Faking Volumetric Lighting

## Making the PERFECT Demo Reel

TIPS and TRICKS

Two NEW Columns!

Bison

This Thornton Dynamics AV-12 Bison was built in LightWave by Foundation Imaging's visual effects director, Ron Thornton. Bison was constructed as an experiment in photorealism and is the end result of several weeks of tedious R&D to discover what makes a model seem authentic. Look closely at where the wings meet the body: the subtle specularity at every joint, which adds greatly to the plane's realism, can only be achieved with smooth surfaces. More details on the creation of this object will appear in a future issue of *LIGHTWAVEPRO*. *Copyright 1995 Ron Thornton*

Combustion

An exploding wave effect produced with Fractal Noise, Luminosity and Transparency. See "Explosion Shockwaves," page 6.

Copyright 1995 Lloyd Eldred

Hit and Miss

A spaceship escapes a large explosion. See "Explosion Shockwaves," page 6.

Copyright 1995 Lloyd Eldred

EDITOR'S MESSAGE

by John Gross

L IGHTWAVEPRO is going through some changes.

With the introduction of LightWave to the PC world, *LWPRO* will have an expanded audience so we'll be making some changes to help meet the needs of additional readers, while striving to maintain the level of information you've come to expect. What these changes are will be determined down the road. If you have any suggestions or comments, let us know!

Starting with this issue, you'll see two new regular columns. Dan Ablan's "lwpro@internet.online" brings you the latest LightWave questions, tips, tutorials and information from the Internet's various LightWave sites. "LightWave Tech," by Philip Hice, Amblin Imaging's resident technical guru, addresses hardware-related LightWave concerns. Both columns should help you get the most out of LightWave.

If you subscribe to the *LWPRO* disks, you'll notice that starting with last issue, we supplied the content on a PC-formatted disk. We went with this solution because of the growing numbers of LightWave users using PCs. Since anyone running AmigaDOS 2.1 or greater can access PC disks using the built-in version of CrossDOS, PC-formatted disks have now become the common denominator between Amiga and PC users of LightWave. The question we have is whether to supply content on 720K disks or 1.4MB high-density disks. Ideally, we would always like to use the high-density disks, but we realize that many non-Amiga 4000 owners may not be able to read a high-density disk. A short survey on the Internet revealed that the majority of users prefer to have the *LWPRO* content supplied on 1.4MB disks.

Since I realize that many users are not on the net, I would appreciate any feedback that you have in regards to the question of disk formats. In the meantime, we will be supplying the content on 720K PC-formatted disks, the lowest (if not the best) bridge between the Amiga and PC.

Speaking of Amigas and PCs, many users have come up with a variety of

TABLE OF CONTENTS Vol. 3, No. 6

See the Light
Faking Volumetric Lighting, Part I

by Alan Chan

Being the digital visionaries we are, we all remember, sometime in our past, seeing that one marvelous sunrise that was actually worth a thousand words. That wonderful, warm, glowing ball of plasma, under a slight staccato of clouds, brushing away the cold of the night, the rays of the sun shimmering in the golden air...

As you sit outside on your lawn, LightWave running on the laptop in your hand, you wonder what causes that magnificent display. Particles in the air, you conclude. The rays of the sun are reflected from particles in the air and into your eyes, making the rays appear brighter than the other areas. And as the clouds shift and move, they block out the rays of the sun from some particles and reveal others. It is this Brownian movement that creates the marvelous display of shimmering rays of light.

Hmmm—rays of the sun, huh? Shimmering rays of light. That's just what I need to spice up my flying logo, you think. But how do you go about doing it? Well, short of going out and shooting a frame of it as a background image, you could try a couple of different techniques.

Volumetric Schmolumetric

The 3D-ish techie animator's term for "rays of light" is Volumetric Lighting. Simply put, the presence of particles in the air (smoke, for instance) intersecting with the path of a lightbeam gives it the illusion of having volume, hence the term "volumetric." The more interesting side effect of this phenomenon is that if an object travels within this "volume" of light, it casts long visible shadows in the air. Imagine a volumetric light behind a logo, casting shimmering rays across your screen, and you can see why volumetric effects can sometimes be an effective element in your animation.

We cannot fake volumetric lighting with LightWave by brute force alone. It is theoretically possible to simply fill the scene with millions of half-transparent particles and let them pick up the light, but you'd probably run way out of memory, and your render times would go through the roof (and your neighbor's as well). So let's be smart. Let's only quantify the *edges* of any given volume, where the light meets the shadow. What am I talking about? We'll go through some tutorials, and

hopefully the meaning of what I just said will become more apparent through some hands-on work.

The Tutorials

There are several ways of faking volumetric lighting, and in your eventual use of it you will be utilizing some, or all, of the following techniques, depending on light and camera positions, what happens in your animation and what you want to achieve. In rough order of complexity, they are:

(1) Visible lightbeams and oscillating lightbeams
(2) Shimmering logo shadows
(3) Objects with volumetric shadows

The key to good volumetric lighting is a firm grasp of the basic foundations of lighting and surface settings mimicking volumetric shadows, so this month we'll spend our time understanding the concepts and creating some lightbeams. Next month, after you've had some time to play with your ideas and everything's started to soak in, we'll tackle shimmering logos and volumetric shadows. As a primer, you may also want to re-read John Gross's article on creating realistic lightbeams, called "Spotlight On Reality," in the July 1994 issue of *LIGHTWAVEPRO*.

Let's begin.

(1) Visible Lightbeams/Oscillating Lightbeams

The first part, the visible lightbeam, is easily achieved. In fact, there is a lightbeam object included in the Objects drawer of your LightWave package, based on a concept courtesy of the honorable Mark Thompson, of Fusion Films fame. Revisiting this object provides the basic principles we need.

Load up the object (found in Objects/Tutorial/Lightbeam.obj), keyframe it so that its length is visible and test render a frame. Then examine the surface settings. Note the surface parameters on the lightbeam cone:

Surface Color	255,245,180
Luminosity	0%
Luminosity Texture Type:	Grid
Texture Falloff:	0,0,34.0
Texture Value:	85%
Line Thickness:	1.0
Additive	On
Diffuse Level	0%
Edge Transparency	Transparent
Edge Threshold	1.0
Smoothing	On
DoubleSided	On

All other values are at their default.

The two critical items on this list are the **Luminosity** texture and the **Edge Transparency.** The Luminosity texture here (Figure 1) is in fact a grid texture, but where's the grid? Notice that the line thickness of the grid is 1.0, which means that the grid lines
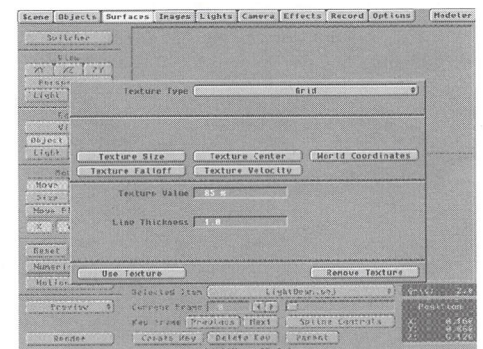
Figure 1: The Luminosity Grid Texture panel

are in fact so thick that they touch end to end and form a uniform surface. This crucial fact now allows us to add a **Texture Falloff** value to the texture, causing it to gradate slowly from 85% (from the **Texture Value** setting) to 0% (the setting on the main **Surfaces** panel) along the length of the light beam. Note that **Texture Falloff** values are noted in percentages over one unit (usually one meter in Layout), which means that a falloff value of 25, for instance, will cause a 100% luminous surface to gradate to the main surface panel setting for Luminosity (default: 0%) over four meters. Therefore, the settings described above designate a lightbeam that is 85% luminous at its source, slowly fading away to darkness in a little under three units. The **Additive** button then instructs LightWave to treat this surface not as a normal opaque surface, but as an additive function, where the computed surface
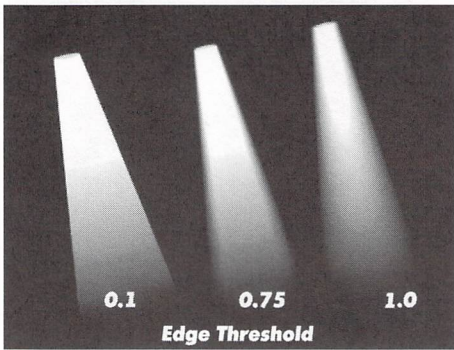
Figure 2: Lightbeams with differing edge threshold values

values are simply added to the RGB values behind the additive surface (like a lens flare).

Note that the **Edge Transparency** is set to **Transparent.** This makes the edges of the light cone transparent, giving it more of an ethereal quality. The sharpness of the transparency is varied by changing the **Edge Threshold** value. Figure 2 illustrates the effects of different edge threshold values. Note that values above 1.0 or below 0 do not make any additional effects.

Lightbeams 102

Now, with basic lighting behind us, let's try something a little harder. Like an oscillating lightbeam, the effect you'd get if you tried to shine a huge spotlight through a clump of trees. For this we'll need to go into our favorite paint program and make up a texture map such as the one in Figure 3. We'll save this as "Cookie.iff." (The term for a shape, or cutout pattern, used to make patterns with lights is a "cucaloris," or cookie for short.)

Now we'll replace the grid pattern with the cookie pattern, so that the alternating light and dark patches affect different parts of the lightbeam's luminosity. Make the texture move and you'll have shifting light-streaks, yes? Ah, but if it were only that easy. Sorry. First we have to map the texture along the length of the lightcone, for which we'll fall back on one of the earlier LightWave morph tricks (a variation of a Mojo-of-*Babylon 5* trick).

