LIGHTWAVEPRO

Inside:

More Dirty Tricks from the 'Imperfect World'

From the Net: Full Steam Ahead

MORPHING WITH METANURBS

Desireable Dents

Arnie Boedecker achieves the "lived-in" look through use of a bump map to create scratches and dents at the base of the door and along the baseboards.
Copyright©1996 Arnie Boedecker

Wash Me!

More of Boedecker's "dirtying-up" tactics: simulated water spots and hand prints on a window pane. Too often, drinking glasses or glass coffee tables are rendered perfectly clean.
Copyright©1996 Arnie Boedecker

EDITOR'S MESSAGE

by John Gross

So what was new and exciting at SIGGRAPH? LightWave-wise, I feel that the D1 version of the Perception card from Digital Processing Systems, Inc., was pretty cool. There also was a nice showing of plug-ins and programs for LightWave. Among the offerings were Dynamic Realities' Particle Storm, ATR's WaveNet Pro, Blevins Enterprise's Vertilectric and Vertitilities, Positron's MeshPaint and Unlimited Potential's Wave Filter.

All of these plug-ins/programs are top-notch and should definitely find their way into your LightWave arsenal. There are many other plug-ins in the works for LightWave, so you should expect to see some pretty amazing things in the near future.

NewTek also showed their Morph Gizmo which is an amazing little tool that allows for interactive morphing of objects. It was written by Dan Milling and Fori Owurowa of NewTek and hopefully will be shipping as a separate program by the end of the year.

Speaking of shipping by the end of the year, the LightWave team is hard at work on both Sun and Power PC versions of the program. Expect to see some version of 5.0 ship for the Power Mac by the end of the year. (Look for more news on the Mac version at this year's VTU Expo, Nov. 5-8 at the Universal City Hilton in North Hollywood, Calif.)

And, if things go according to NewTek's plan, you can expect the next version of LightWave (6.0?) to be available for workstations from Sun Microsystems, Inc., another new ally.

The addition of these two platforms will basically round out the cross-platform capability of LightWave. LightWave is the only program in existence that I know of that runs on such a wide range of platforms .

I believe a nice little side effect will be the common use of LightWave in advertisement benchmarks of system comparisons. This will certainly help to make LightWave a "universal" program. Another nice side effect is that all

TABLE OF CONTENTS Vol. 4, No. 8

LIGHTWAVEPRO
a Miller Freeman, Inc. journal

"We provide the most valuable information to people who use technology to create messages with impact."

Are you interested in writing for *LIGHTWAVEPRO* or submitting images? If so, contact us at our office or electronically.

About this month's cover: The image is from an upcoming episode of the Fox Network television show, Sliders*. It is a mechanical spider designed to electrocute captives caught on its magnetic web. It was modeled by John Daniel and animated by David Lombardi, both of Digital Muse, an animation studio in Santa Monica, Calif.*

On the Cutting Edge
How to Simulate a Torch Cutting Through Metal

by Adrian Onsen

Have you ever seen something cool in real life, and wondered, "Is there any way I can recreate this effect in LightWave 3D?" Well, I wonder that all the time. One of the things I recently challenged myself to do was using LightWave to recreate the effect of cutting through a chunk of metal using an acetylene torch. You can see this effect in many movies (especially sci-fi ones) and at construction sites. I couldn't think of any movies off hand that had this effect in them, so I decided to find someone with a torch and shoot some video footage of them cutting through a piece of metal. I didn't want to do it from my imagination, for the effect would lose it's realism. After quite a bit of work in both Modeler and Layout, I generated some good looking effects.

This tutorial will give you a step-by-step instructions on how to achieve this cool effect! This can be used in many situations. Right now, for example, I'm doing some logo work, where the letters are cut out of a chunk of metal, after which they fall to the ground.

All you need now is any version of LightWave (yes, even version 3.1 will work), and a tasty beverage of your choice, because this is not a short tutorial.

How It All Works

First, we will model the entire piece of metal with the cut created by the torch. Secondly, we will create a clip map of the path that the torch will take. The clip map will then be mapped onto a plane right in front of the metal piece. This will give the illusion of the metal being eaten away over time. Then a combination of sparks and lens flares will be used to fake the pyrotechnics of it all. Not too difficult, is it? So let's start!

