existing photographs and have them put on a disc. No photo albums to buy, no sifting through shoeboxes of photographs—just one simple compact disc. Sounds great, right? Well, it seems that the price for a PhotoCD unit was just too much for the average consumer (anywhere from $600-$800), so there wasn't the hype that Kodak had hoped for.

The Technology

It's been said that one person's misfortune (in this case, one corporation's misfortune) is another's fortune. The PhotoCD never made it big in the home consumer market. Sometime last year, the prices on consumer PhotoCD players dropped dramatically, and I was able to take home a brand-new, top-of-the-line Kodak PhotoCD player for $179. Big deal, you say? To me it was. Before I moved into animation and before broadcast, I was studying to be a photojournalist. Even though I had found my niche in 3D, including photography into my animations was always a great interest. Plus, I knew the advantages of using image maps in LightWave. By using a PhotoCD from the Kodak PhotoCD player, I was able to use the video out directly into input number one on my Toaster, without a TBC, and I didn't have to deal with shuttling through video tape, frame-grabbing, then stabilizing my images. The PhotoCD unit lets me crop, enlarge or scale my images before I grab, which is just one less step to go through in the computer. Any roll of film that you take pictures on can be developed onto a PhotoCD. Your local drugstore or photo shop can develop the CD for you. It's not very expensive, and although prices vary, you can expect to pay around $15 for a roll of 24 exposures. Don't feel that you need to go to the best camera shop in town to have your film developed: every store sends your film out for processing to PhotoCDs, so shop around. WalMart usually does a pretty nice job, and turnaround time is about seven to 10 days. You can also take existing negatives or prints and have them transferred to a PhotoCD. Each disc holds up to 100 images.

What About My CD-ROM?

Most of today's CD-ROMs, whether internal or external, will read a PhotoCD. With the right CD-ROM software, such as ASIM CDFS (a PhotoCD manager that makes your CD-ROM readable) you can grab images. This is a great way to go, especially if you are using the standalone LightWave version. The standalone LightWave, obviously, has no Toaster, which means no frame-grabbing capabilities. Using ASIM CDFS, you can have a "contact sheet" of the images on your PhotoCD. From there, you select the particular image you want to load and tell the computer where to save it. This process works just fine; however, I've found that if you have frame-grabbing capabilities, you're better off.

Through the CD-ROM, loading an image takes a few minutes, and saving it adds more time, because the image is uncompressed, sometimes consuming nearly 5MBs of disk space. Then, you need to crop, scale and resave the image before using it in LightWave. ASIM does let you flip or rotate the image, but cropping needs to be done in a paint program, or more likely, through Art Department Professional or ImageFX.

When you use the Kodak PhotoCD player and frame grab the image, the whole process takes less than one minute, and as you may know, the more time you can save, the better. Also, the PhotoCD unit lets you crop and zoom into any area of the image you desire. A remote control is provided, with directional keys that allow the user to position a rectangular box anywhere on the image. Once you've selected the area you want to enlarge, just hit the button, and you'll see the blown-up image right on your monitor. If you want to scale an image down for use as a diffusion map or an image in the distance, you can hit the full-view button on the PhotoCD player and get an image that is scaled down 50 percent but maintains the proper aspect ratio.

Image Quality

Either way you choose, frame-grabbed or CD-ROM, there are bigger advantages to using PhotoCDs in your work. Making scenes appear realistic in LightWave is always a challenge, and the proper lighting and surfacing are keys to creating that perfect look. I use image maps in just about every animation, for simple reflections, image maps, bump maps, etc.

By taking just an hour or two, using a 35mm camera and a good roll of film, you can have a supply of texture maps better than anything you could buy, and these images will make your work more original and unique. For example, in the warehouse image in the color pages, I used a PhotoCD image of my windowsill as a diffusion map for the walls, wood beams and floor. By taking a single PhotoCD image and manipulating it a bit—cutting it up, darkening certain areas through ToasterPaint or some other paint package—I now have different-looking PhotoCD images. I used a PhotoCD image of some rocks for the bumps in the tank. That particular rock texture has been used countless times for roughing up a surface, bump maps, specularity maps and helping add realism to my scenes. Figure 1 shows the actual image of the rocks that I grabbed from my PhotoCD.

Probably the single best advantage of using PhotoCD technology in your animations is the quality. The image resolution of film looks so good that if you

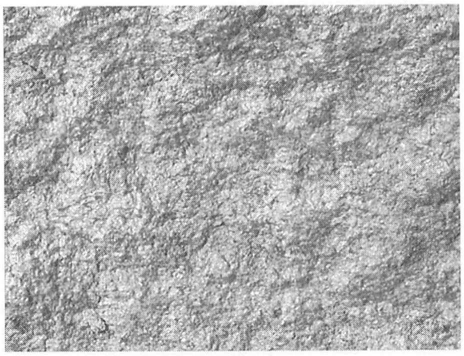

Figure 1: This image of rocks started as a photograph and was grabbed into the Toaster through a Kodak PhotoCD unit.

begin using CD images on a regular basis, you'll have a difficult time when taking your images from video or pre-packaged sets. Also, with the image quality being so crisp, you can scale down the image, saving memory in LightWave, without much degradation to the final output.

For example, if you were building a front hallway and wanted the camera to fly across the room, ending up on a closeup of a picture on a bookshelf or wall, and you grabbed that particular image from video, you would have to either (a) use a large image taking up too much memory, or (b) not travel in so close. By using PhotoCD image quality, you can have an image that takes up only 200K and still be able to travel up close to it.

Figure 2 shows a front hallway. Total image memory used is 1.13MB. The floor is a variation of the rock image (Figure 1), and each photograph and texture is from PhotoCD, taken with 200 speed film and developed at a small camera shop around the corner. The closeup image in the color pages shows pictures on the wall that still look clear and sharp when the camera travels by, though each one is only about 200K.

One thing to think about with PhotoCDs is putting your own pictures into your scenes. Carry your camera around with you and take pictures of whatever catches your eye. Take a picture of your blue jeans and use it for a nice background for text. Or how about the bricks or concrete on that building you always walk past? Wouldn't that be a nice texture to use in LightWave?

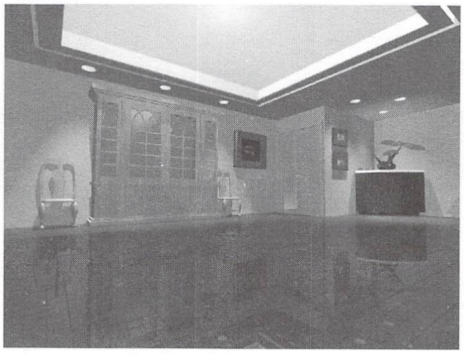

Figure 2: This image of a front hallway uses a total image memory of only 1.13MB. The floor is a variation of the rock image and each photograph and texture is from PhotoCD.

An interesting thing will happen when you begin using image technology like this. You'll develop your eye even more, paying closer attention to textures, surfaces and light in the real world. And you know that creating a 3D computer-generated world that looks like the real world is the ultimate goal.

Results

Experiment with textures as often as you can, for image maps, bump maps, clip maps, etc. Perhaps you could mimic a $50,000 ADO (digital video effects unit) and use PhotoCD images mapped onto flat polygons slowly moving across the screen. If you've seen the TV series *Viper*, you've observed LightWave's powerful Front Projection Image Mapping (FPIM) capabilities in scenes such as those of the vehicle launching a probe out of its back end. Think about using a PhotoCD for your next FPIM project. Take a photo of the street you live on. Get the PhotoCD image of it and bring it into LightWave. Now use front projection image mapping to make a dinosaur or spaceship fly out from behind a car or building.

The more you play with all of LightWave's buttons, the more familiar you'll be with what they do, and you'll open the doors to creativity, thinking of new and exciting ways to achieve that look you want. LightWave is a very powerful 3D program and is constantly improving. By adding elements from the real world through a process such as PhotoCD technology, your animations can become even better.

LWP

Dan Ablan is the president of AGA graphics in Chicago.

Digital Cinematography

continued from page 7

depth of field does not function in the exact way real-world depth of field does. Though LightWave is based on values found when using a real camera and lens, there is no correlation between the amount of light in your scene and the depth of field.

As described in most technical manuals, Depth of Field is a direct correlation between the amount of light reaching the film plane, the size of the aperture opening and the size of the lens. To break it down to basics: the smaller your aperture (and the larger your F-Stop), the greater the depth of field (more in focus); the wider your opening (smaller F-Stop), the less depth of field (less in focus). In addition, the longer the lens (greater the zoom factor), the less depth of field; the wider the lens, the greater the depth of field.

Many fashion photographers strive for as little depth of field as possible. To achieve this, they use extremely long lenses (e.g., 1000mm) and use as little light as possible, forcing them to open the aperture extremely wide (e.g., F1.3). The combination, as you might predict, creates a very small zone of focus, sometimes as little as a few inches, thus creating a softening effect that works well with more artistic photography. Make sure to experiment with your settings before final rendering to avoid any unnecessary surprises. Keep in mind that a wide aperture means a smaller F-Stop number and a decreased focus zone. A small aperture means a larger number and a larger area of focus. Also note that a background image is not affected by depth of field.