Figure 3: An IFF file used as a cutout pattern

- Begin by going into Modeler and, using the Disc tool, extruding a cylinder, making sure that the center of one of the ends of the cylinder is at 0,0,0. For this example, let's extrude a cylinder four meters in length. Select both disc polygons at either end of the cylinder (Figure 4) and cut them off, leaving just a hollow tube. Set the surface name of this object to "OscLightSurface" and save, then load it in Layout. (Amiga users can export it to Layout). Call this object "OscBeam.lwo."

Figure 4: Preparing to cut off the ends of the lightbeam cylinder
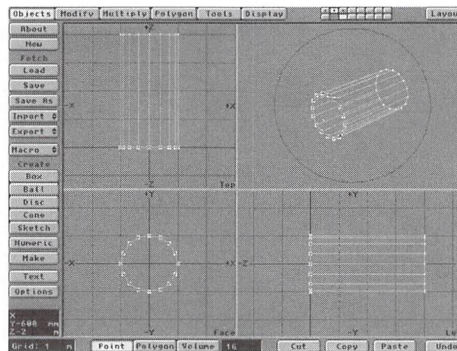
Figure 5: Selecting points on the zero plane

- Select all the points of the cylinder that sit on the zero Z plane (Figure 5) and hit (Shift+H) for the **Size** tool. Click on the **Numeric** button and enter the following values into the requester:

    Size Factor: 0

Center XYZ: 0,0,0

- Hit OK. You will see that all the points you selected are now converged at the 0,0,0 point, and we have, once again, a cone representing a lightbeam. Do not use **Merge** Points on this object! Simply bring it to Layout as a new object, calling this one "OscBeamMorph.lwo."
- Returning to Layout, enter the following values for the surface "OscLightSurface":

Surface Color	255,245,180
Luminosity	0%
Luminosity Texture Type:	Grid
Texture Falloff:	0,0,25.0
Texture Value:	85%
Line Thickness:	1.0
Additive	On
Diffuse Level	0%
Transparency Texture Type:	Planar Image Map
Texture Image:	Cookie.iff
Texture Axis:	Z Axis
Automatic Sizing	On
Texture Center:	0,0,0
Texture Velocity:	0,0.02,0
Edge Transparency:	Transparent
Edge Threshold:	1.0
Smoothing	On
DoubleSided	On

- From the Objects menu, make the OscBeam-Morph.lwo object 100% dissolved. Then select OscBeam.lwo as the current object, and enter these settings into the Metamorph section of the **Objects** menu:

Metamorph Target:	OscBeammorph.lwo
Metamorph Level:	100%

Figure 6: The completed oscillating lightbeam

Explosion Shockwaves

by Lloyd B. Eldred

I t's a peaceful day in deep space. Suddenly, an incredibly intense flash appears, accompanied by a very deep thud. Soon, an expanding, burning shockwave begins racing away from the source of the explosion.

Science-fiction movie fans may recall the spectacular explosion of Praxis in *Star Trek VI*. Or perhaps the explosion of the pyramid ship in *Stargate*. A similar effect also appeared in the *X-Men* animated series. This two-dimensional burning shockwave effect is quickly becoming a science-fiction staple.

The physics of such a wave are dubious at best, but the effect is still stunning enough to add to your repertoire. This article will cover the details of the basic effect and hopefully give you lots of ideas for using and improving on it.

Strips and Rings

The construction of the wave is rather similar to the Saturn's Rings tutorial in the LightWave manual (pages 23-25). On the off chance that you, like me, didn't get it to work on your first few tries, I'll run through it and attempt to clarify surfacing and morphs at the same time.

Figure 1: Setting up the wave strip

The first step is to create a rectangular strip. Use the **Box** button (**Objects** panel) and the **Numeric** options shown in Figure 1 to create a subdivided strip.
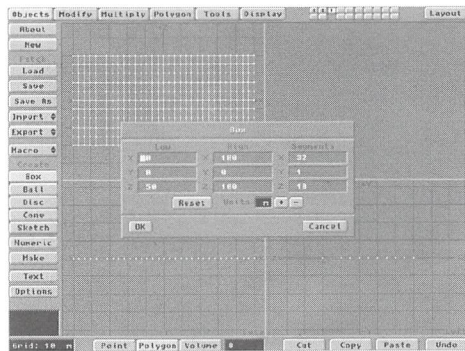
The number of segments in the X direction determines how round the ring will appear. The number of

segments in the Z direction control how smooth a gradient our textures will have from the leading edge, or outside, of the wave to the trailing edge, or inside. The polygons will be facing downward, but that's OK.

The next step is to assign surface names to the strip. Use the volume selection tool to individually select each of the 10 horizontal rows one at a time. Assign each a different surface name using the **Surface** button (**Polygon** panel). Figure 2 shows the topmost row selected, to which I assigned the name "Ring.LE." I chose the name "Ring.LE-1" for the next row in, "Ring.LE-2" for the third row, etc. The bottommost row (trailing edge) will be called "Ring.LE-9."

Once you're finished, choose the polygon selection tool, then hit the **Stats** button (**Display** panel or w). Select the surface name pop-up menu that (probably) reads "Default." Verify that there are 32 polygons with each of the 10 "Ring.LExx" names. If there are missing polygons, pull down to the problem name, then hit the (+) beside the "**with Surface:**" option. That will select and highlight the polygons with that name. Identify and fix any problems before proceeding.

Because we only have 10 different surfaces, there will be fairly visible lines between each section on the finished frames. I decided to break up these lines somewhat by jittering the internal faces. If you choose to do this, use the Volume selection tool to select everything except points in the leading edge (+ Z side) and the sides (including the trailing edge is fine). Use the **Jitter** button (**Tools** panel) to move the inner

Figure 2: Surfacing the strip

points around some. I used the Gaussian option with the radii set to X = 2 m, Y = 0 m, Z = 2 m. Figure 3 shows the result of this operation. Save the completed strip object. I named my strip "WaveStrip.lwo."

The time has come to make our ring-shaped morph target. Copy the strip to another layer using the copy and paste tools. Hit the **Bend** button (**Modify** panel). Press the **Numeric** input button (n), and configure the settings to match those shown in Figure 4.

Figure 3: Jittering the internal points
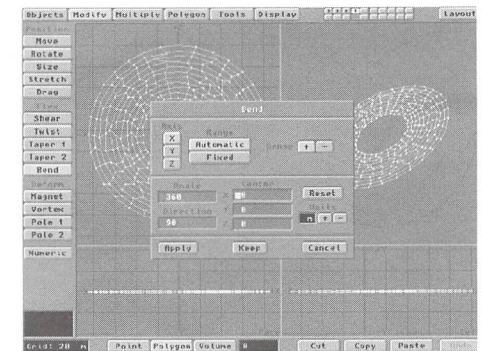
Figure 4: Bending the strip into a ring

The result should be a nice ring. Using the Center macro command, center this object, then save it as "WaveRing.lwo." [Editor's note: Modeler versions before 3.5 do not keep the point order of an object intact when you copy and paste in another layer. For those with pre-3.5 Modeler, perform the Bend opera-

tion in the same layer on the original object (after saving, of course).]

Now, switch over to Layout, and load the two objects from within the **Objects** panel. Adjust the camera to a position above the objects with both visible, so you can watch what's going on (remember to set a keyframe for the camera position). For the wave strip object, set the **Metamorph Level** to 100% (**Objects** panel) and the **Metamorph** Target to "WaveRing.lwo" (make sure **Object Dissolve** is 0%). If you will be setting up a scene using ray-traced or cast shadows, you can select the shadow options you want for this object to save rendering time. For the wave ring object simply set the **Object Dissolve** to 100%.

At this point, you should see just one copy of the ring, which is actually the strip morphed into the ring shape.

Burning Surface Questions

The main trick to this special effect is proper surface settings. The leading edge is bright and luminescent. As we move to the trailing edge, the surface gets slightly darker and dissolves further and further out, using fractal noise for transparency.

The most important thing to remember is to set the Fractal Noise textures to use **World Coordinates**. This setting allows each section of the rings to closely match its neighbors. Plus, when the ring is animated, the features of the texture will flow smoothly from one section to the next.

For the Ring.LE surface, I used the following settings:

    Color: 46, 20, 255
    Luminosity: 80%
    Diffuse Level: 10%
    Specular Level: 0% (Texture Mapped)
    Glossiness: Medium
    Reflectivity: 0%
    Transparency: 0% (Texture Mapped)

Set the two texture-mapped channels to the same values:

    Texture Size: 10, 10, 10
    Texture Center: 0,0,0
    World Coordinates: On
    Texture Falloff: 0,0,0
    Texture Velocity 0,1,0
    Texture Value 10%
    Frequencies: 3
    Contrast: 1.0

Save these surface settings as "Wave.surface." Load it into the other nine surfaces. We'll be using it as a basis.

The explosive wave object that appears in the color pages varies smoothly in color down to 0,0,165. The **Luminosity** decreases by eight on each surface, until it reaches a value of eight at the trailing edge (Ring.LE-9). I decided to use a little extra transparency on the inside region, so Ring.LE-6 has a **Transparency** setting of 20%, Ring.LE-7 has 30%, Ring.LE-8 has 40%, and Ring.LE-9 has 50%.