The Modeling

As for any project, we need to make some object which we will later animate in Layout. So sit back in your favorite chair, relax and take a sip of the beverage (you did bring one, didn't you?).

• To start off, we need to make a plane. Make it 1 m by 1 m with the center at 0, 0, 0. In other words, X low = -500 mm X high = 500 mm Y low = -500 mm Y high 500 mm with all segments = 1.

A plane should appear in the XY view (Face View). Now you need to know what you want to cut out of this. I decided to do a nice smooth horse-shoe shape (Figure 1). Keep in mind that this shape can be as simple or as complex as you want. I chose a simple shape for the sake of simplicity.

• First go to Layer 2 (2), Here you need to set down points (Click Polygon, then Points) in the shape of the curve you want to cut, and then make it a curve (CTRL-p).

When setting points, make sure that they are at somewhat equal intervals from each other, otherwise you run into problems later (Read the Beginner's Sidebar for further details on spacing out points on the curve).

Now that you have a curve in the shape of your cut, you have to decide how

Figure 1: Shape of the curve that needs to be cut.

wide you want the cut to be. I made my cut 25 mm in width.

• Switch to Layer 3 (3) and set down two points. In my case, these are 25 mm apart. Then click Polygon/Make (or p) to create a polygon between the two points. This will create a line between the points.

• Put Layer 2 (Alt 2) in the background, move this polygon to the beginning of the path (curve) you created earlier, and rotate the single-sided polygon until it is perpendicular to the start of the curve (Figure 2). Another thing to check before going any further is this single sided poly you just created is at the beginning of the curve. You can determine the start and end of the curve by selecting the curve as polygon. The curve will turn yellow, and one end will have a box or "diamond" around it. That's the beginning of the curve, and that's where the single-sided polygon should be. If it's not, just flip (f) the curve polygon and the start will jump to the other side. Otherwise, you'll get interesting side effects in the next steps.

Now we will use another one of LightWave's under-utilized features: Rail Extrude (in the Multiply menu).

• Make sure that the curve object is in the background layer and the single sided polygon is in the foreground, and after clicking the Rail Ext button (CTRL r) select Uniform Lengths, and make sure the Oriented box is checked. Now depending on how much detail you want, you can change the default value in the text box. I set mine to 50. Click OK to extrude the polygon

• You should now have a ribbon shape, bent around the curve. If it looks too wide, click Undo, go back to where you created the single-sided poly, bring the points closer together, then Rail Extrude again.

• We now want to cut this shape out of original plane. Copy the plane from Layer 1 to Layer 4. Put Layer 3 in the background (Alt 3) and Layer 4 in the foreground (4) and select Drill (Tools menu).

• Select Tunnel as the operation. Because the front view is the XY view, I selected Z as the axis. This will give you the plane with the curve cut into it.

• Now Extrude (Multiply menu) this object to the thickness of your metal piece. I made mine 30 mm thick.

• To give the cut a somewhat jagged look, select all the points that make up the curve (an easy way to do this is to click on the Volume button at the bottom of the screen, then draw a box around the points. Hit W, and select + for Points in Volume). Now select Jitter (Tools menu), and use the default Gaussian Jitter with a radius of X = 10 mm Y = 10 mm Z = 0 mm.

• This operation will create non-planar polygons so hit W, and select + for Non Planar. If you did everything right up to now, the only non-planar polys that you should have, should be along the inside of the cut (Figure 3). Triple them (Polygon menu).

• Go back to Layer 1 (1) and offset the plane by -5 mm in the Z direction.

Now we need to surface our objects.

• Give the object in Layer 1 the surface name "Front". In Layer 4, give the front face polygon the same surface name. Name the polygons inside the cut "Cut" and the edges of the metal piece, "Edges".

Since we're done with the major part of the modeling, it's time to save our work. You need to save Layer 1 (the offset plane), Layer 4 (the metal chunk with the cut in it) and Layer 2 (the path for the Rail Extrude). Save Layer 1 as FrontPl.lwo, Layer 2 as Cutcurve.lwo and Layer 4 as Cutmetal.lwo.