Many of the functions discussed above relate directly to compositing rather than pure CGI images. Though full computer-generated images will always play a huge role in the effects world, the exciting new frontier is in compositing with live action. You may not believe it, but you own an extremely professional and powerful compositing tool already: LightWave 3D.

At Amblin Imaging, we have begun immersing ourselves in some very involved composite shots over the recent months. Though we have experimented with many other software packages, we usually return to LightWave for most of our work. Like Amblin, it is in your best interest to stay as close to the cutting edge as possible with your abilities. Pick up books on motion picture production and buy the occasional *American Cinematographer* magazine. Though your interests may not be in photography, expanding your knowledge of the techniques will only enhance your ability to create.

LWP

After graduating from the University of Southern California's School of Cinema/Television, John F.K. Parenteau began his career as a cinematographer for motion pictures and television. Discovering computer graphics in 1991, he quickly found the versatility of visual design lacking on live action sets. Parenteau is currently vice president and general manager of Amblin Imaging, and one of the digital effects supervisors for seaQuest DSV.

Statement of Ownership

Post Office Notice: Statement of Ownership, Management and Circulation for second class privileges as required by 39 USC3685: 1A. Title of Publication: **LIGHTWAVEPRO** 1B. Publication No.**10767819** 2. Date of Filing: **September 15, 1994** 3. Frequency of issue: **Monthly** 3A. No. of issues published annually: **12** 3B. Annual subscription price: **$72** 4. Complete mailing address of known office of publication: 273 North Mathilda Ave., Sunnyvale, CA 94086 5. Complete mailing address of the headquarters of general business offices of the publisher, editor and managing editor: **Publisher: Michael D. Kornet, 273 North Mathilda Ave., Sunnyvale, CA 94086 - Editor: John Gross, 273 North Mathilda Ave., Sunnyvale, CA 94086 - Managing Editor: James Plant, 273 North Mathilda Ave., Sunnyvale, CA 94086** 7. Owner: **James Plant, Laura Plant, sole proprietorship** 8. **N/A** 9. **N/A** 10. Extent and nature of circulation given in this order, number of average copies each issue during preceding 12 months followed by actual number of issues published nearest filing date, A. Total no. of copies **2,561 2,600** B. Paid and/or requested circulation 1. Sales through dealers and carriers, street vendors and counter sales: **N/A N/A** 2. Mail subscription **1,644 1,970** C. Total paid and/or requested circulation: **1,644 1,970** D. Free distribution by mail, carrier or other means, samples, complimentary and other free copies: **N/A N/A** E. Total distribution: **1,644 1,970** F. Copies not distributed 1. Office use, left over, unaccounted, spoiled after printing: **917 630** 2. Returns from News Agents: **N/A N/A** Total: **2,561 2,600** 11. I certify that the statements made by me above are correct and complete. Michael D. Kornet, Publisher.

Animation Techniques

The Walk

by Mark Glaser

The human form is appreciably a sophisticated object, and putting it into motion is no easy task, unless as an animator you are content with any old shuffle. The genuine animated walk is a result of identifying and blending a myriad of individual movements. Since the act of walking is so familiar to us, failed attempts to capture its complexity can be rather discouraging. Take some time to go through this tutorial and you will discover the basics by giving life to a simple collection of polygons.

Do We Have any Volunteers?

Enter Modeler so that we can build a simple figure. Referring to Figure 1, proceed to build the six different parts, using a separate layer for each. Start with the head, using a default size ball with eight sides and four segments. Move it up so that the bottom point of the ball is on the origin (0,0,0).

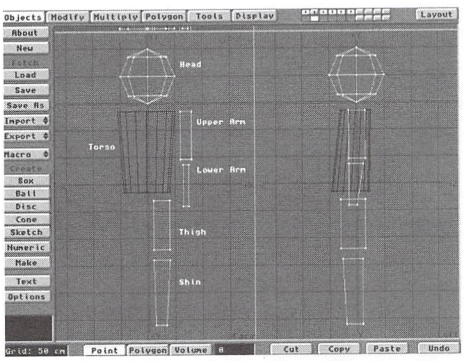

Figure 1: A simple figure means fast screen updates.

Continue adding parts in their natural positions in relation to the head. The torso is a slightly modified cylinder; upper arm, lower arm, thigh and shin are just simple boxes with some tapering.

Export each part into Layout using the same names that I have used. Next, enter Layout and make a clone of the upper arm. Enter the **XY** view and slide the original over so that it is in position to form the right arm. Do this for the rest of the missing parts so that original objects, referred to with "(1)," are on the figure's right and clones, referred to with "(2)," are on the left. If you wish to rename these left and right, simply save the objects from layout using the appropriate names.

Next, add two null objects, calling one "Global" and the other "Local" (note: you can use the "+" key on the numeric keypad to add a null if you are using LightWave 3.5). Parent the objects as shown in the following table. The Rotate column gives suggested rotation restrictions:

Object	Parent	Rotate
Global	(none)	H P B
Local	Global	H P B
Torso	Local	P B
Head	Torso	H P B
Upper Arm (1)	Torso	H P B
Lower Arm (1)	Upper Arm (1)	P B
Upper Arm (2)	Torso	H P B
Lower Arm (2)	Upper Arm (2)	P B
Thigh (1)	Local	P B
Shin (1)	Thigh (1)	P
Thigh (2)	Local	P B
Shin (2)	Thigh (2)	P

Once everything is parented, Figure 2 shows where to place these null objects. Also, select **Move Pivot Point** for each object and place them as shown.

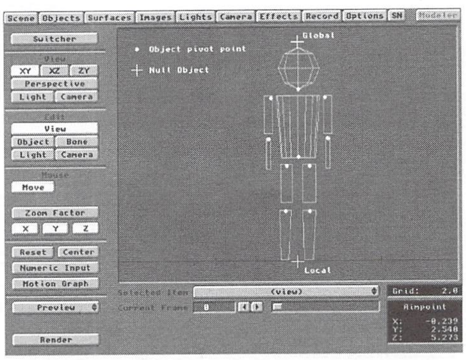

Figure 2: How to set up your model in Layout.

You should now save this group or hierarchy setup as a scene. By saving objects set up as a hierarchy this way you can easily bring them into any scene, ready to animate with the **Load from Scene** button in the Objects panel.

5 Easy Steps

Enter the **ZY** or side view. Zoom into or out of the view so that our figure is fully visible. Take a look at "position 1" in Figure 3. This is where the arms and legs should be in terms of rotation at frame 1. In my figures, the left arm and leg have been enhanced in order to avoid confusion between left and right. The heavier outline designates objects that are on the left or closer to you. If you wish to set your objects up this way, simply go to Modeler and import the left parts. Use the bevel tool with some small settings and temporarily replace the objects by exporting back to layout. Make sure your original objects are saved so you can easily use the **Replace Object** button later.

Set your keys for the eight objects that make up the arms and legs on frame 1 according to position 1 in Figure 3. Proceed in making keys for all five positions as shown in Figure 3. Remember that at this point we are only dealing with the arms and legs. When you are finished, go to frame 25 and create keys for the arms and legs in the same position as on frame 15 except switch left for right. Work backward until you set your last keys on frame 40, which match those of frame 1. For a looping cycle that doesn't stutter, the keys on 1 and 40 should be identical. Simply create keys at 40 using frame 1 as the current frame. Before you continue, realize that creating keys every five frames is not in any way a rule to follow. Although it keeps things simple, remember that the elements of most hierarchal animations don't move to exact timing.