The most important surface setting for this effect is the **Fractal Noise** Texture Value for the **Transparency** Map. Though it's not well documented, many of LightWave's settings can be set as high as 400 percent. The wave makes good use of this feature

to vary smoothly from a solid wavefront to a trailing edge that has just a few wisps of smoke. Moving from Ring.LE to Ring.LE-9, my settings for the Texture Value were 10, 60, 100, 140, 180, 220, 260, 300, 340 and 380.

Now would be a good time to save your two objects. Hit F9 to do a quick test render of the current frame. If you don't see anything, your camera is on the wrong side of the wave. Move to the other side and re-render.

Animation Extras

It sure is pretty, isn't it? But there are still a number of things missing.

You certainly need a nice starfield for a background. LightWave comes with two that will fit the bill. In the Objects/Space directory, you'll find the ActualStars and the RandomStars objects. Select one and load it.

That helps. Now, what is it that's exploding? A spaceship? A planet? A star? A whole galaxy? Design something, or load one of the spaceships included with LightWave. Parent your wave ring object to your victim. Next, use the **Size** feature to reduce the scale on the wave strip object to 0,0,0. Point the camera at the victim and set up a few moments of animation of just your victim object: the calm before the storm. Why does the victim explode? You may want to add a laser beam or some other weapon striking it just prior to the explosion. Or, your poor ship may just suffer a spontaneous core breach. It's up to you.

The time is here. Our victim must go, in the most spectacular way possible. One important addition is a brief but very bright flash of light. LightWave's lens flares are the thing to use. Add a new light (**Lights** panel). Rename it "explosion." Set the intensity to 0, unless you have some other objects in the scene that you want the explosion to light up, in which case it should be a point light with an intensity envelope like that described for the lens flare.

Hit the **Lens Flare** button and select **Flare Options**. For starters, use the default settings, except for the **Flare Intensity**. Set up an envelope (**E** button) that keeps the flare invisible (0 intensity) until the explosion occurs. Set a key for the frame, also with a 0 value, immediately before you want the explosion. Make a key to start the explosion with a large value. Something like 300 percent will turn the screen nearly completely white. Quickly ramp the flare back down with a few more keys. You may want to leave it flickering around a value of 10% to indicate a small burning ember. Examine your envelope and adjust the tension, or set the Linear option (**Spline Controls**) to fix any odd features in the curve.

Exit the **Flare**, then **Lights** panels. Parent the explosion light to your victim. Go into the **Objects** panel and set up an **Object Dissolve** envelope for the victim that takes it from 0% dissolved just before the explosion to 100% dissolved as the explosion starts. Finally, set up some keyframes for the size of the explosion wave. Take it from 0,0,0 as the explosion starts to something that fills your shot and moves past

the camera over the course of a few seconds. The exact values to use depend on your camera settings and the size of your victim object. Save your scene, then experiment away!

I recommend doing several test runs in low resolution with Antialiasing off. This allows you to quickly experiment with all of the various surface and timing issues without waiting too long. My stock Amiga 4000/040 takes about eight minutes to render a Medium Resolution, Low Antialiasing frame of the wave. When you're ready, open all the stops. Turn on Antialiasing and Motion Blur. Depth of Field might also make a good addition. Depending on how your camera is moving, particle blur may also be appropriate. Hit the Render button, check that you have all the options set the way you want (including Automatic Frame Advance), then head off to bed. Don't worry—there won't be anything to see for hours.

Explosive Possibilities

The finished product looks pretty amazing. Before I leave, I thought I'd share a couple of ideas to get you started on spiffing it up even more.

The Texture Velocity setting is worthy of experimentation. The X velocity causes some spiraling of the flames. The Z velocity setting controls radial movement of the texture. The Y velocity setting causes the flames to evolve or boil over time. Don't go overboard. A little movement goes a long way. Remember the dimensions of our object and that the velocity is the distance the texture moves for each frame. If you're animating at 30 frames per second, it doesn't take much velocity to be noticeable.

The breaks between each band of the wave are somewhat distracting if you get too close to the wave. This problem could easily be lessened by adding more rows of polygons in the original objects and thereby having more variations in the surface changes. Perhaps an easier route is to create the original objects with only one surface name and use a combination of **Texture Falloff** and **Fractal Noise** applied to **Luminosity** and **Transparency** to create the wave effect.

Burning debris flying from the explosion past the camera could be a nice addition. Secondary waves were used in the *Star Trek* shot, as well as a very rapidly expanding 3D sphere. This wave shot also had some thickness to its leading edge.

Add in other space animation suggestions that have appeared here in *LIGHTWAVEPRO*. Mojo's tips for improved starfields, easy laser beams and burning starships would all make wonderful additions to this shot.

Above all, remember that this tutorial is designed as a jumping board for experimentation. Be creative!

LWP

Rocket scientist Lloyd Eldred is head of an amateur film group, Galtham Films, in his spare time. Its current project, Space Rogues, T.S.E., *may be finished by the turn of the millennia. Eldred can be reached at galtham@universe. digex.net.*

Digital Cinematography

by John F. K. Parenteau

Here at Amblin Imaging, the pattern is almost predictable. Every once in a while—usually on special occasions such as somebody's birthday—we gather and watch demo reels. There's a certain level of contradiction happening when we gather in the "Big Room" (our editing and screening room) for a round of reels from unsuspecting animators. Everybody is so eager to rip into the new tapes sent in, almost as if they want to find bad stuff. But, in truth, I think there is a sense of fear in the air as we have our laughfest, cheering, booing and throwing things at the screen. Most of us realize that, at one time, we were on the outside looking in, hoping somebody would take a liking to the images we sent in to Amblin. Many of our animators know that in the beginning their stuff wasn't all that great either. So, as we do our impersonations of Joel on *Mystery Science Theater 3000* to the demos, somewhere in the back in our minds, we are really looking for something not to make fun of, perhaps something positive among the sea of the unacceptable.

Actually, not everything is unacceptable. Quite a lot of work we receive has some considerable merit and is worthy of further consideration. The problem usually arises in the mode of presentation. Unless you have the opportunity to see a large number of demo reels, it is impossible to evaluate your tape with an objective eye.

Creating a demo reel isn't extremely tough. All of us see a type of demo reel as we sit in a theater or watch television at home every day. Think of your reel as a type of commercial or theatrical preview. Both are trying to sell something. Both are put together in an organized fashion to be entertaining and informative. Your work is much harder, however, since you cannot use dialogue to sell yourself. It is important to construct both an entertaining and brief example of the product. Let's examine some of the more general mistakes some animators make when constructing their reel.

Length

There is a reason why most reels are about two to three minutes long: that is the average length of a song. The last thing you want is to have someone sit through a five- or 10- minute reel, no matter how great you feel about your work. Imagine *Jurassic Park* with no

A good example of a bad example. This title card is difficult to look at and too confusing.

action—just a long series of shots of fantastic computer-generated dinosaurs. For about two minutes you would be impressed. "Look at that great work! Can you believe how realistic they look?!" But after a short time you'd be checking your watch and thinking about getting that free popcorn refill. No matter how interesting the work, you always want the audience yearning for more, especially when it is your reel. Yearning for more means they want to meet—and potentially hire—you! So keep the tape short and sweet. But remember, too short is just as bad. Try to have at least a one-minute reel, even if this requires you to create more animations.

Variation

Many animators have DPS' Personal Animation Recorder and use it to output their work. Many of those people have the horrible habit of looping animations and playing them over and over again. Not only is this far too repetitious, but that's what a rewind button is for! If viewers want to see the animation again, they can stop and rewind. The first impact is the most important. Make sure your animations have enough diversity so the viewer can see a few different moods you have created. Just as a cameraperson's reel has exteriors and interiors, day and night, it is crucial to show off all of your talents.

Music

And how about that two- to three-minute song? Never, ever use a song with lyrics. Words give the audience something else to think about, when you would rather have them thinking about your great work. We often find a soundtrack that has a good variation of highs and lows, fast segments and lulls. Try to avoid music that is too familiar. For example, many people like to use classical music. Though it is pleasant to hear, it doesn't carry the impact of a hard-hitting new piece, something unfamiliar to the viewer. Music carries a lot of weight with any visual, and more so with a demo reel. Pop in one of your favorite action moves at home and try to watch the best scene with no sound. Kind of loses something, doesn't it? Good music can raise the excitement level up high, and some of those

little mistakes you had to leave in the reel can be easily overlooked. Movie themes seem to work best, if carefully chosen. Most music scores in motion pictures are hidden beneath dialogue and sound effects. If you choose the right piece, viewers will feel a sense of familiarity, but not enough to be distracting.

Title Card

I've always felt that the name card at the head of the reel should be as simple as possible. So many young animators try to create some fantastic animation just to display their name. In truth, the card at the beginning is usually ignored. It's the name at the end that is important. Assuming the viewer liked your work, it is from that card that they extract your phone number or address for further contact. Personally, I like a simple card: white letters on black, in a cool and tasteful font. Others may try something a bit more interesting, like a fade up or wipe on. The important point here is to avoid spending all your time animating your name card. Your effort is better spent making your animations in between the cards shine.