Creating the Clip Map

I've thought long and hard about how to explain how this animated clip map can be created, because people on all platforms will be reading this article. So I had to come up with a way which would only require LightWave, and no other packages. Here's how it works:

We're going to convert the curve we used as a Rail Extrude to a motion file for Layout.

• Back in Modeler, load up the Cutcurve.lwo object (the curve along which the Rail Extrude was performed). You need to convert this to a motion file in order for Layout to read it. You either need the Motions.p plug-in or, on the Amiga, the ARexx script PathToMotion. If you have LightWave 4.0 and up, and have the Motions.p plug-in for Modeler, select the PathToMotion from the Custom (Objects) menu, and enter a name for the motion. (See the sidebar for more information on the plug-in.)

• You also have to create a small disk which will be used in Layout as a mask. Give the disk a radius of whatever the width of your cut was. I made my disk with a 25 mm radius in the Z axis, with 16 sides and with Top and Bottom both set to 0. Change the surface name from Default to Disk, and save this object as Disk.lwo.

Now that we have a motion file we can use in Layout, let's go and use it.

• Open up Layout, and load up Disk.lwo, FrontPl.lwo and CutMetal.lwo

Figure 2: Align the single-sided polygon perpendicular to the curve.

objects. Now select the Disk.lwo object, go into the Motion Graph menu, and click Load Motion, and load the motion file saved by the motion.p plug-in back in Modeler.

• Upon exiting the Motion graph menu, you should see the points that make the curve. But wait a minute! There's a problem. I don't see a smooth curve. I see points joined by straight lines. Well, the plug-in saves the motion file with the Linear Spline Control option enabled. You have to go to each frame and disable the Linear feature from the Spline Control menu. You should then have a smooth curve. The plug-in also saves the size as 0 at each key frame. In Layout, you have to go to each key frame, hit Reset while Size is selected, and then Create Key. This will set the size back to 1 for each key.

If you made a simple shape like I did, with only 6 points, you'll end up with 25 frames worth of animation. We need more frames to make a smoother/longer animation. In the Scene Panel, select Scale All Keys, and scale them by a factor of 10. This will give you 250 frames or about 8 seconds worth of animation. If that's not enough, scale the keys by a larger factor (15 or maybe 20).

If you looked closely, you probably have already noticed that the motion path doesn't exactly fall on the cut created in Modeler. This is because Modeler and Layout handle their curves a little differently. You need to go in and add key frames where ever you feel necessary, in order to correct for this. With a little patience, this can result in quite close approximations. The reason we made the disk twice as wide as the cut, was to give us some leeway in how close the motion came to the actual cut.

• Save this motion because we need to use it later. Save it as Curve1.lmot in the motions directory.

• Once you're satisfied with the motion path, we need to start rendering the clip map. But first, in the Surface panel, select the Disk surface, and give it the color white (255, 255, 255), and a Luminosity and Diffuse level of 100%. Then in the Effects panel, have the background color set to black (0, 0, 0).

The rendering of the clip map will be a two-step process. The first step is to render the motion of the disk, and secondly to extract out of the full frame only the portion we need to map onto our object.

• Bring your camera close to the Cutmetal.lwo object, but not too close. You want to have a little bit of space between the top of the object and the top of the preview screen. Now create a Key at frame 0, to save this camera position.

The Imperfect World

Shut Your Windows, Close Your Doors

by Arnie Boedecker

Even shutting your windows and closing your doors won't hide our world's everyday imperfections. In fact, windows and doors are just another place to find examples of these imperfections. Oddly enough, doors and windows are just the sort of things that are generally modeled and rendered ever-so-perfectly within the 3D environments we create. But look around you ... doors and windows are rarely perfect. Toes and heels tend to mark up the bottom of doors, while weather and children's fingertips ensure that windows are rarely clean. Let's take a look at how to create nice models of windows and doors, and then how to surface them to look as though they came from our imperfect world.

Slam!