If you haven't created a preview by now, do so over frames 1 to 39. The first thing you should notice when you playback the animation is that the figure seems to be dangling from the head. This is not good. It is here that the null object comes into play. Select the "Local" object and restrain its movement for X

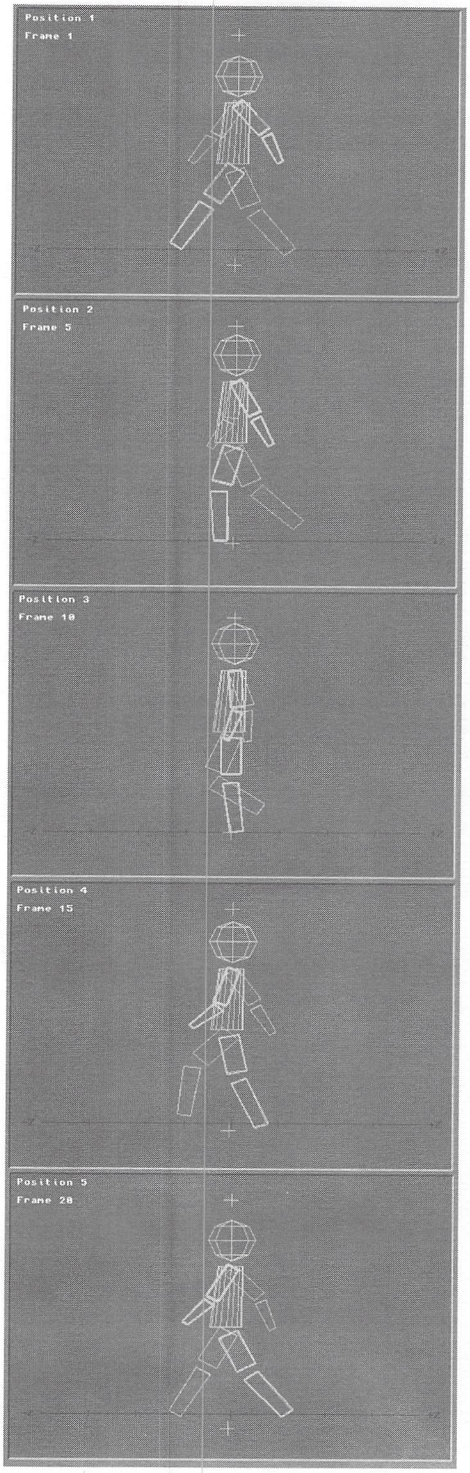

Figure 3: The five easy steps of the basic animated walk.

and Z. Go through all the key frames (frames 1, 5, 10, 15, 20, 25, 30, 35 and 40) and slide the Local null up or down so that at least one foot is firmly planted on the "ground" or grid. Look at Figure 3 again and you will notice the different positions of the local "handle" and where the legs make contact. The next preview you make should be somewhat improved.

Weight and Balance

Switching views to the front, create another preview. You should notice that the figure seems stiff and unnatural. This is because we have not considered weight. Weight is an extremely important item to the animator. Evidence that it's lacking here is our figure, while spending most of the time on one leg, isn't falling over. To give it the weight it needs we can have it fall right over or have it balance. To get an understanding of the two options, stand in front of a mirror with your feet slightly apart. Now, while staying rigid and perfectly vertical as our figure currently is, lift one foot. Hopefully, you put it down again in time. To keep from falling you would naturally lean to the opposite side to shift your center of gravity over the planted leg. Try it in front of the mirror and you will catch on to what your figure needs. Using the "Local" object, bank the figure appropriately throughout the 40-frame cycle using the previously established keys. You don't want to make this leaning too extreme. Experimenting with the degree of tilt can be one way of establishing the identity of a character. Try some perspective previews to see how everything is coming together.

At this point you have completed the most basic elements of the walk. Of course, there is much more that can and should be done. You might not perceive these things when they are in place, but if absent, the animation won't quite "feel" right.

Often referred to as secondary animation, the details you add to dress up the basics are what really give life to a character. Rehearsing any of these movements in front of a mirror can be a great help.

There are some limits to what you can do, given the simplicity of our model, such as the lack of pelvis and shoulders. But consider the arms. As the lower arm swings back and forth, it should cross slightly in front of the body. The amount of swing can be used to say something about the character. Energetic characters use a larger, faster swing. Phlegmatic (or calm) characters would have much less arm movement with a slouched head.

Give the legs some banking or sideways movement. As the leg is raised and brought forward, swing it out a little, bringing it back in for contact with the ground. A character who has a limp would keep most of the weight on one leg, with one leg swinging out more than the other.

From A to B

The simplest way to make your figure actually advance while taking steps is to set a beginning and ending point using the Global handle. The distance between the keys depends on the speed and duration of the leg movement. With this approach you will always have some sliding occurring between the feet and the ground. In some animations it may be adequate, such as when the feet are not in view or are some distance from the camera. When greater accuracy is a concern, movement should be piloted using the Local handle. This should always be the last step in creating a walk because it is based on the way you have established the leg motion.

To use this method, enter the side view and zoom in on the feet. Determine which foot is in contact with the ground and proceed in creating keys that keep this foot stuck to the ground until the next foot comes down. You may want to begin with making keys every five frames and add more if any unwanted sliding remains. Note that the Global null can be used to position your walk cycle anywhere in your animation.

The most important things you should learn from this tutorial are the breakdown positions of the standard walk. You can always refer to Figure 3 to start any type of walk. It is always a good starting point that you can add to. The simple scene that you have set up is also useful for quick and easy motion sketches to test ideas for more ambitious projects, such as the walking knight on the cover of this issue. Another important thing to remember is setting up a hierarchical figure for animating, keeping in mind strategic parenting and the use of null objects. Look for further insight into animation techniques in future issues of *LWPRO.*

LWP

Besides the dismal fact that digital media has de-objectified art, Mark Glaser has no regrets in pursuing his artistic career with computer graphics. He has worked with LightWave as a free-lance illustrator and as an animator for RoboCop: The Series. *Mark recently made the big move to Hollywood.*

In the December Issue:

In next month's issue of *LWPRO*, we will wrap up the year with a technical look at raytracing, an end-of-the-year index to past issues and a few other LightWave goodies.

LightWave 101

Course 1D: Animating with Splines

by Taylor Kurosaki

To animate is to bring to life. However, creating life is often an incredibly painstaking process. Just imagine the work involved in hand painting thousands upon thousands of cells for an animated film. Now imagine a cartoon animator only having to paint every other frame of an animation. How about every fifth frame? Every tenth? Every hundredth? The labor and time saved would be enormous. Now apply this principle to computer animation. Imagine how tedious it would be to keyframe the camera, every object, and every light at every frame of your scene. Thankfully, this is not the case. The computer calculates the positions of the components of the scene based on their respective keyframes immediately preceding and following the current frame. In LightWave, interpolation, as this mechanism is called, comes to us in the form of splines. Certain characteristics attributed to keyframes affect the spline and determine how the object's position is interpolated. In layman's terms, animating with splines in LightWave is not unlike having your own overseas animation department:.

This installment of "LightWave 101" will cover the use of splines in LightWave animations. More specifically, effective use of the **Motion Graph** and **Spline Controls** will be illustrated, along with the **Align to Path** feature.

A few words of general advice: add keyframes as discriminantly and strategically as possible. Don't get me wrong, I'm not suggesting skimping on keyframes in any way. Rather, be wary of ending up with a messy animation full of unnecessary keyframes. This usually occurs when an animator tries to fix a problem in an animation by adding extra keyframes rather than first trying to adjust those which already exist. With very complex animations involving multiple, interacting objects, it may seem easier to "tweak" a motion by adding keyframes to counteract the negative effects of others rather than altering existing keyframes. This only serves to further complicate the animation when you invariably must go back to make later changes. Get enough of this and you may end up with a monitor lying in pieces on the sidewalk below. Don't be seduced by the dark side...fix it right the first time.

The Motion Graph

The Motion Graph window is an incredibly helpful feature that debuted with the 3.0 release of LightWave. It provides LightWave animators with a visual representation of their animations and allows interactive adjustment of an object's, camera's, light's, or bone's movement. The main area of the window contains a graph which plots movement along the three axes, heading, pitch and bank angles, object scale along any axis, and velocity. Values for the various channels are plotted vertically, while time is plotted horizontally. Through this interface, one can **Add, Delete, Drag, Shift** and **Scale** keyframes, and observe visually the results of changes made. While you can do all of this in the main Layout window, it is often easier to use the motion graph controls. Select the motion graph for an object, bone, light or camera by selecting the item in Layout and clicking on the **Motion Graph** button or by pressing the "m" key any time you have an item selected. Remember that dragging a keyframe with the left mouse button allows you to move the key up or down (except while in the Velocity channel) and dragging with the right mouse button allows you to move a keyframe right or left.

Spline Controls

One of the most useful purposes of the Motion Graph is the ability to visualize the effects of **Spline Controls** on a motion path. The Spline Controls available in LightWave are **Tension, Continuity, Bias** and **Linear**. Valid settings for the first three controls are -1 to 1, with 0 as the default. Linear causes an object to move directly from the previous keyframe directly to the selected keyframe without any spline deviation. When Linear is selected, it overrides all other spline controls from the keyframe on which it is selected to the keyframe immediately preceding it. It's important to note that spline controls are still valid to use at a keyframe with Linear selected since they can be useful in the path leading to the next keyframe.

Tension controls the velocity at which an object travels through a keyframe. A negative tension setting will cause an object to accelerate through the

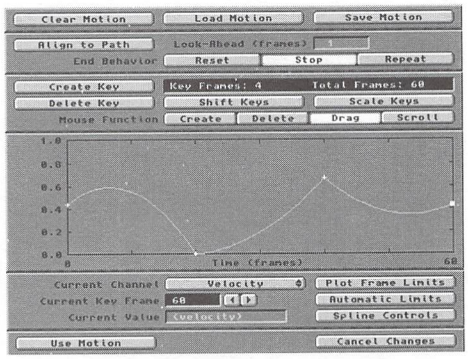

Figure 1

keyframe, while a positive tension setting will cause an object to slow as it passes through the keyframe, then accelerate gradually away from it. A tension setting of 1.0 will cause an item to come to a momentary stop at a keyframe. Figure 1 shows the effects of tension on the Velocity of an object. Frame 20 contains a tension of 1, while frame 40 has a tension of -1.