Laydown

No matter how great your work is, if you put it on a bad tape with glitches and hits in it, your prospective employer or client is going to quickly shut it off. If you are located in a city that supports any form of video production, it is possible to buy short VHS tapes (five-minute tapes, for example). Packages of 120-minute tapes can be costly, forcing you to recycle to save money, but recycling creates bad tape spots. Try using a tape only once to prevent this; five-minute tapes make it much less expensive. If you do not have access to an editing system or industrial dubbing rack, you may be forced to "hard record." Imagine you have two VCRs hooked up to each other. When you press play on the source and record on the recording deck, you have just performed a "hard record." On VHS, the first 15 seconds of the tape will have a rolling bar of color moving across the screen. It's important that your material doesn't start until well after this bar has passed and the image has stabilized. Make sure the head of your tape has at least 30 seconds of black to get past this glitch. In a perfect world, a professional facility would be making your dubs from a D1 master, but most people can't afford D1. Since you are working in a digital medium, try to lay off your work to the highest tape medium possible. Preferably, you have access, through a rental

house somewhere nearby, to a Betacam deck. At least try to master on 3/4-inch. An investment in higher tape quality will pay off in cleaner-looking demos with a more professional appearance.

Labeling

It may seem trivial, but a proper label can make all the difference when a client goes to view a demo tape. One's first instinct is to view the tape that looks the most professional, and leave the unlabeled or handwritten 120-minute cassette for last. It all takes money, but you must invest to profit, and this is no exception. Any local printer can make up good-looking labels with your name and telephone number for a reasonable price. If you manage to buy short tapes, generic sleeves usually come with them.

That about covers the basics. Though most of the above sections outline the presentation side of demo reels, there is a much more fundamental reason why many of them fail. Ask yourself this question: What is my favorite CGI shot from TV or movies? Once you've selected something, examine the reason why you chose it. Most likely, the work is slick, with no errors, has smooth camera motion and is efficient in its purpose. These simple reasons should be the goal you strive for with all your work.

Is It Slick?

It's very easy to fall into the "slick is impossible on my budget" trap. But you must remember that slick doesn't mean expensive. Slick means doing the work as best as possible. For example, we often receive reels that show objects with low polygon counts or no textures. Though memory is an important exception to the rule, nobody is going to be impressed with a wild shot with primitive objects. If memory is a problem, use procedurals. Many *Babylon 5* textures rely heavily on Allen Hastings' built-in surfaces, and they look great! There is no excuse for lack of textures.

Modeling of your objects is, however, dependent on the RAM of your machine. Work carefully to take advantage of the memory you do have by building your model in the areas that are seen closely. Don't waste too much detail in modeling if the audience won't be close enough to care. That's where image maps can fill in.

Motion

Camera and object motion are often the most essential areas to focus on. Most potential employers quickly

overlook shoddy modeling or texturing, but if the motion is bad, the errors pop out fast! Camera motion in particular is a crucial part of your talent and, regardless of what you feel your strengths are, it must be presented in a good light. Remember, you may not be able to see a tape glitch before it's too late, but with the preview function in LightWave, you can always see a camera error. Simple adjustments such as tension at the head and tail of the camera motion prevent abrupt stops and starts. Take the time to finesse the camera.

Object motion is equally important. Remember that you are telling a story with your shot. Don't use unnecessary motions just because they look cool. Try to be as economical as possible while still accomplishing the point of your shot. More often than not we see reels with wild camera and object motion that may look great to the animator's eye, but don't make any sense to the audience. Though you may think that a crazy 180-degree pan and zoom in is cool, it looks gratuitous and unnecessary to the people who want to hire you for their television show. Work on telling a story, whether it be a one-shot piece or a sequence, so employers feel confident you can produce the shot they need.

For the most part, what we look for in a reel is some sense of an accomplished talent. Nobody, even those of us at Amblin, ever started off with a reel that was perfect. That only comes with years of careful study and experience. So give yourself a break. Try to focus on a particular area that you feel your talents lie in, while allowing for some error in other areas. This doesn't mean blowing off textures if your strength is animation. What it means is simply being accomplished at those things you aren't great at. The textures or modeling may not knock anybody's socks off, but if there is a good sense of motion, your point has been made. An attention to detail is obvious, even if that detail isn't great.

Back in our "Big Room," we tend to make fun of everything. In a way, it's how we make our own shots better. Naturally, no matter how good a reel is, we still critique it. But deep down inside, we take note of the talent we can see on the screen, and log it away. When the time does arrive when we need a new animator, the jokes are forgotten and the quality reel stands out.

LWP

John F.K. Parenteau is a vice president and general manager at Amblin Imaging. His first demo reel wasn't so hot.

Reader Speak

by John Gross

It's been quite awhile since we've had a "Reader Speak" column in *LIGHTWAVEPRO*. The fact that I have been very busy these past few months is mostly to blame, but I felt that this month we needed to address some of the questions that have been coming in.

Q: A while back, I received LightWave version 3.5 for the Amiga. I loaded it onto my Amiga 4000 and proceeded to be amazed by the sheer excellence of this superb rendering package.

Unfortunately, I have a great deal of learning to do. Layout isn't much of a problem, but Modeler is. There is much to learn, and very few guides to assist me. I'd rather not rely on videotapes and can only hope that a LightWave tutorial book is released soon. As it turns out, *LIGHTWAVEPRO* is one of the only good LightWave information sources available.

I really enjoy *LWPRO*, and I realize that, as the name suggests, it is a magazine that outlines professional techniques. Despite this, I am still requesting that more time be spent on basic tutorials showing how to effectively and efficiently use Modeler and Layout. Hopefully, a number of extra pages could be added to the magazine in order to cover this material. After all, this is likely the most expensive computer magazine in the world (per page).

Your instructional format is quite good, yet it isn't the clearest. The best format I have ever seen in any magazine is the one that Amiga Format (Future Publishing) used in order to give instruction for the Image rendering package. You can find these articles in issues 53-57, the first of which is dated December 1993. Please adopt this format.

And now, one last request:

I am trying to duplicate the type of Cray-rendered images that appeared in the motion picture *Tron*. My first project was supposed to be simple: build a recognizer. This object is quite simple in form. The problem lies in rendering the object with a solid dark tone and a thin, luminescent/glowing red outline of light around it. I can't seem to get LightWave to do this. The same effect is used on the *Tron* city, and a more advanced effect is to place a "trace" pattern over all of the city objects near the end of the film.

How do you create this outlined polygon effect, exactly the same as the original Cray renders?

Thank you for your time. By the way, keep *LWPRO* going. This is an absolutely fantastic magazine. Also, to Mojo, keep those excellent articles coming. You're a demigod!

W. Jared Brookes
Calgary, Alberta, Canada

A: Your letter brings up some interesting topics. One of the things that we at Avid Media Group have been discussing lately relates to the direction of *LWPRO*, especially with many new LightWave users coming on board. From fairly early on, we realized that the majority of *LWPRO* readers may not be "high-end" users, and we have tried to balance the content of the newsletter to reach everyone. At times we've succeeded; other times we've failed. With the advent of many "newbie" users, we would like to have informative content geared toward all users without compromising the goal of making eventual LightWave pros out of everyone.

In regards to this, *LWPRO* will be undergoing some changes soon. How it changes partially depends on the readers. We welcome feedback of all kinds. If you have e-mail access, we can be reached at jgross@netcom.com and avid@cup.portal.com.

In regards to your *Tron* question, there are a couple of ways to achieve a similar effect. If you happen to be running LightWave 4.0, the best way is to simply assign a glow effect to the surface of the objects of your recognizer, and then select **Enable Glow Effect** in the **Effects** panel. You can then input a **Glow Intensity** and **Glow Radius** (in pixels) that you wish the glow to contain. Remember that because Glow Radius is measured in pixels, if you are rendering tests in **Low Res** and then bumping up to **Medium Res** or higher for your final image, you need to increase the Glow Radius to achieve similar results, since the pixels become smaller as you increase your **Basic Resolution**.

If you are not running 4.0, you could parent a second, slightly larger recognizer object (with different surface names) to the original. Make sure the surfaces of the parented object are mostly transparent and you are using **Transparent** for **Edge Transparency.** While this approach isn't as effective as 4.0's glow effect, it can look pretty good.

I won't tell Mojo about the demigod part, as I'm sure it will just go to his (omnipotent) head.

Q: I just read your Editor's Message in April's *LWPRO,* and though this publication is immensely helpful and I love my subscription, I have to respond to some things said about the pre-release version of LightWave. You said that it is close to the actual commercial version that is coming out in a couple of weeks, but have you actually seen it? I own the pre-release version and I certainly hope that the commercial version packs substantially more than this one. This so-called "upgrade" not only lacks the ability to view the rendered image, but you cannot render to useful IBM formats besides Targas. Who uses IFFs on PCs?! Where are the macros in Modeler that we have grown fond of and even dependent on? Why can't the image file requester recognize or load standard IBM format files such as GIFs, JPEGs and BMPs? Why, when you click (in the Record panel) on the Save Anim button, does it say "only available on Amiga version"? I know that Anim format is an Amiga format, but, in that case, why is it there taking up space? Shouldn't it have been changed to some kind of Flic format saver?

I want to take time to thank you for advice you've given me, and I don't want to sound ungrateful (just frustrated and disappointed), but when you praise the new LightWave for its great beta version that you have access to, all I can say is, try the pre-release. You pay nothing for being a beta-site, but we users paid $150 for a sizable downgrade!