Let's start by modeling a basic door. Modeling a door may sound pretty basic (and it is), but it's detailing the door that makes it stand out. Adding a cleanly modeled doorknob, and hinges go a long way. (I've seen way too many "Ball and Disc" doorknobs).

• The basic All-American door is approximately 7 feet tall, 3 feet wide, and 2 inches thick, so as my mind drifts into the metric system, create a box with these settings:

Axis	Low	High	Segments
X	-457.2 mm	457.2 mm	1
Y	0 mm	2.136 m	1
Z	0 mm	50.8 mm	1

•Assign a surface name (q) of "Door".

Easy enough so far? OK, now for the doorknob. To make the doorknob, I actually just made a disc, then beveled it a number of times to create the doorknob's shape.

For those of you familiar with creating complex shapes with bevels, go ahead and create your own doorknob. For those of you who have only used the Bevel Tool to detail your favorite logo, follow along here, and get a feel for how beveling an object like a

Figure 1

doorknob is a bit more interactive than plotting points and lathing one.

• First let's create the starting disc. Go to a new layer and make a disc on the Z-axis with 32 sides, Top and Bottom values both set at 0, a center of (355 mm, 892 mm, 50.8 mm), and radii of (56 mm, 56 mm, 0 mm).

• Next, go to the layer with the door in it, and put our new disc in the background. Use the Drill Tool (Tools Menu) in Stencil Mode on the Z-axis to stencil in the disc with a new surface name, DoorKnob. With polygon as the selection type, hit the W-key to bring up Polygon Statistics. Select only our new disc stencils by choosing the DoorKnob surface and hitting the + button for that surface.

You should now have two polygons selected. From here we start our series of bevels. Do not ever deselect the polygons, you should always have two

Figure 2

selected. If you are using LW5.0 ignore the random bevel (+/-) settings, leaving them at zero.

To make a bevel, hit the B-key and enter the values listed below.

Bevel	Inset (mm)	Shift (mm)
1	0	2
2	4	2
3	0	3
4	2	2
5	28	0
6	4.5	3
7	0	12
8	3	1.5
9	-3	1.5
10	0	15
11	-7	0
12	-2.5	1.5
Bevel	Inset (mm)	Shift (mm)
13	-9	26
14	-3	17
15	-2	18
16	2	2
17	2	-6
18	3	7
19	26	-7
20	2	4
21	3	0
22	0	-10

• When you are finished with the last bevel, deselect all polygons. If done correctly, your doorknob should look similar to the one in Figure 1. If you have LightWave 5.0, go ahead and give the doorknob a basic surface within the Modeler Surface Panel (q).

Next we need to create the hinges for the door.

• In a new layer create a disc with 16 sides on the Y-axis, with a Bottom value of 300 mm and a Top value of 450 mm, and Center of (-458 mm, 400 mm, -8 mm) and Radii of (8 mm, 0 mm, 8 mm). Assign a surface name of Hinge.

Now we need to make this solid cylinder look like a segmented interlocking hinge. To do this, go to a new layer and make a box with these settings:

Axis	Low	High	Segments
X	-480 mm	-440 mm	1
Y	445 mm	447 mm	1
Z	0	0	1

• Next, use the Clone Tool (Multiply Menu) to create four clones of this box with a Y-Offset of -35.5 mm.

• Next, put the Hinge cylinder in the foreground, and the boxes we just created in the background, and use the Drill:Stencil Tool again to stencil through the boxes on the Z-axis, assigning a surface name of Hinge-Black.

When rendered, the Hinge-Black surface will be colored black, to appear as spaces in between the interlocking hinges of the door.

This next step is only really necessary if you plan on having your door open. We are going to create the hinge plates which connect the hinge to the door. I would recommend building these as well, so that you can have the door open if you wish. After all, in our imperfect environment, how often do we actually close doors? Leaving closet doors open a crack, and other doors wide open can give your 3D home a lived-in look.

• To create the hinge plate, go to a new layer and make a disc with 16 sides on the X-axis, Top and Bottom values of -456 mm, a Center of (-456 mm, 433 mm, 30 mm) and Radii of (0, 15mm, 15mm). Then hit the (k) key to kill the polygon. You should just have points left.