Continuity affects the smoothness at which the keyframe is incorporated into the spline path. With a continuity setting of 0, the spline curves fluidly through the keyframe. With continuity set to an extreme -1, the spline abruptly changes course directly at the keyframe, causing an object to alter its direction immediately as it reaches the keyframe. A continuity setting of +1 causes the spline to over-compensate before and after the keyframe and "stutter" at the keyframe. Positive continuity values are generally not common, unless continuity is used in combination with another spline control.

Bias allows one to offset the spline curvature before or after the keyframe. With a positive setting, the object will overshoot the keyframe, while a negative setting will result in a spline path which anticipates the keyframe.

All these Spline Controls can be altered directly in the Motion Graph using the keyboard. Hold the t, c, or b key and click and drag right or left to adjust the spline by eye, or enter the values numerically by clicking the **Spline Controls** button in the Motion Graph or from within Layout.

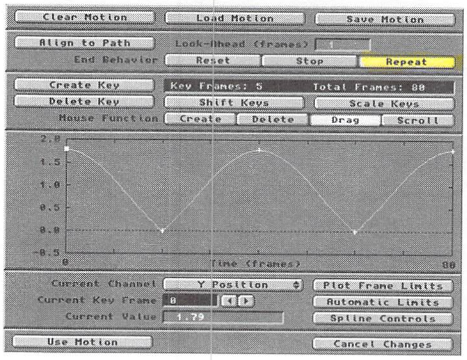

Figure 2

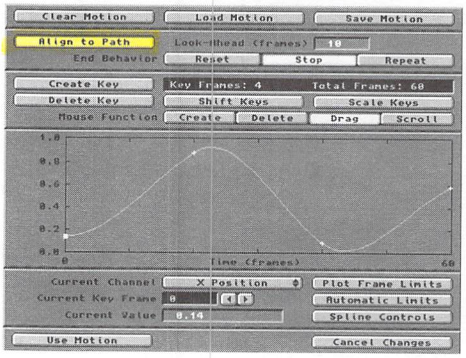

Figure 3

Follow the Bouncing Ball

Using Spline Controls is an excellent way to easily animate a bouncing ball. Figure 2 shows the spline path for the Y Position of a bouncing basketball. To reproduce the way the ball slows as it reaches its apex, surrendering to gravity, use a Tension setting of 1 for the keyframe at the top of the bounce (Frames 0, 40, 80; Figure 2). A tension setting of 1 literally causes the ball to stop briefly (for the duration of one frame, to be exact) before heading earthward. Conversely, to accentuate the force of the ball slamming solidly to the ground, use a continuity setting of -1 (Frames 20, 60; Figure 2). To keep the ball continually bouncing select an End Behavior of **Repeat** from the Motion Graph. In order to make a smooth loop, just make sure the first and last keyframes have identical values. Keep in mind, though, that a ball would not endlessly bounce unless someone were dribbling it, and upon each bounce, the ball would not rebound as high.

Driving School

Bias is a useful way to simulate an object with a lot of velocity and mass. Examples which come to mind are a speedboat or a race car. In these examples, the mass of the object, combined with its forward velocity, causes the object to overshoot the keyframe as it begins to turn. Use a bias setting of 1 to offset the curve in the spline path to after the keyframe. Figure 3 shows the X position of a race car as it drifts around corners, or slides on an icy road. Notice how the spline remains nearly straight until after the keyframes at frames 20 and 40. Whereas positive bias values would be used for a skidding race car, negative bias values would be used to simulate a perfectly in-control car. With a bias of -1, the curve in the spline path occurs before the keyframe as the car sets up and completes its turn before the keyframe (Figure 4). In conjunction with bias, **Align To Path** can be used to keep the car's heading aligned with its direction of travel. In the case of the skidding car, a **Look-Ahead** setting of five to 10 frames accentuates the sliding motion. For the in-control car, use a look-ahead value of 1 frame. Note: When using Align To Path, the object will reset its heading when it reaches the last keyframe, due to the fact that it has no more frames to look-ahead to. If you wish to use Align to Path and you experience this "last frame resetting," you may wish to extend the motion path past the last frame you will see in order to keep your object aligned.

Terminal Velocity

The **Velocity** channel is an especially useful component of the Motion Graph. It is an effective way to

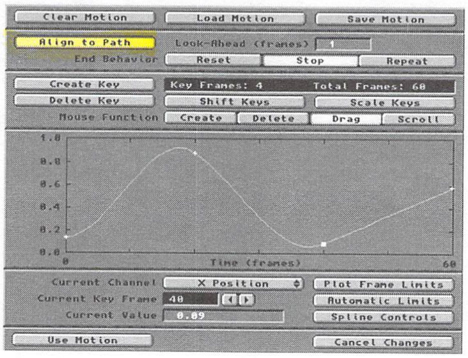

Figure 4

diagnose and smooth jerky animations in a number of ways. Assuming you are trying to produce an animation with a fairly constant velocity, the motion graph in Figure 5 would reveal some significant problems. Dragging the mouse while holding the right mouse button, the spline can be smoothed by spacing keys farther apart to lower the velocity or bringing keys closer together to raise the velocity. Alternately, tension can be used to raise or lower velocity values to smooth the spline. Figure 6 shows the path from Figure 5 smoothed out by adjusting tension and sliding keyframes. Ideally, however, inconsistent velocity values can be averted altogether by establishing a basic velocity by creating the first and last keyframes

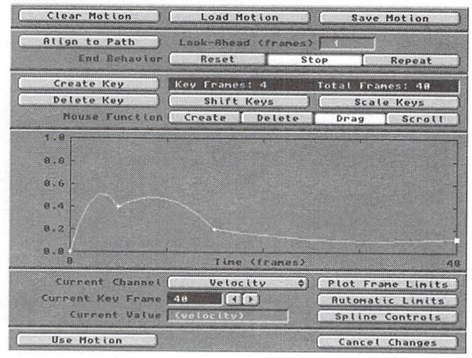

Figure 5

of your animation first. Keyframes in-between are then created subsequently to conform somewhat to the original two-keyframe spline. This is easy to do by going to the frame you wish and moving the item slightly from its original path. The resulting final

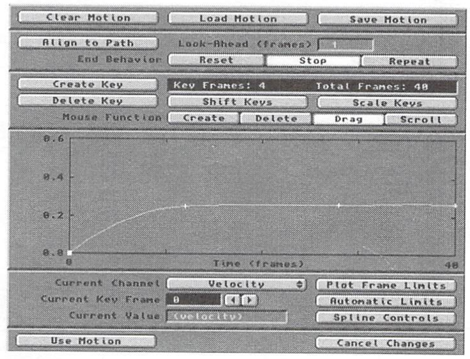

Figure 6

velocity stays in the neighborhood of the previously set values.

While these are very rudimentary examples of spline control use, the principles remain true as you tackle more complicated animations. Character animating, in particular, can be a very frustrating task. The Motion Graph becomes indispensable as it allows you to smooth out jerky motions visually and interactively. In general, however, spline controls enable you to create complex motions with minimal key-framing. Not only does this save time, but it keeps animations streamlined and simple. When you find yourself working on an animation containing several hundred keyframes, you gain a new level of appreciation for this fact. Truth be told, shattered CRT tubes are a pain to clean up.

LWP

Taylor Kurosaki is a visual effects artist at Amblin Imaging. Topic suggestions and comments should be sent to his attention at 100 Universal City Plaza, Bldg. 447, Universal City, CA 91608, or to John Gross on-line.

Corrections:

The Water Surface image in the July issue was credited to John F.K. Parenteau. It was actually done by Greg Teegarden.
In the September Issue, image descriptions were mixed up in "Metaform Basics" for images 6, 7, 8 and 9.

Lens Flare Madness

by Mojo

For decades, if a director wanted cool lighting effects in a film, it took a team of rotoscope artists to accomplish it. The glow of a spaceship engine? Easy. Simply create photographic blowups of every frame in the sequence, have the roto guy rough out in pencil the glow for each frame, then airbrush the finished product on acetate and photograph it against a light table. Of course, the finished film has to be sent to the optical department to be re-composited into the original footage.

These days, we just press the "lens flare" button. The history of LightWave's lens flare effect dates back to 1992, when Ron Thornton felt it would be an important feature in the production of the original *Babylon 5* pilot. Although it certainly helped contribute to the expensive look of the effects, the early lens flares required a lot of tedious envelope setting and manipulation. With the release of version 3.5, the lens flare now has a multitude of features that make it easier to use than ever.