I'm not one of the people claiming to have lost business for depending on the new product. I switched from my Amiga because I had a golden opportunity fall into my lap when someone who needed an Amiga offered to buy mine for $700 more than dealers were normally selling them at, because Amiga dealers had dried up. I thought, great time to move to a Pentium. LightWave is coming out soon on PC, so I won't wait too long. To be honest, the wait wasn't long—I just wish it had been worth it. I have a great opportunity to submit some animations to a well-

known expanding game company in Cambridge called LookingGlass, which used to work with Origin Systems. Maybe you've heard of them—they made SystemShock, a best-seller. They are interested in me and want to see more. I wish I had my Amiga. The pre-release, with all its shortcomings, also crashes often! Tech support has no answers, either.

(Sigh) I guess I'll have to switch to the mega-expensive 3D Studio until LightWave becomes the incredible, versatile, pleasure-to-use tool that it once was.

Tearing my hair out in Amherst,
David Bryant
Amherst, MA

A: When I wrote April's Editor's Message, I mostly meant that the *tools and functions* that shipped with the pre-release were pretty much finished, and the only things that needed adding were plug-ins and a few extra little things. Re-reading that column, it's apparent to me that I wasn't very clear. It's also apparent that I wasn't very correct.

At the time I wrote it I was using the pre-release version. Now that I look back (with the additional hindsight that new beta releases provide), I can safely say that most, if not all, of your concerns have been, or are being, addressed.

Full-sized rendered images are now viewable. Elastic Reality's HIIP architecture will allow for multiple save and load image formats. Modeler plug-ins will support the important macros that you've gotten used to (it's amazing how much I took Center for granted). To top it all off, Bones and Inverse Kinematics have been greatly improved and some other new features (such as distributed rendering) have been added.

It's important to remember that the pre-release of LightWave was just that, and that you will be getting a free upgrade to the final shipping version (plus manuals, etc.) when it is released. Hopefully, when you receive it, you'll have renewed faith.

LWP

John Gross is a supervising animator for Amblin Imaging and the editor of LIGHTWAVEPRO. Questions and comments can be sent to LWPRO, Avid Media Group, 273 N. Mathilda Ave, Sunnyvale, CA 94086-4830. Send e-mail to jgross@netcom.com or avid@cup.portal.com.

Editor's Message

ways to transfer images/objects and scenes from one platform to the next. Now that LightWave will be using a "Content Directory" structure (see last month's Editor's Message), things will be a bit easier as far as loading scenes. But what about getting files from one system to another? Networking systems, null modems and sharing PC-formatted media have all been used. Each has its own set of problems. Perhaps one of the easiest solutions I've seen is Dan McCoy's rework of lha into LhaNT. Lha is a common archive used on the Amiga and the PC. McCoy modified it into LhaNT, which will work on an NT-based Intel machine. Now you can use lha to archive scenes/images/objects on the Amiga side using their normal long filenames. Then you can use LhaNT on the PC side to unarchive everything into the same directory names that were used on the Amiga side. Because it's NT, long filenames are kept and your scenes will load just fine (assuming all the paths are correct).

For those of you that don't know McCoy, he is (among other things) the brave soul that has taken on the responsibility of maintaining the LightWave mailing list on the Internet. This often thankless job is not something I would want to deal with. His LhaNT was included on May's *LWPRO* disk and will be featured on this month's as well. You can also find it on-line at the tomahawk LightWave ftp site. Thanks Dan!

While I'm on the subject of the Internet, the LightWave mailing list and newsgroups have changed once again. See Dan Ablan's new column (next page) for the latest names to subscribe to.

From the Troubleshooting Department

The way LightWave works with sequences has changed from version 3.5. You no longer need to edit the name of an image to get LightWave to find the sequence. The software now detects the base name of a sequence and figures out the extension and number of digits automatically when you click on any image in the sequence. Because of this auto-detection, there are a couple of things that you need to be aware of when creating/naming sequences:

(1) Avoid saving sequences with a numeric digit at the end of the base name. For example, a sequence of the name "Flyby2" would result in files named Flyby2001, Flyby2002, etc. When selected in LightWave, this sequence would read as "Flyby," and LightWave would correctly assume four digits for the sequence. In this case, you wouldn't see your image being used until you rendered frame 2001. Of course, you could use a Frame Offset of 2000 to enable image 001 of the sequence to render on frame 1 of your animation. If you must use a digit in your base name, ending it with a hyphen or other character so "Flyby2" would become "Flyby2-" will make things easier.

(2) Existing sequences named Flyby.001, Flyby.002, etc., will not be recognized correctly by LightWave's auto detection. LightWave will, however, recognize a sequence named Flyby001.iff, Flyby002.iff, etc. If you need to use sequences named with the sequence as the extension (e.g., "Flyby.iff"), you will need to rename all images in the sequence so they correspond to the proper format.

LightWave now allows for the automatic addition of a three-digit extension to saved images by selecting the Name001.xxx or Name0001.xxx options in the Output Filename Format pop-up menu (**Record** panel). LightWave will inform you of how images will be named when you make any changes to the Output Filename Format and/or RGB Image format pop-ups. Remember, if you are running LightWave from Windows (not NT) or an Amiga hooked up to a Novell network, you need to make sure the name of the saved images (including the numbered sequence) will be no longer than eight characters. In this case, a saved name of Flyby with Name001 selected for Output Filename Format would be allowed. But Flyby with Name0001 selected would cause nine characters in the image name, and this image would not be saved. If you are working on an NT machine/network, there is no need to limit the names to eight characters.

By the way, by the time you read this, expect LightWave 4.0 to be shipping anytime now (if it hasn't already).

LWP

In the July Issue:
Look for Part II of Alan Chan's discussion of volumetric tricks and a tutorial by Joe Dox on proper wheel rotation acceleration/deceleration formulas.

lwpro@internet.online

by Dan Ablan

It seems that anywhere you turn these days, the phrase "surfing the net" keeps popping up. Everyone is talking about the information super-highway and asking you for your e-mail address. What is all of this, anyway, some of you might ask? And what does it have to do with LightWave 3D?

Many of you are well aware of the Internet, and the benefits of being on-line. I've received a number of messages and comments from readers via the Internet, and am constantly learning more about this useful service. And I'd like to share some of that knowledge. This new column is dedicated to the Internet and LightWave. Among other things it will help familiarize you with, FTP sites, World Wide Web pages and newsgroups related to LightWave. In coming articles I'll cover highlights of Internet happenings, newsgroup topics and interesting sites, all related to LightWave.

Where Did It Come From?

The Internet is essentially a giant computer network. It is now the fastest growing source for services and information for computer users. This technology has revolutionized the way the world communicates. Information on this network is freely exchanged between the over one million computers connected to it.

The Internet began as a project started by the Department of Defense in 1969, to link the DOD and military research contractors. Universities were also linked to this experimental network. In the beginning, it was called ARPANET. One of the important developments during this time was something called dynamic rerouting. This innovation basically kept the network up and running if there was an enemy attack.

In 1991, then-senator Al Gore sponsored the National Research and Education Network which hooked up all primary schools, universities, government agencies and so on to one big super computer. What this was supposed to do was let a small school have access to the same information as the big universities, and vice versa.

Once APRANET grew to the point of unmanageability, it was broken up into two parts. The first was a military site, MILNET, the second a smaller APRANET.

These two networks were connected by Internet Protocol, or IP, which basically allows all networks to speak to each other. Later, the National Science Foundation started NSFNET, which was highly successful. This network became so successful that APARNET was eventually shut down.

NSFNET is now specifically for research and educational information, and IP networks are being used for other types of information. Within 10 years, one half of the homes in the United States will have a personal computer, and most of those users will have access to the Internet. Every day, more people join the masses and hook into the Internet, often communicating with people around the world to learn about and discuss a variety of topics. And there are topics for just about everything. Even LightWave.

The LightWave Newsgroup

One of the great things about the Internet is the newsgroups. Basically, these are areas of discussion on any number of topics. And, if you haven't figured it out by now, there is a newsgroup dedicated to LightWave. The way it works is quite simple. By entering the correct address via your Internet software, you can link to the newsgroup area. From there, you select the area of interest, such as LightWave, and locate the particular newsgroup. Most software applications will remember your newsgroup choices, so you don't have to search for them the next time you log on.

Once on the newsgroup, you'll see the topics. The topics are posted by anyone who is on line. You, for instance, can sign on to the newsgroup and post a message about the lighting question you had. When you post a message, anyone in the world who has an Internet connection can read it. Imagine a huge electronic bulletin board, like you'd see in a school hallway; anyone walking by can read what you posted. The Internet works the same way. Figure 1 is what my Macintosh interface looks like when I'm in the LightWave newsgroup. I use software available for the Mac or PC, and there are a number of other software packages you can use.
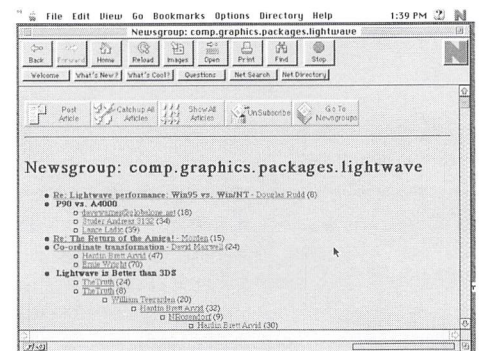
I bet you're wondering what people talk about on

Figure 1

the LightWave newsgroup, and who talks on the newsgroup, for that matter. Well, the topics are very diverse, covering everything from how to create a gear in LightWave's Modeler to rumors about industry manufacturing. I've been monitoring the newsgroup for a couple of months now, specifically for this column. What I've seen posted has been very interesting. One of the most common topics of discussion is LightWave's release for the PC. Most of the time, there are many factual posts from people in the industry, such as Stuart Ferguson, John Gross and Mark Thompson, though rumors are always posted as well. What is truly interesting is the way messages posted on the LightWave newsgroup spark other message posts, regardless of the truth of the original message. It is, in the most basic sense, a discussion group.