• Next in the Left View, select the 11 points as shown in Figure 2, and hit the (z) key to delete them. Then, still in the Left View, select only the top left-most point and hit the Copy button. Then use the Move Tool (Modify Menu) to move this point -32 mm on the Z-axis. Then hit the Paste button to put the original point down again.

• Deselect all points, and use the Mirror Tool (Multiply Menu) on the Y-axis with a value of 375 mm. This completes the point outline of the hinge plate. Select these points in a clockwise manner, and hit the p key to create the polygon. Use the Extrude Tool (Multiply Menu) on the X-axis with a value of -6mm to give the plate depth. Assign the surface name, Hinge, to this as well.

Depending on how close to the door your camera will be in your renders, you may want to add screws to detail your hinge plate.

• Next, Cut and Paste this plate into the Hinge layer. To make the other needed hinges, use the Clone Tool to create two clones with a Y-Offset of 725 mm. Finally, Cut and Paste the Hinges into the Door layer, and save the object as Door.lwo (Figure 3).

You may also want a set of hinge plates to attach to the inside of the doorframe, once a room is created. Also, if your door will be open, you may need to create a door latch, for the door, and a latch hold, for the doorframe.

Out the Window

Next, we'll model a basic window. A good window is more than just a frame with a pane of glass. Most house windows have two panes of glass, each having their own frame, as well as a screen, all boxed by one large frame.

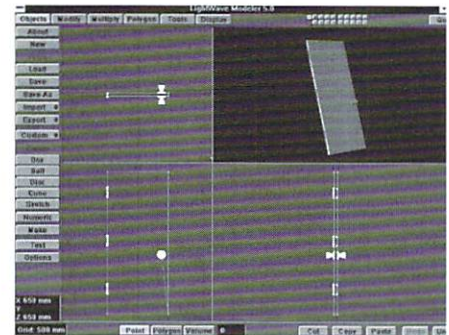• Let's start by making a box with these settings:

Figure 3

Figure 4

Axis	Low	High	Segments
X	-600 mm	600 mm	1
Y	800 mm	2 m	1
Z	0 mm	200 mm	1

• Hit the (f) key to flip all the polygons inward, and assign a surface name of WindowBox. Grab the backside polygon in the Left View and hit (Z) to delete it. Then grab the very front polygon and assign it the surface name, WindowFrame.

We will first create the main frame, boxing the entire window. As with the doorknob, this will be done with a series of bevels, to create a smooth frame. Again, if you are familiar with using the Bevel Tool, feel free to design a fancy frame. If not, follow along while we make a standard house frame.

• Hit the (b) key to create the following bevels, and remember not to deselect the polygon, you should always have only one polygon selected.

Bevel	Inset (mm)	Shift (mm)
1	0	30
2	-5	10
3	-10	5
4	-100	20
5	-5	0
6	-10	-5
7	-5	-10
8	0	-50

• Once you finish the bevels, hit the (Z) key to delete the polygon you used to create bevels. Your window frame should look like the one in Figure 4.

Figure 6

Figure 5

• Hit the (w) key to bring up Polygon Statistics again and select all of the WindowFrame polygons. Then deselect all of the polygons that run horizontally. Once only the vertical WindowFrame polygons are left selected, resurface them with the name WindowFrame-V. Then selecting only the remaining horizontal WindowFrame polygons, rename their surface, WindowFrame-H.

Next we need to make the two smaller windows.

• We'll start modeling the lower window by going to a new layer, and creating a box with these settings:

Axis	Low	High	Segments
X	-600 mm	600 mm	1
Y	800 mm	1.425 m	1
Z	0 mm	0 mm	1

Use the Bevel Tool on this box with an Inset of 50 mm, and a Shift of 0 mm. This creates the template for our window frame and its pane of glass. Select the center polygon, which will be the pane of glass, and Cut and Paste it into a new layer.

•Go back to the lower frame layer, select the horizontal polygons and assign the surface name WindowFrame-H. Do the same with vertical polygons, assigning WindowFrame-V.