Up and Down

Back in the early years, lens flares were simple, additive effects layered over the entire image after it was rendered. If we wanted a lens flare to go behind an object, it was necessary to carefully examine the scene frame by frame to determine when the flare should be obscured. An envelope then had to be created that ramped the flare up and down over the appropriate frames.

Now you can press the "fade behind objects" button in the lens flare menu and it's completed for you. When this is active, LightWave will figure out what objects are between the camera and the lens flare and ramp them automatically. However, circumstances may arise when this is undesirable: a lens flare should show through the base of a lightbeam, or a transparent or gaseous object. In this case, by clicking off the "cast shadow" option, the lens flare will be instructed not to fade behind the desired object. If the object needs to cast shadows, you may need to resort to the old-fashioned method of hand-ramping a flare.

The biggest lens flare complaint has always been the need to manually adjust the size of flares that move to and from the camera. Thanks to the new "fade-in distance" feature, this tedious enveloping is now a thing of the past.

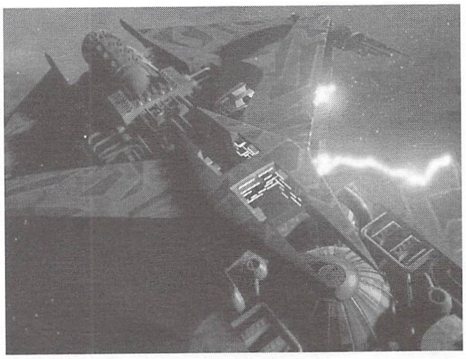

Figure 1: The damaged Narn cruiser from *Babylon 5*. Perfectionists may want to create envelopes to make each flare shimmer a little to add realism.

To use this feature, set up a flare as you would normally. Give it a size you feel would be appropriate based on the effect and the flare's current distance from the camera and then type that distance into the flare's "nominal distance" box. Don't worry if you're having trouble figuring out the distance. Just type in any number and adjust either it or the flare size until the distance is correct. Once the flare looks fine, you won't have to worry about it again: LightWave will increase and decrease the size to maintain your settings whenever the flare (or the object it's parented to) is moved closer to or further from the camera.

The "fade-in fog" button has a similar effect, but obviously only works when fog is turned on. Once the flare has reached the maximum distance of the fog, the flare will be completely ramped down. I find that dissolving the flare out with an envelope can help or even replace this effect when necessary. Again, using the example of the engine exhaust, the flare would appear to shrink inside the engine object if fade-in fog was utilized. This would look wrong, while dissolving the flare out as the object dissolves into the fog would appear more natural. I would recommend fade with distance and a simple dissolve envelope on the flare instead of using fade-in fog at all.

Putting On the Squeeze

"Anamorphic squeeze" and "anamorphic streaks" mimic the look of photographing a bright light source with a wide-screen movie lens. When a wide-screen movie is made, the lens scrunches up all the information and the resulting image on the film is actually stretched vertically. The projector is fitted with a sort of "anti-lens" that stretches the image horizontally to correct it and fill a wide screen. Since lens flares are actually created by the lens itself and not in the scene, they end up being round on the finished anamorphic film. It is the expanding of the image when it is projected that gives this flare its ovalic shape. The blue streaks are simply an odd characteristic created by the lens.

Given this information, if realism is important to your animations, the anamorphic options should only be used in letterbox mode (or when compositing lens flares into existing anamorphically shot footage). Sure, few movie watchers know the whys and wheres about anamorphic lens flares, but these effects are never seen on anything other than big-screen movies (TV shows never have them since they are not shot with anamorphic lenses). In addition, they should only be used to simulate lens flares from bright light sources, such as flashlights or aircraft lights, and not effects (like engine exhausts). Of course, this only applies when realism is a goal—there really are no laws when all you want to do is make stuff look cool.

Tricks of the Trade

Headlights, engine exhausts, laser beams, explosions—these are some of the more obvious lens flare applications used by just about everyone who owns LightWave. However, after using it almost every day for the last several years, some of us old timers have come up with a few neat lens flare tricks that we feel can now be passed down to the young folks.

Ever want to rotate a lens flare? On a few occasions I have needed to make those neat ethereal streaks spin around in a heavenly fashion. The only way to accomplish this is to put a lens flare on a polygon and physically rotate it. You'll need to ren-

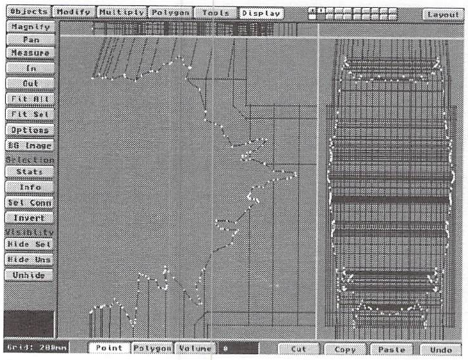

Figure 1A: Although not necessary, using point lights for lens flares helps you keep track of what's what, since point lights look like lens flares.

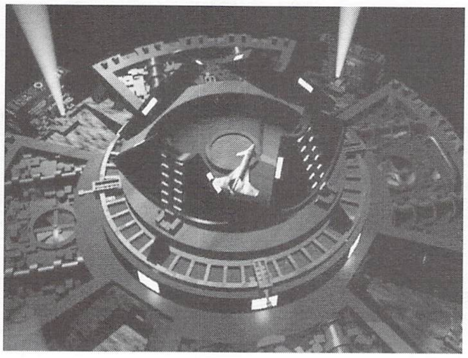

Figure 2: A spectacular image to be sure, yet somehow devoid of a certain something.

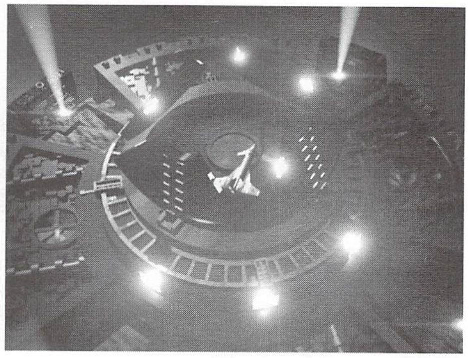

Figure 2A: Ah! Lens flares over the glowing panels do just the trick. What panache!

der a flare and save the RGB file to map onto a polygon. Make sure the flare is relatively small (maybe 60 percent) to keep the edges of the screen completely black, since you'll need to use this image as a transparency map as well. Unless the edges of the map are completely black, the edges of the polygon will be visible (use negative image for the transparency map). If the flare will get close to the camera or fill the screen, you will probably have to render the image as either high-resolution or print-resolution to avoid jagged edges in the flare streaks (for very pronounced streaks, make the light color zero and increase the size of the flare). This technique was used to create the Switcher's popular "flashbulb" effect, in which (you guessed it) a lens flare quickly spins into and out from the camera. This trick may also be desired to rotate the points of a star-filter lens flare.

When lens flares get big, it becomes much easier to see through them. This is a nuisance when trying to cover up a large area with a flare (such as an engine exhaust) and you always end up with an uncontrollably huge glow. How do you fix it? Double up your flares. Two flares in the same space can be controlled to create a strong center with a soft, manageable glow and not be too large.

One flare needs to be designated the hot spot (usually white) and the other serves as the outer glow (providing color). The hot spot should generally be about half the size of the outer glow and dissolved out as much as 50 percent. (The outer flare will add to the central one and bring up the apparent intensity—two 100 percent flares would create a flare with a 200 percent center and appear much too strong.) I usually turn off every flare effect except central glow—the deletion of especially the random streaks option creates a much softer, generally more pleasing effect. Also, remember that the incredibly useful Lightswarm macro allows for the creation of double flares, so use it to save time when creating multiple sets of flares with similar attributes.

An incredibly cool, although time-consuming, lens flare effect is the heat fissure. By lining up dozens of low-intensity flares in close proximity, you can simulate a hot, glowing tear, welding streak, lava flow and many others. The idea is to put so many flares next to each other that they lose their circular shape and blend into one another, appearing to become a straight (or curved) line of light. Figure 1 shows a scene from *Babylon 5* in which a damaged Narn cruiser has just been hit by an energy weapon. I wanted the metal to appear burning hot and knew the best way to do that was to line the damaged area with enough flares to simulate the effect. Since more than 100 flares were needed for the shot, the Lightswarm macro quickly became my best friend.

In Modeler, I isolated the points around the damaged polygon and copied them to another layer (Figure 1A). Using the drag points tool with the grid snap off, I positioned them to be more or less equidistant from one another. I then ran the Lightswarm macro and guessed at a lens flare setting. The key is to make the flares just large enough to cover the edge of the object.

The first rendering showed that the flare size was good, but individual circles were still apparent. This meant more flares. Back in Modeler, I added close to double the number of points I started with and ran another Lightswarm. This time the circles disappeared but the huge number of flares adding to each other created a center line that was too hot, so I had to re-run Lightswarm and give all the flares a 30 percent dissolve.