Another great aspect of the LightWave newsgroup is what it can do for an animator. Remember, people all over the world can link into this, so when you post a message, anyone can read it. Let's say you've been working on an animation containing a magnifying glass, and it's just not working for you. You've called the people you know, checked the manuals, but can't get it the way you want. So you post a message on the LightWave newsgroup, asking your question. Within a short time, you'll have messages posted in response to your query, and many times, you'll receive private e-mail. Often you will gain the answer you are looking for.

FTP Sites

When you're cruising around the LightWave newsgroup on the Internet, you'll often see people refer to FTP sites. FTP stands for File Transfer Protocol. Basically, this means you can copy files from one system to another. If you have a friend in California and you are in Chicago, and you want to send an image file but don't want to overnight it or spend money on a long-distance modem call, you can upload it to an FTP site. From there your friend can download it. There are some great FTP sites out there for 3D, the Amiga, LightWave and graphics. I'll have listings for you in upcoming articles.

The World Wide Web

You've probably heard this term thrown around quite a bit lately. A Web Page is based on a technology called hypertext. Some people are now labeling it hypermedia, because of its ability to handle graphics and sound files. The World Wide Web connects different bits of information from different systems all around the world. Many companies, including NewTek, now have WWW pages, and you can visit those pages to find out information, look at images, and more. If you

are thinking of attending the SIGGRAPH convention this summer in Los Angeles and need information, for example, checking out the web page on the event is easy, fast and informative.

Regular users can also have web pages. Those pages can consist of anything the author wants, and often contain links to other web sites. It's mindboggling how many places you can visit on the Internet, especially when it comes to 3D animation and graphics. I've recently come across a couple of WWW sites that have tons of images, and some even have sample animations to download.

So Now What?

This column will certainly change with each issue of *LIGHTWAVEPRO*. Next month, plan on seeing a list of interesting World Wide Web sites, explanations of animation topics talked about on the LightWave newsgroup, and a list of good FTP sites to visit. In the following months, we'll continue to list interesting Internet locations and LightWave users' web pages, including mine. If there happens to be a good question on the LightWave newsgroup about a particular animation technique, we'll talk about it in detail, and print a

step-by-step tutorial based on it.

If you are new to the Internet, find out what you can in the coming weeks. Then, and if you've surfed the net before, make sure not to miss this column next month.

For now, if you want to subscribe to the LightWave newsgroup, the name is comp.graphics. apps.lightwave. your internet access software should give you instructions on subscribing to a newsgroup. There is also a mailing list, which I will go into more next month. To receive mail from it, simply send mail to lightwave-request@webcom.com. In the body of the message, include the following: subscribe <optional address>. Optional address can be an address other than the one you are posting from. Don't include the brackets. Remember, this is optional. In a day or so, you will start receiving messages directed toward the list. They can build up fast, so if you don't get on the net very often, it's better to use the newsgroup, as you can pick and choose subjects there.

LWP

Dan Ablan can be reached via the Internet (no kidding) at dma@mcs.com.

Moj-O-Rama II

by Mojo

So many people responded positively to the first "Moj-O-Rama" (October 1994 *LWPRO*) that I knew a sequel was inevitable.

After almost a year, I have finally amassed enough new helpful LightWave hints to put together a few new pages of tips and tricks. Though no one did the first time, I again encourage readers to send in their own little secrets so they can be included in a future installment. The reward? The satisfying knowledge that fellow animators have been aided in their quest to make life a little easier.

Easy Debris

We all know what a pain it is to Boolean an object into a billion bits for an explosion scene. Perhaps a trick using the **Sketch** and **Drill** tools may help quicken the pace. Let's say we want to slice a section of the wing from the popular DC-10 object. With the wing section in a secondary layer, use the **Sketch** tool to draw a curve in a rough shape over the section of wing you want to slice off


Figure 1

(Figure 1). Switch layers as you would to Boolean, but instead use the Template Drill (**Drill**) feature's **Slice** command. Presto! The wing now has a section with a squiggly slice through it. You can easily select and cut these polygons, saving it as a piece of debris (Figure 2). In addition, unlike with **Boolean,** the rest of the object remains, ready for you to chop off the next piece.

Better still, you can easily draw a new curve for each cut, making the edges of each chunk more random. Keep in mind, however, that this technique will not create polygons along the cut edges as Boolean does. If you're going to be close to your cut sections, their absence may be noticeable and Boolean should be used where necessary. As always, consult a physician first.
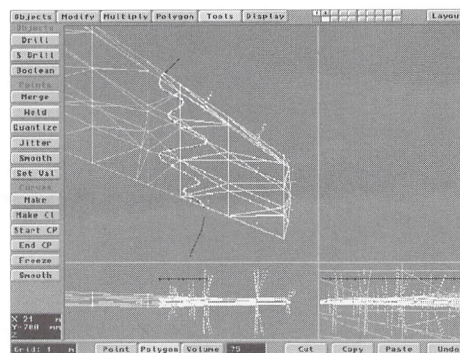

Figure 2

Better Debris

While we're on the subject of debris, any good animator knows that the job doesn't end once the piece is cut. Oh no. If we're talking about an explosion, the debris that is strewn forth from an object needs to look as burnt and damaged as possible. One problem I've seen in many explosions is


Figure 3

too much clean debris. Of course, you could Boolean little holes in the debris itself to make it look really good, but this process would take forever. You don't have forever!

So use a **Clip Map** of **Fractal Noise** on your debris sections. Figure 3 shows this on the sliced nose section of the jet. The clip map makes the object look like metallic swiss cheese—kind of how wreckage would look! As with Drill, the clip map shows no inside faces, so it is best used on smaller pieces that tumble away quickly. For added nastiness, those using the latest version of LightWave may want to consider adding the new **Crumple** procedural in the **Bump Map** channel. If used subtlety, it can really help debris look badly damaged. Try this especially on the edges created with Booleaned pieces.

Flare on a Polygon

Glow behind lens flares *hates* transparent objects when there is a solid object behind both of them. Since the flare doesn't know an object is transparent (only if it's there or not), it won't shine properly through them and often creates black edges along them. On *Babylon 5*, we encounter this problem particularly during explosion scenes, when the glow behind flares used for the nebulas turns spark particles black.

To fix this, we simply rendered a full-screen lens flare against a black background and transparency- mapped it onto a polygon. The polygon was then positioned way in the background, occupying the same space the flare would have. Problem fixed! The polygon can easily be sized to change the flare intensity and dissolved to match the look of any flare.

In addition, if you map the flare image into the polygon surface's **Luminosity** channel, the color channel is open for tweaking. By turning off the texture **Antialiasing** (unnecessary for such a soft image), the lens flare polygon will also render faster than an actual flare. So, even when it's not causing a problem, sticking a flare on a polygon can be advantageous compared to the "real thing."

Wide-Angle Tricks

How often do you adjust the **Zoom Factor** on LightWave's camera? Most people almost never do. However, this feature does mimic real lenses by compressing or elongating space, depending on the length of the lens. By typing in a very low setting, weird distortion effects can be achieved as they are with wide-angle camera lenses. Figure 4 shows what the director's chair object looks like with a 4mm lens. As the camera pans and tilts, the object appears to smear across the screen like a bad acid trip. While these effects aren't very useful for every-day work, they can spice up a logo animation by stretching and bending text in impossible ways. Remember that the zoom can be enveloped, allowing you control when the distortion effects take place. Experiment by making a wireframe preview with any object. Try this feature with particles! At the very least, I guarantee you'll find it interesting.

Faster Wireframes

No matter what you do, even with stand-in objects, wireframes can sometimes take forever. To cut the time in half, create your preview "on twos." Simply make it with a frame step of 2 and play it back at 15 frames per second (fps). It won't be as smooth as 30 fps, but it will be more than good enough to preview your action and save a lot of time!

Speedy Modeler Redraw

Turn off the polygon **Normals** display option in Modeler (d). Selected polygons will still turn yellow (or white) when selected but will redraw almost twice as fast. I find that the majority of the time I simply don't need to know what direction my polygons are facing, and the Align feature easily fixes errors that may occur.

Quicker Tedious Input

You know the story. You've just LightSwarmed 100 lens flares to a spaceship so they all have the same settings. Now you must assign a **Fade With Distance** number to each and every one individually. Or you have to change the diffuse value for 50 surfaces. No matter how you look at it, you're going to spend a lot of time pointing, clicking and entering numeric values. Want an easy way around
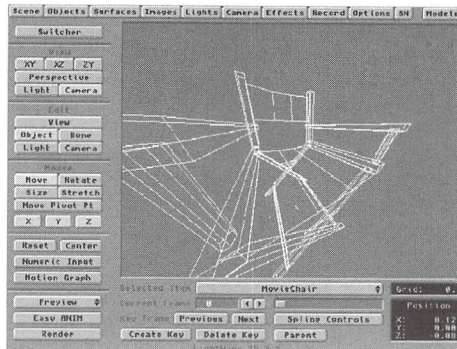

Figure 4

this? Ha! Wouldn't we all! However, there is a way to do this that can save you some time. I find that moving the mouse and clicking in the requester box repeatedly takes the most time. How many of you remember that simultaneously pressing the left ALT and Amiga keys simulates a left mouse button click? By placing the pointer directly on top of the requester and pressing these keys, you'll see the cursor appear, ready for input.