• Then, after deselecting all polygons, use the Extrude Tool on the Z-axis, with a value of 67.5 mm to give the lower frame depth. Then go back to the layer where you Pasted the glass. Move the glass polygon 25 mm on the Z-axis, and then Extrude it 14 mm on the Z-axis. Assign the surface name, Glass, and Cut and Paste it back into the lower frame. If you want your window to open, save this lower window as a separate object, so you can shift its position in Layout.

• To create the upper window, Copy the lower window and then use the Move Tool to shift it up 576 mm on the Y-axis and 80 mm on the Z-axis. This now becomes the upper window. You can then Paste down the lower window back into place, to see how they line up (Figure 5).

•Cut and Paste the upper window into the Main window frame (as well as the lower window frame, if you don't plan on opening the window), and save the entire object as Window.lwo.

The last thing you need to do is model the

screen. This is not absolutely necessary, as not all windows have screens. Even those that do, the screen size will vary. Some screens fit the entire window, while others fit only half. For simplicity, we'll model a full window screen, but feel free to modify or remove the screen as needed.

• We'll make the screen in the same manner as we did the lower window, by first creating a box:

Axis	Low	High	Segments
X	-600 mm	600 mm	1
Y	800 mm	2 m	1
Z	180 mm	180 mm	1

• Bevel this polygon with an Inset of 30 mm, and a Shift of 0 mm. Select the center polygon (the screen), and Cut and Paste it into a new layer.

• Go back to the screen frame layer, and Extrude it 20 mm on the Z-axis, and assign the surface name, ScreenFrame. Then go to the layer in which you Pasted the screen polygon, and move the polygon 10 mm on the Z-axis and give it the surface name, Screen (turn on Double Sided if using LW5.0). Cut and Paste the screen polygon back into its frame, and save this as a separate object if you think you may want to remove it.

Finally, you may want to create a basic lock for the window, depending on the amount of detail you need. The completed window frame, with both windows, a screen, and a lock will look similar to Figure 6.

On the Surface

Creating the right surfaces, more often than the detail of the models themselves, is what creates the illusion of realism within our 3D world. The basic kinds of surfaces we need to be concerned with, for our door and window, include: Wood (door, window frames); Polished Metal (doorknob, window lock); Dull Metal (hinges, hinge plates, screen frame); Glass; and Screen Mesh. Let's take a look at each of these more closely.

Against the Grain

How to surface the wood on your door and window depends on how you want your wood to be "prepared". What I mean is, doors and window frames made of wood are usually one of two types: Natural wood that has been stained and finished, or wood that has been painted over.

Stained wood is pretty straight forward. After choosing the color you want your wood to be (actually the color of the stain), pick out your favorite wood image. Using a wood image as a diffusion map keeps the color for the stain, but brings out the texture of the wood. For those using LightWave 5.0, a more subtle look can be achieved by using an opacity setting of less than 100% for the wood diffusion

Figure 7: The Bump map used to create dents and scratches.

Figure 8: This image is used as a transparency map to 'dirty-up' glass.

map. Using the same image as a specularity map may add a nice look, depending on whether or not you want your wood to have a glossy or dull finish. The glossiness of varnish will tend to vary over the wood grain, while wood that is only stained will not have much, if any, specularity. Adding a Bump Map, again using the same image, with very low settings (around 5%) can add a more textured look.

Painted wood is done a bit differently. You don't want to use a diffusion map for painted wood. If you have white painted wood frames, for example, applying a diffusion map would make your white painted surface appear as a series of gray wood grains. Using a Specularity Map for painted wood is much better. The difference in specularity between wood grains creates the illusion of painted wood. Again, a slight Bump Map may add to your look.

Remember, when surfacing the wood of the door and window frames, to use an image of vertical wood grain to surface the vertical beams, and an image of horizontal wood grain for horizontal beam surfaces. To save RAM, it is also a good idea to reduce 24bit images to grayscale, since Diffusion,

Editor's Message

types of computer owners will be able to use LightWave. Maybe it will help with the 'my system is better than yours' mentality (but probably not)!