This procedure worked great. Then came the tough part—manually giving over 100 flares proper fade with distance parameters since Lightswarm doesn't incorporate new features yet. Also important was making sure that only a few flares had light intensities. A nice orange haze over the surface of the ship was important for realism, but only a few evenly spaced lights with falloff were necessary. Absent-mindedly leaving effects flares with 100 percent light intensities is a common mistake, so keep an eye out for it.

The entire effect took perhaps half a day to get right, yet the results are clearly worth it. Always test out your effects with just a few flares to get a ballpark idea of how many you'll need and what their settings will be. This will save you a lot of rendering time and mean fewer visits to Lightswarm.

Believe it or not, lens flares can create subtle effects as well. My favorite is to place a dissolved flare (about 40 percent with outer glow only) over a visible lightsource or window. Figure 2 shows several luminous slits in a futuristic building's hatchway. However, Figure 2A displays the same slits with several lens flares placed over them; the result is much prettier and effective. The look is quite similar to photographing a scene with a fog filter and takes the harshness off a flat, luminous polygon. Use this to add just a touch of class to any scene.

I hope this gives some food for thought to all you flare-crazy animators. Lens flares look great, yet are so easy to use that few people give any thought to when they should be used. Like any other special effect, restraint is the key. After your 10th shot of a flare-encrusted UFO, the novelty begins to wear off. Use flares to add a little zest to your scene—not to hide poor modeling. If your shot doesn't look good without lens flares, you need to brush up on some other skills before you hit that oh-so-tempting button.

Mojo has a thing for engine exhausts.

What's on the Disk?

This month's disk includes samples of Grant Boucher's ADProConvert.Rexx program, scene files from Glen Miller's "Your Friend the Null" and Ken Stranahan's "Metaform Magic" project file from last issue. Scene and object files from this month's issue and modeler patch to fix some bugs in Modeler 3.5 are also included. To order your *LWPRO* six-disk subscription, call toll-free 1-800-322-2843; rates are $30 U.S., $40 Canada/Mexico and $50 overseas.

And the Winner is...

by Grant Boucher

One of the big advantages of being at Amblin Imaging is the fact that we are one of the premiere beta test sites for LightWave. Because of this, we get to try out new features before they are released to the general public, and we get to moan and groan over bugs that you, the reader, have never even heard about. With LightWave 4.0 preparing for release on multiple platforms, we have been beta-testing numerous versions of ScreamerNet (LightWave's render-only engine) on the wide variety of CPUs now becoming available to the LightWave animator.

After having passed these benchmarks on to various hardware vendors to spur on the "spirit of competition," we at *LWPRO* thought our loyal readership might want to see how their current system stacks up, as well as take a peek at what juicy render engines are lurking in the future.

CPU MHz	M68040 28	MIPS R4400 100	MIPS R4600 133	DEC Alpha 275
Render Time				
RenderTestTypical	3460s	1036s	654s	409s
RenderTestNastiest	4944s	1594s	846s	353s
040 Time Comparison				
RenderTestTypical	1.0x	3.3x	5.3x	8.5x
RenderTestNastiest	1.0x	3.1x	5.8x	14.0x

Benchmarks of popular CPUs (data in seconds and ratio to 68040)

Availability

Only the MIPS version of the ScreamerNet render kernel is generally available to the public at the time of this writing. Contact NewTek directly about possible availability of the render kernels mentioned in this article. However, assume that as LightWave and Modeler 4.0 are made available around December 1994 for these CPUs, so should the ScreamerNet modules.

Your Mileage May Vary

While all of these render kernels behaved flawlessly in our tests and we believe these versions to be stable and truly representative of LightWave's behavior across these computer systems, programmers Allen Hastings and Stuart Ferguson are adding new features daily. Compilers, features and optimizations will surely change the behavior of LightWave substantially by the time of formal release. While I believe that these changes should affect all systems relatively equally, variations must be expected.

Note also that the benchmark times for all CPUs except the 68040 include the time it takes for the ScreamerNet render node to return the rendered image to the host Amiga. LightWave 4.0, running native on these CPUs, will certainly not have this added delay.

Updates

As the pace of hardware development escalates (don't you love it?), we will certainly update these benchmarks as needed here in *LWPRO*. This data is valid only for this test, of course, so if you have any special requests, write us or send e-mail.

Scenes Tested

I chose two example scenes of my own design for my benchmarks, based on a number of factors. The first decision I made was to create two scenes. One would be a nightmare for floating-point intensive renderings like those involving ray-traced shadows, reflection and refraction (RenderTestNastiest). This sample scene, included on this month's *LWPRO* disk, takes a plane of polygons and sends them through a ripple displacement map. The plane has a glass surface, and I turn on ray-trace refraction, reflection, and shadows in the **Camera** panel. It may not be pretty (see illustration in color pages), but it proved to be an excellent "worst-case scenario."

My other scene (RenderTestTypical) should have just tested just the integer side of these CPUs, using a typical LightWave rendering. However, I chose a typical *seaQuest* scene (the EVA suit walking at the bottom of the ocean, from the first season episode "Abalon") with ray-traced shadows on. I chose to ray-trace the shadows because I believe that this scene represents a real-world benchmark. Your best and most professional work most likely has shadows in it, so I felt this benchmark truly represented the day-to-day work of a large amount of the *LWPRO* readership.

I chose not to use the Texture Examples scene (a typical benchmark scene) because the newer render engines render such images so fast, it takes longer to transfer the image across ScreamerNet than to render them.

CPUs tested

Motorola 68040—The standard for today's LightWave animator comes in 28MHz, 33MHz, and 40MHz varieties from many Amiga third-party developers. Our tests were with Progressive Peripherals and RCS 68040 accelerators running at 28MHz.

Note: If you are using a stock 68040 Amiga 4000, you should realize that the lack of burst and cached RAM on your machine means that these third-party accelerators render anywhere from 75-100 percent faster than an Amiga 4000. Adjust the 68040 benchmarks below accordingly if you have an A4000. I chose not to follow the lead of most hardware vendors who use the A4000 as a benchmark to inflate their system's rendering ability. I feel that the A4000 is a crippled 68040-based machine and not the LightWave-community professional standard.

MIPS Technology R4400 and R4600—These are the same CPUs used in the popular Indy and Indigo lines of computers from Silicon Graphics, the Raptor series of LightWave render engines from DeskStation and the ShaBLAMM! PC accelerator cards from ShaBLAMM! Computer Inc. The R4600 is the newer, faster, consumer version of the R4400 processor, but is otherwise similar in architecture.

We will be testing the ShaBLAMM! and the SGI Indy in the near future, but for now, assume that similar processor and MHz ratings provide comparable render results. While the ShaBLAMM! and the Raptors use CPU caches to

see And the Winner is..., page 18

Reader Speak

Subscribers' Questions and Comments

by John Gross

This month's Reader Speak takes a look at some of the LightWave questions in the queue...

Q: I've been finding your newsletter to be extremely useful. You've answered a lot of questions for me. I especially appreciated the articles on macros, but really, every issue has been full of terrific information.

One thing I wish I could learn more about is the use of fog. I've never been able to get it to work. When I use it, it blocks out the whole scene, except for some lights that are aimed into the camera (with Intensity set high enough).

A question you might also deal with is how fast different systems are at rendering the same scene. I have a 2000 with an 030 (33MHz) and 12MB of RAM. A scene with a few objects, three or four lights, bump maps, smoothing on and some surface mapping takes about 40 minutes to render per frame. If I add Motion Blur or go to Low antialias instead of none, it might take an hour and 20 minutes. If we did exactly the same thing on a 4000 with an 040 or a "Warp Engine," how much time would I save per frame? Lastly, why is it that when you create an envelope for a light that has Lens Flare on, it will never go off? I've tried many times to place a light inside or behind an object, and if Lens Flare is on, no matter how opaque the surface it is behind, no matter if it's set to absolute zero intensity or even negative values (both Intensity and Ambient Intensity), it still shows when rendered (the central glow at least). Can you help with this?

Pete Wagner
MinneHA! HA! Studios
Minneapolis, MN

A: Hi Pete. It's good to hear from a fellow Minnesotan again.

To start with, let me discuss the system speed tests you mention. In this issue, Grant Boucher runs some tests on some of the ScreamerNet CPUs that are available (or will be soon). We have not run any official tests on the Warp Engine yet, nor on 030 systems. We are including a scene that Grant used for his "nasty" benchmark tests on this issue's disk. Perhaps some readers with these systems will perform the tests and report back. Also expect more coverage of speed tests in upcoming issues as more options are available.

The other two questions I believe I can help with.

Let's talk about fog first. The Fog options (found in the **Effects** panel, Figure 1) have been greatly modified since earlier versions. LightWave 3.5 now includes two types of nonlinear fog and **Minimum** and **Maximum Fog Amounts** (a *seaQuest* lifesaver).

It's important to realize what's happening when you enable fog. First and foremost, turning fog on does not automatically create a foggy haze around your objects. It works a little bit differently than that.