Input your number, hit the down arrow to advance to the next light (or surface or object or whatever), *make sure you don't move the mouse* and click those two keys again. Presto! The cursor appears, ready for input.

Now all you need to do is keep hitting the arrow, (left ALT+ Amiga) and the numeric value. It may seem close to the same amount of work, but trust me, by avoiding the mouse, you can get those values entered much more quickly—no more than a second each, if you're good!

Mirror Object in Layout

Yes, you heard me. Imagine being able to mirror a mapped object! Let's say you're building a left airplane wing and it has 25 surfaces. To make the right side, you can mirror it in Modeler, but then you need to map 50 surfaces. Instead, load the finished, mapped wing into Layout and scale it to -1 on the X axis. This will literally turn your object inside out and create a mirror image, maps and all! Be careful to make your surfaces **Double Sided** (remember, you'll be flipping all the polygons inside out as well) and keep in mind that any maps with text or numbers will be reversed. In addition,

entire objects can be inverted to suit your tastes. If you had an asymmetrical object and decided you liked it better with the features reversed, this would save you the trouble of having to rework it in Modeler.

Motion Removal

Did you know that if you size a light or camera to zero, they cannot move? Well, you can't size a light or camera, but you can parent them to a null object that's been sized to zero. We once had a long, (two-minute) shot at Foundation Imaging that involved the camera flying smoothly through a valley. Everything was fine, but a moment later the camera had to begin shaking and vibrating for the rest of the animation. Making the camera shake for 90 seconds is easy—simply create a two-second-long repeating envelope. But to keep it steady for 30 seconds and then make it shake? Envelopes don't let you repeat a limited range, so it would mean handmaking a shake envelope over 2,700 frames! Not very appealing.

To remedy this, the camera was given a short, repeating shake envelope as first suggested. It was then parented to a null object, which provided the motion path through the canyon. However, for the first 30 seconds of the move, the null was sized to .0001, and scaled up to 1 afterward (you can't size the null all the way to zero or else the camera becomes inverted). While the null had a size of nearly zero, the camera shake was imperceptible. However, as the null sized up, the shake became more obvious until the null reached 1 (when the shake was normal).

Remember that this technique applies to motion only, not rotation. The shake could not have been done with tilts and pans. You may never find yourself in a similar situation, but isn't it neat to know that problems like these have such bizarre solutions?

That's all for now!

LWP

Mojo works as an animator/technical director for Foundation Imaging.

What's on the Disk:

This month's *LIGHTWAVEPRO* disk contains scenes and objects for Alan Chan's and Lloyd Eldred's tutorials on volumetric lighting and exploding waves. Also included is Ernie Wright's DOF Toy, a program that allows you to change values for LightWave's Depth of Field settings and near and far in-focus limits, giving you the numbers you need to plug into LightWave's DOF settings for proper results. The program comes in both PC and Amiga form, and the PC version includes source code. A readme file gives all the directions for this easy-to-use but extremely powerful utility. It should be found on every LightWave animator's system.

LightWave Tech

by Phillip Hice

This new column actually starts at the next paragraph. Because John Gross runs the *LWPRO* editorial department like a dictator, I was ordered to include a brief introduction, so here it is:

Hi. You are probably aware that John Gross, Grant Boucher and John Parenteau (among others) not only write *LWPRO* articles, but are also animators at Amblin Imaging. I also work for Amblin Imaging, except, rather than animate, I handle all of the technical aspects of 3D work. These areas include, but are not limited to, building file servers and disk array sub-systems, network administration and design, and on-site workstation service.Anytime John has a technical question, he comes to me. With my joining the LWPRPO writing staff, you now have the same opportunity. It is my intention, with this column, to evaluate new hardware, inform users of potential bugs and how to avoid them, and to assist with anything that has to do with your system in general.

Well, as you know, LightWave for the PC is just about here, and, as with every type of change, both good and bad come from it. On the good side of this change we have faster rendering times, inverse kinematics, HIIP (host independent image protocol), distributed network rendering, and more colors available for the interface, just to name a few. Now for the bad side. I am sure that many of you, even the ones who can't currently afford a new system, have been looking into pricing on various PC- type systems to run LightWave, such as 486s, Pentiums, MIPS and DEC Alphas.

Although each of these systems runs a different processor, they all have certain things in common. Some of these items—like floppy drives, hard drives, mice and keyboards—are obvious similarities, but there are more important, not so obvious similarities that you need to be aware of. Of these, not-so-obvious similarities, two of them, bus architecture and video display, stand out as areas that require immediate understanding in order to make the best purchase decision.

Bus Architecture

Bus architecture refers to the style of adapter card slots your motherboard contains. In the Amiga, you had your Zorro slot, your video slot, a few Amiga slots, and even several PC-style slots. In the world of PCs, however, you normally have two main bus types on the motherboard. The first type of bus, which happens to be the very first PC-type bus, is the ISA slot. ISA stands for Industry Standard Architecture, and includes both XT and AT bus types. The ISA-type slot is the same PC slot found in your old Amiga. The XT slot refers to the small 8-bit slot not normally found on new systems. The AT slot refers to the larger, 16-bit slot, which is an XT slot with an additional 8-bits of data path included. As stated earlier, the ISA bus type is found on all PC systems, because it was the first design and has, as its name implies, become the

industry standard for all systems. Because of this broad-base support, the ISA bus has the largest selection of peripherals to choose from. For certain items, the ISA bus is just fine, and sometimes recommended. Such is the case with I/O devices. The design of serial and parallel communications cards was done back when there were only eight bits to use, and thus they have only 8-bit memory addresses. If the cards only use an 8-bit memory address, there is no reason to use a 16-Bit card/slot to handle the communications. Also, the ISA slots timing is based on a fraction of the clock speed, usually somewhere around 8 to 12 MHz depending on the overall system speed.

The next step in bus design is the **EISA** type. EISA stands for Enhanced Industry Standard Architecture, and once again, as the name states, this design is an enhancement of the ISA type. Like the ISA bus, the EISA bus utilizes a fractional method for determining clock speed, also somewhere in the 8 to 12 MHz range. Where the enhancement comes in is that the EISA bus breaks the 16-bit ISA barrier with a full 32-bits of data path, effectively doubling the data transfer capabilities. But that's not all. Along with 32-bit data transfers, the EISA bus brings in another feature to further enhance its abilities—Bus Mastering. Bus mastering is a system function that allows multiple cards to utilize the same physical hardware settings without causing a conflict within the system. This is highly desirable in file servers, as it allows for multiple network cards to share settings that normally could only be used by one card. Another performance enhancement over the ISA bus is in improved DMA (direct memory access) transfers, as they are done at 32-bit wide request rather than 16- or 8-bit.

The bus structure that won wide approval after the introduction of the EISA bus was the **VESA-Local Bus** (VLB). VESA stands for Video Equipment Standards Association. This group which determines

the standards computer video devices should be built to meet, such as the pin-out configuration for an SVGA monitor or the standard refresh rates for specific resolutions. The reason the VLB bus came about was that with the introduction of Windows and other graphical software, something was needed to allow for performance gains in the video display area of PC computers. Thus VESA stepped in with the Local Bus standard. Like the EISA bus, the VLB is based on the original ISA bus, with additional connections to allow for 32-bit data transfers. The significant performance increases come from a combination of two things. First, the local part of Local Bus refers to the electronically close proximity to the CPU, which allows for quicker data exchange. Secondly, in order for the VLB to communicate directly with the processor, it runs not at a fraction of the system speed, but at the true system speed at a maximum of 33MHz. This combination of 32 bits and 33MHz allows for data transfer rates that are approximately 3 to 4 times that of EISA. The only drawback from VLB is that it doesn't allow for the Bus Mastering capabilities of the EISA bus, making it not as functional as an EISA machine designed for file server usage.

The fourth, and currently the final, bus structure to discuss is called **PCI**, or PCI Local Bus. PCI stands for Peripheral Component Interconnect, and was developed by Intel, the company we all know is the world's largest manufacturer of PC processors, which helped convince the industry to accept PCI. Currently, the PCI bus design is the bus of choice for PC users, for both obvious, and not so obvious reasons. First the obvious: the PCI bus is basically all of the best things from the VLB and EISA bus types in that it has 32-bit data paths, can communicate at 33MHz clock speeds, has improved DMA performance, and Bus Mastering features. All of these make the PCI bus the best choice for performance in both workstation and file server applications.

Well, now that we've spoken about the obvious reasons for going with the PCI bus, I would like to mention a couple of other primary reasons for its acceptance. First is the fact that the current PCI conventions have provisions not only for a 32-bit data path, but it also calls out specifications for a 64-bit version of the bus, which with no other changes will give double the performance of the original 32-bit version. Although the specs for the 64-bit version exist, I have only seen one system with a 64-bit slot, the Carrera EV5 tested in last month's benchmark

article. Currently manufacturers have devoted manufacturing of both motherboards and peripherals to the 32-bit version. But, that leads us to our second, not so obvious reason for choosing PCI, expansion. The 64-bit slots are merely enhancements to the 32-Bit slot, and will handle 32-bit cards without any problems (as a matter of fact I ran the 32-bit video and SCSI cards for the EV5 in 64-bit slots without fail). Another aspect of expansion comes from cross-platform expansion. This is to say that not only is the PCI bus the primary performance bus for x86 and Pentium systems, but is also the case in RISC machines as well. As an example, my personal Cobra 275 machine from Carrera utilizes my old Matrox PCI video card that I previously had in my 486DX2-66, as well as my SCSI card. This cross platform usage allows for significantly less expensive upgrades, as you can utilize some of your existing hardware, thus keeping down the costs.

Video Display Cards

Now that we have decided that the PCI bus is the best way to go, let's talk about another confusing area of the PC world, video cards. (Although I am specifically discussing PCI cards, the information discussed concerning video cards can apply to all different bus types.) Unlike the Amiga, PCs and similar machines have no inherent video capabilities. All systems require a display card to interpret the digital data and convert it into visual data. Back in the early days of computers, a black screen with a single or 'monochorme' display color was more than sufficient for the tasks of the day. Now that we have Windows, Windows NT, OS/2, and various other graphical type operating systems, we require the ability to display millions of colors at resolutions that even five years ago were merely dreams. It is this color/resolution ratio that needs to be understood in order to select the proper video card for your application.

As you are now aware, the current standard for the PCI bus is 32-Bits wide. This means that the CPU/system can send data to the card 32-bits at a time (VLB and EISA also send 32-Bits at a time, while ISA only sends 16-Bits at a time.) This is where the confusion starts for some people. Someone might ask, "If a true color image is only a 24-bit image, then a 32-bit bus is overkill." WRONG! The 24-bits that are talked about in regards to an actual digital image such as an IFF file is bits-per-pixel of resolution. Which means that for every pixel of on screen data, you have 24-Bits of

information in memory to determine its color. On the other hand, the 32-bits referred to by the PCI bus is the number of information bits that can be sent to the video card, from memory, at any one time. So, if a 32-bit PCI display card is sent 24-bit per pixel color information, the card is capable of receiving approximately 1.5 pixels worth of information at a time (1 pixel =24 bits of data, 1.5 pixels = 36 bits of data).

If you understood the last paragraphs explanation of the difference in image data, and video display data then you should be ready to tackle the next question, "What is a 128Bit PCI video card?" Well as we have touched on several times, only 32-bit PCI cards are available, so what is a 64-bit card then? What this means is that the PCI card communicates with the CPU/system in 32-Bit blocks. Then once the video card has those blocks, it puts them together and processes them in 64-Bit wide chunks. The way that this is done is by VRAM, which is short for Video-RAM, is a type of fast memory used in video display cards. What the card does is load the data from the system through the 32-bit bus as fast as it can into VRAM, then the video processor, the chip on the video card that converts the pixel information into displayable color information, retrieves the data from VRAM and processes it out to the monitor in 64-Bit blocks. What this does is allow for faster screen updates as the card can process more data to the screen in the same amount of time. Currently, the graphics accelerator cards are available in 64-Bit, and 128-bit versions. What these allow for is two to four times the display processing power of a standard 32-bit display processor connecting to the system through a PCI bus.

Understandably there are many other decisions to be made during the purchase of a system. Most of these decisions are easy to make with little or no help. I'll be here to help you with the tough ones.

Stay tuned for upcoming articles on how to get your old Amiga information to your new system, how to set up your network for distributed rendering, and the latest and greatest systems available for LightWave.

LWP

Phillip Hice, Amblin Imaging's technical director, has been involved with this aspect of computers for longer than he cares to remember.

Morph Surfaces: Off

- You will note that the OscBeam object has now turned into a cone. This is good. Set a keyframe and test render a frame (Figure 6).

What did we do? We texture-mapped the cookie luminosity map to create streaks down the length of the cylinder, then we morphed that cylinder into a cone so that the streaks would look as if they were originating from the point of the cone. Over time, the cookie texture slowly moves through the surface, causing the oscillating effect. Render a test anim and see for yourself.
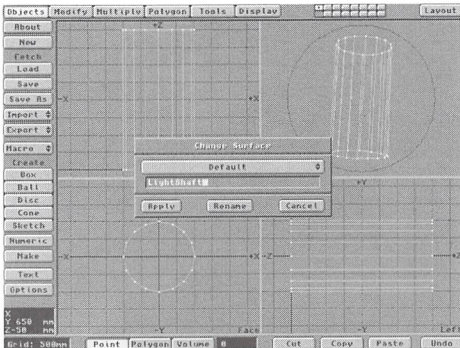
Figure 7: Creating a shaft of light in Modeler

Figure 8 :The OscBeam object with tripled polygons

This sort of oscillating light source is best used in conjunction with other elements, such as tree objects or foliage. For instance, if you set up a tree scene, parent a light to the tip of the lightcone and move it, you'll have what looks like a lightbeam rifling through the trees. For further realism you could set a dissolve envelope on the lightbeam to cause it to be occasionally occulted by tree fauna.

As with anything of complexity, there are some limitations to this technique, but they are not impossible to work around. For instance, it can be very effective when this lightbeam is pointed toward the camera, but because of its falloff settings, the apparent intensity diminishes dramatically. To increase the intensity, you can either increase the Luminosity **Texture Value** of OscLightSurface (the maximum limit is 400% in 4.0), or, if that's still not enough, parent a second light cone to your first. The second option, of course, adds to your overall render times.

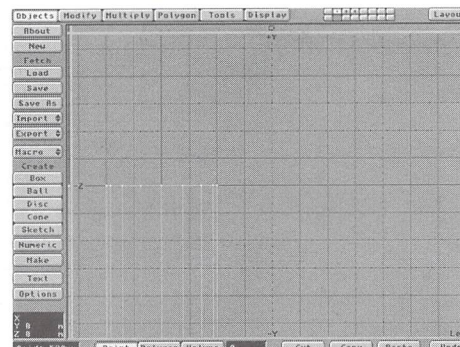Figure 9a: The Shaft object before a 90-degree rotation

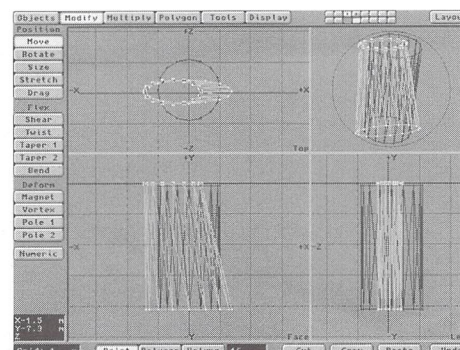Figure 9b: The Shaft object after a 90-degree rotation

Figure 10: The morphed version of the shaft object

Getting the Shaft

Armed with this concept it is possible to extend it a step further, and construct another variation on volumetric lighting: the shaft of light that pierces a dark, dusty environment, say in a mining tunnel, or the abandoned house with the derelict roof letting in slivers of light. Instead of a cone, simply extrude the shape needed and use that instead.

- In Figure 7, I went back into Modeler and reloaded the OscBeam/lwo object. After changing the surface name to "LightShaft," I rotated the entire object by hitting (r) while the cursor was located at the crosshairs in the Left view (Figures 9a and 9b). This orients the lightbeam in the Y axis. Finally, I went to the Polygon menu and tripled the polygons (Figure 8). This object was then saved as "Shaft.lwo."

- Next, the points of the top and the bottom of the cylinder, respectively, were selected, and using the Stretch tool (**Modify** menu), I pulled and stretched the cylinder into an oval shape (Figure 10). This shape was then saved and loaded in Layout as "ShaftMorph.lwo."
- Again, I made the ShaftMorph.lwo object 100% dissolved, and made it the **Metamorph Target** for the Shaft object. The Shaft object should be set to 100% morphed without **Morph Surfaces** on.
- From the Surfaces menu, I applied the following surface settings to the "Lightshaft" surface:

Surface Color	255,245,180
Luminosity	0%
Luminosity Texture Type:	Grid
Texture Falloff:	0,0,25.0
Texture Value:	60%
Line Thickness:	1.0
Additive	On
Diffuse Level	0%
Transparency Texture Type:	Planar Image Map
Texture Image:	Cookie.iff
Texture Axis:	Y Axis
Automatic Sizing	
Texture Center:	0,0,0
Texture Velocity:	0,0,0
Edge Transparency	Transparent
Edge Threshold	1.0
Smoothing	On
DoubleSided	On

- Finally, I keyframed with the thin edge facing the camera and test rendered. For a more realistic look you may want to remove the cookie pattern, or add a slight Fractal Noise as a Surface Color texture. Add World Coordinates and a very slight Texture Velocity to simulate particles in the air, such as you would see swimming around in a bright shaft of light.

Combined with proper lighting, this technique can be quite convincing.

Same BatChannel

Next month we'll get into stuff that's a bit heavier, such as creating volumetric shadows for logos, spaceships and the like. Between reading this issue and next month's you might even experiment a bit, try your hand at different surfacing options and see what exotic lighting schemes you can create. Render on!

LWP

Alan Chan is volumetrically light in the head. Be nice and send him e-mail at alan.chan@oubbs.tele-com.uoknor.edu, or call Vision Digital at (405) 447-7075.

Jet Nose

The nose of a jet has been sliced and the use of a clip map gives it that metallic look. See "Moj-O-Rama II," page 14.

Copyright 1995 Mojo

Leaves and Light

Use LightWave to fake the effects of volumetric lighting. See "See the Light," page 4.

Copyright 1995 Alan Chan