When you enable fog, you are instructing LightWave to "tint" an object the color of the fog, depending upon certain parameters that you input in the fog options. Fog color can consist of either a single solid color selected with the **Fog Color** requester, or **Backdrop Fog**, which will use whatever values the background is, whether it is a **Background Image**, a **Solid Backdrop** color or a gradient color spread (between the **Zenith, Sky, Ground** and Nadir colors).

LightWave will systematically replace the object's color with the fog color depending upon how "deep" the object is in fog. This is where the minimum and maximum distances as well as the minimum and maximum amounts come into play. For all of these distances, it is important to remember that the measurements are always measured in units from the camera.

For our example, let's take three urns located at an increasing distance from the camera. The first urn is approximately .7 meters from the camera. The second is located 1 meter away and the third is 1.3 meters away.

Think of fog as a sphere surrounding the camera. If we had a Maximum Fog Distance of 1.5 meters, any object past this measurement would take on 100 percent of the fog color. Therefore, if you had a light blue background and a fog color of yellow, any objects located farther than 1.5 meters from the camera would be yellow, regardless of their previous colors. However, the real power of fog works when you select **Backdrop Fog** as the color. With Backdrop Fog, the object will seem to disappear the further it is in fog until it is past the **Maximum Fog Distance**, where it will become "invisible." Selecting a maximum distance of 1.5 meters would place the fog just behind the urn farthest from the camera. If you choose **Show Fog Radius** in the **Options** panel, you will see a circle representing the Maximum Fog Distance sphere in the three orthographic views (Figure 2).

Now, whenever you move the camera closer to the urns or the urns closer to the camera, they will appear to "come out of the fog." Likewise, an object moving farther from the camera will recede into fog. Now you get the idea of how underwater scenes are effectively created.

The **Minimum Fog Distance** tells where (in relation to the camera) the fog effect begins. Usually this is left at 0, which means it begins right at the camera. For certain special effects, you may want the fog to begin at a distance from the camera. Any object between the camera and the Minimum Fog Distance setting will take on no fog values.

The **Maximum Fog Distance** instructs LightWave where to create a total fog effect. As stated above, any objects located past the Max setting will take on the total fog colors. Of course, any objects in between the two settings will have some value of the fog color added to their surfaces, depending upon their distance from the camera.

The amount of fog effect at the minimum and maximum distances was always calculated at 0 percent and 100 percent, respectively, until LightWave 3.1 added the ability to determine how much of the fog effect would be calculated at each distance value. The **Minimum Fog Amount** refers to the percentage of the fog effect at the Minimum Fog Distance. Before this ability, it was a common practice for *seaQuest* animators to use a negative Minimum Fog Distance in order to create some "fog" right at the camera. Now, inputting a value for Minimum Amount assures that you can be close to objects yet have the objects partially obscured in fog (you know, like a nice smoggy day in Los Angeles).

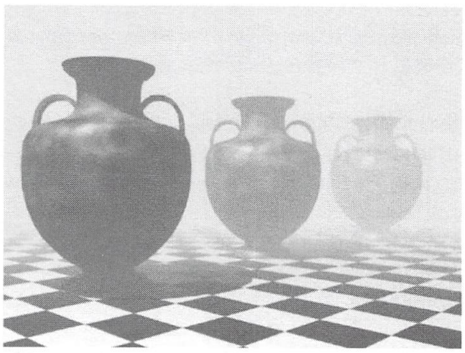

This image demonstrates a Nonlinear 1 fog with the maximum distance placed just behind the farthest urn.

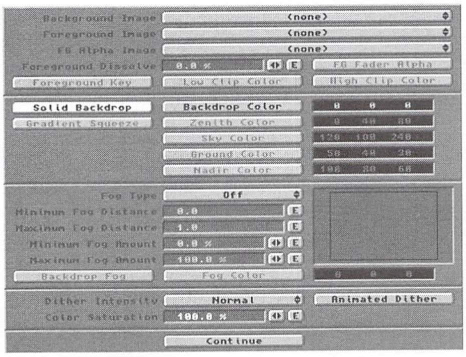

Figure 1

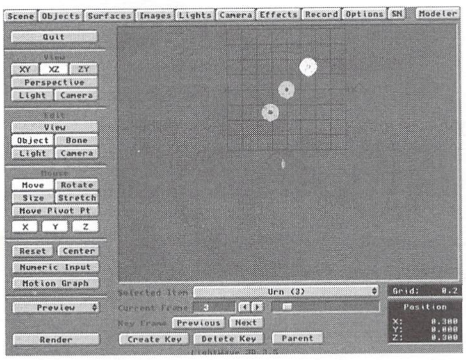

Figure 2

Of course, you can adjust the maximum amount of fog to assure that objects will never be completely obscured by the fog color as well. The fog graph shows the amount of the fog effect dynamically change as you adjust the Minimum and Maximum Fog Amount values (try using the mini-sliders). By creatively adjusting these values, you can achieve a fog that grows thicker the closer you are to the camera and gets less and less thick as you move away from the camera (Figure 3) for interesting effects.

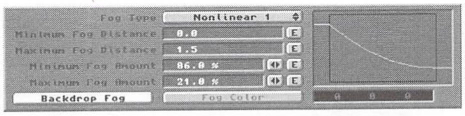

Figure 3

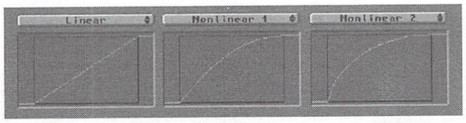

Figure 4

The fog graph also shows you how the different types of fog affect the way the fog is applied. **Linear** fog is evenly ramped up as you get further away. **Nonlinear 1** is a bit more realistic look using an uneven ramp of fog that is "thicker" closer to the camera than Linear. **Nonlinear 2** is even thicker near the camera. Figure 4 shows side-by-side comparisons of the different fog types. Also note that all of the fog values can be enveloped for interesting effects (such as fog rolling into town).

Now that you know how fog works, I bet you can figure out what is happening when you are using it. Chances are you are leaving the default values of fog in place, which will give you a **Maximum Fog Distance** of 1 meter and a Fog Color of black (0,0,0). Even if you are using **Backdrop Fog**, a 1-meter maximum distance is probably a bit small for many scenes. This would cause anything past this distance from the camera to take on the fog colors, whatever they may be (assuming, of course, that you have left your Fog Amount values at their defaults).

So, why aren't your lights showing through the fog? Well, first of all, in order to even be able to see the light source, you must have **Lens Flare** selected for the lights. The Lens Flare panel has defaults of its own (Figure 5), the least of which is a little button called **Fade In Fog**

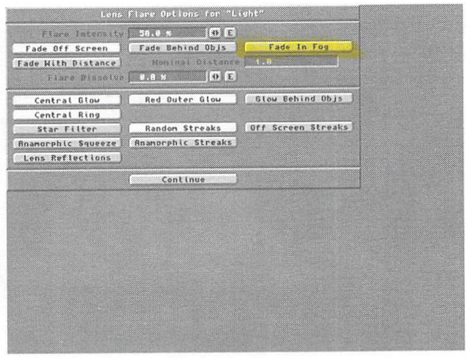

Figure 5

being deselected. With this button in the default state, any lens flare using these settings will not fade in fog, and therefore pokes right through.

Along the same lines is the answer to your last question. It sounds to me as if you are mistakenly creating an envelope for Light Intensity, when what you really desire is an envelope for the light's Flare Intensity (Lens Flare panel). These are two different things. The Light Intensity determines the amount of light that is cast onto other objects, while the Flare Intensity determines the brightness of the flare (if any). Either can be used without the other.

If all you wish is to have a flare dim down and disappear as it goes behind an object, simply select **Fade Behind Objects** and LightWave will automatically calculate this for you. When using Fade Behind Objects, the object in front of the flare must have Cast Shadows (Objects panel) enabled in order to calculate when the flare needs to be hidden.

LWP

John Gross is the editor of LWPRO *and a supervising animator for Amblin Imaging.*

And the Winner is...

continued from page 16

increase speed considerably, you must buy this option for an Indy. We will provide benchmarks for both cached and uncached Indy's in our future tests.

DEC Alpha 275—Digital Equipment Corporation's Alpha processor has found increased acceptance with the Windows/NT operating system. Companies like Carrera Computers Inc. and Aspen Systems make PC-style Alpha-based towers and other vendors are privately working on even more Alpha machines.

And the Winner is... the Alpha!

Notice how the ratios of the processor's render ratios remain relatively the same except for the Alpha. That 14.0 figure is not a typo. It really ray-traces that fast. In typical operation I have seen a consistent 10-12x speed increase in my favorite, nastiest animations.

The single processor Alpha is only slightly slower than two R4600 processors (i.e., Raptor Plus) at typical scenes but blows them away at ray-tracing. Because you do not have to buy additional RAM for a second processor and motherboard, the Alpha also costs much less.

Note that the Amiga version is currently priced at $695 (Version 3.5), but the price will increase to $995 with the release of Version 4.0. ScreamerNet costs the same for any of these modules and LightWave will be released at an announced price of $995 for SGI, MIPS, and Windows/NT PC and Alpha.

The Future

With Windows/NT making it simple for any processor to have software available for it (new products are being introduced weekly), and the increasing competition from hardware vendors, the future looks very bright indeed for LightWave render-addicts.

The Motorola 68060 is due out in February 1995 and should be as fast as a single R4400 processor. It also has the advantage of being inexpensive, and your existing 68040 card might even be easily upgradeable (contact your accelerator manufacturer). The MIPS chips have an entire hierarchy of processors at varying MHz ratings, so expect faster and faster processors to become available as the market demands. Preliminary reports on Intel's Pentium processor (90MHz) give it the same rating as a single R4400 (we'll have Pentium benchmarks in a future article) and faster Pentiums are due shortly. As if that wasn't enough, DEC is working on much faster chip that may be available by summer.

I'd also like to be able to provide benchmarks on LightWave Layout wireframe preview and Modeler redraw rates, but that will have to wait for the release of LightWave 4.0...I can't wait!

LWP

Grant Boucher is a supervising animator for Amblin Imaging. He is currently working on CGI effects for the Star Trek: Voyager *pilot. He would like to thank Allen Hastings and Stuart Ferguson for the render kernels used in these tests and Phil Hice, Amblin's hardware master, for his patience and support in getting all of these wonderful machines up and rendering. Grant can be reached on CompuServe at 74237,1146.*

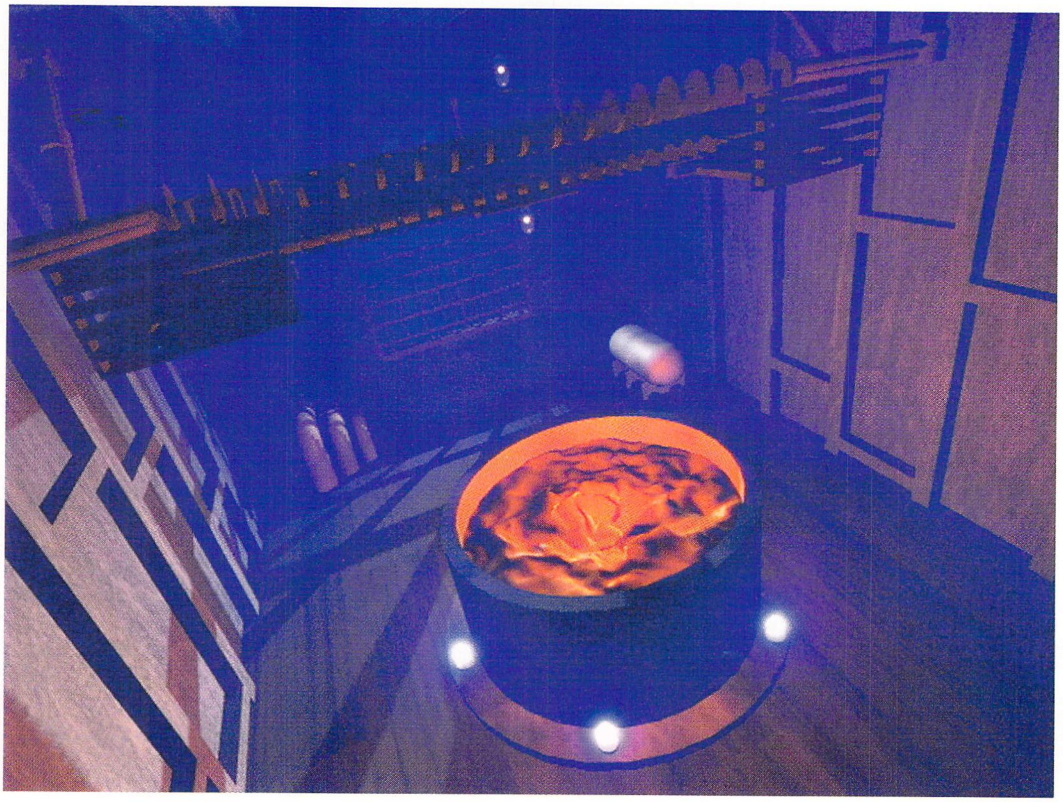

Warehouse

A PhotoCD image of a windowsill acts as a diffusion map for the walls, woodbeams and floors of this warehouse. A PhotoCD image of rocks is used as a displacement map for the liquid in the tank.

Copyright 1994 Dan Ablan

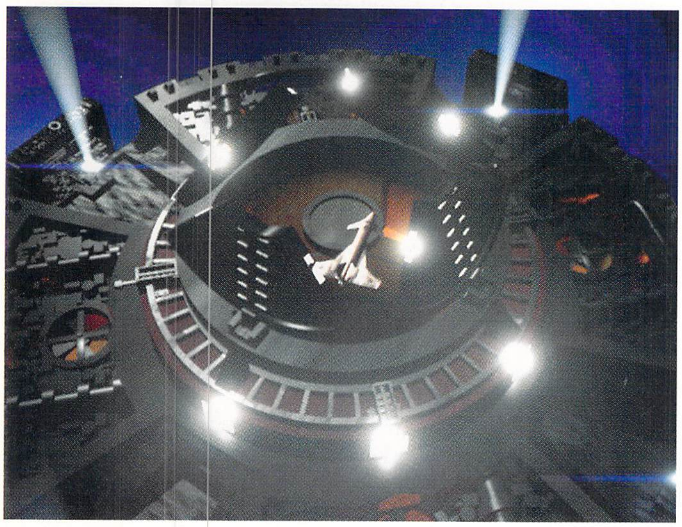

Building Hatchway

Lens flares over the glowing panels do just the trick.

Copyright 1994 Mojo

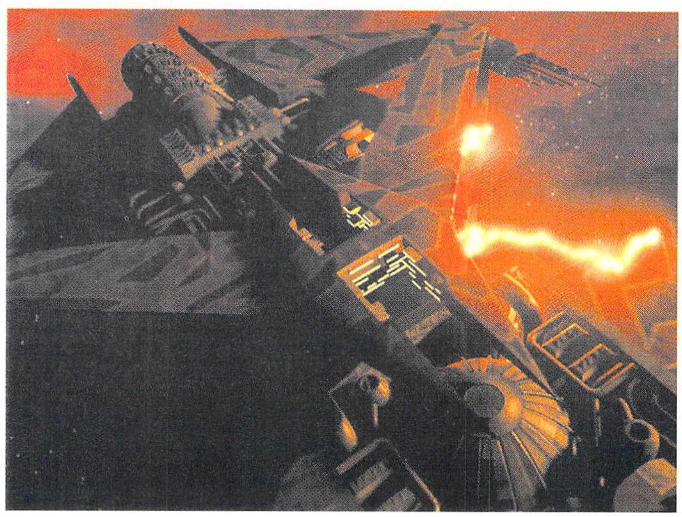

Narn Cruiser

The damaged Narn cruiser from *Babylon 5*.

Copyright 1994 PTEN Consortium, Inc.

RENDER GRAPHICS AT THE SPEED OF ALPHA ON A COBRA AXP 275 WORKSTATION.

ALPHA GENERATION™ There's no better way to burn graphics in LightWave™ or other applications. Introducing the Carrera Cobra AXP 275, the workstation leader in price and performance. Run UNIX® and Windows NT™ frame rendering, animation, multimedia and 3D graphics applications faster than you've ever seen on the power of a 275MHz Alpha™ processor—one of the 64-bit RISC rockets from Digital Semiconductor, a Digital Equipment Corporation business. With the blistering performance of an Alpha-powered Cobra, you'll generate digital images in minutes instead of hours. You'll get more done in a day. Maybe even get home on time for a change. And the Cobra AXP 275 comes in a variety of configurations, loaded with built-in PCI SCSI-2, PCI Ethernet, PCI video, PCI and ISA slots, CD-ROM, and more. Call or E-mail us for details. Then get ready for a workstation that really cooks.

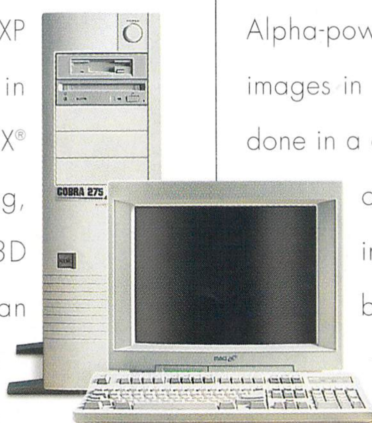

CARRERA Computers, Inc.

23181 Verdugo Dr., Building 105A, Laguna Hills, CA 92653 • 800-576-7472 • e-mail: CARRERA1@DELPHI.COM

© Carrera Computers, Inc. 1994. Digital, Alpha, and AlphaGeneration are trademarks of Digital Equipment Corporation. Other names are trademarks of their respective holders.