It seems NewTek is doing pretty well lately. They have some very good strategic relationships going right now (Apple, Intergraph and Sun) and they've been working hard at improving customer support. Obviously, I'm happy about this as it helps assure continued success for our favorite 3D program.

Until next time,

John Gross
Editor

Specularity, and Bump Maps only need the grayscale information. Also, be sure to use Cubic Image Map surfacing, rather than choosing a particular axis. This will ensure that the image gets mapped onto the faces and depth of the frame, rather than just on the front face, and streaking along its depth. Or, if you prefer, you can use LightWave's Wood Procedural Texture. Remember to change the axis to change the direction of the grain along the X, Y, or Z as needed (the Texture Axis determines which axis the "rings" go around). Be sure to use identical settings for each map, so that Diffusion, Specularity, and Bumps of the grain match up.

Finally, what really makes the wood look "real" is by adding some dents and scratches. Old painted wood gets chipped and painted over again, and stain gets worn away, to reveal the natural wood grain underneath. Doors tend to accumulate these flaws at the base of the door, where our feet are most unkind to its finish. You can achieve these "dents and scratches" by simply drawing some random lines and splotches in your paint program of choice, and using it as an additional Bump Map for the wood. Figure 7 is an image I created in Photoshop, and use for dents and scratches time and time again, by stretching different amounts, to achieve different looks.

Two Faces of Metal.

Surfacing metal depends on the metal itself. I try to stay away from the metal surfaces included with LightWave. While they may get the job done for logos, reflecting the Fractal Reflections image 100% is a bit much, and destroys the "real world" look we are seeking. Polished metal, such as the typical "brass" doorknob we modeled, tends to be glossy and reflective. If your door and window are part of a total environment (in a room of a house, for example) then have the metal reflect its surroundings. Turn on Trace Reflections in the Camera

Panel, and put your reflection amount at around 40-60% (Not 100%!). Many people make the mistake of making all of their metal surfaces reflective. Look at your typical door hinge, can you see your reflection? Unless you recently installed a brand new door, probably not. This metal tends to be dull, and diffuse. At most, you might want to reflect that Fractal Reflection image here, and if so, only about 10-20%, to add some random color to the metal.

Other Surfaces

As mentioned earlier, glass is rarely clean, so be sure to dirty it up in your 3D world. Figure 8 shows an image I used to dirty up a window, again created in Photoshop. Notice it is covered with water spots from the weather, as well as hand and finger prints from unseen CG toddlers. Using this map as a layer of transparency on the glass surface helps make the room appear "lived-in".

When surfacing the screen, remember that people don't usually focus on the screen itself. They are looking through the screen to what lies beyond. Making the screen solid and in focus tends to be distracting and unrealistic. By making the screen a bit transparent, or blurring a grid image as a transparency map, is much more effective than a perfectly rendered, in focus screen.

Want to go that extra mile? How about chipping some paint inside the window box? Or creating some random points to add as the ever popular "dead bugs" in the window cell? Look at your surroundings, there are plenty of ideas around you. A door open just a crack, and a dirty window may be just what your CG world needed.

LWP

Arnie Boedecker can be found insanely examining doorknobs and window frames in the Western Suburbs of Chicago. He can be reached via the Internet at awb@ais.net or by calling (708) 213-6384.

MetaNURBS
A Handy Tool for Morphing

by Patrik Beck

I am predicting that LightWave 3D is about to become one of the most popular character animation software choices in the world. The reason for this is the powerful organic modeling features available with the new MetaNURBS function. MetaNURBS lets you create smooth flowing objects out of rough blocky approximations. More than that, it allows you to interactively mold and shape your model with a near real-time preview that looks embarrassingly close to the final object. Finally, all of us back-room freelancers can break out of the industrial video flowchart jungle and start working on that Saturday morning 3D cartoon show we've been planing to produce.

But wait! There is more to character animation then just building good-looking models; they also have to be animated! Not just animated, but have the ability to show expression, be squashed, stretched, and lip-synced. Unless you are a complete bones master, the most practical way to accomplish this is to use object morphing between various expressions and mouth shapes.

While it seems too good to be true, MetaNURBS makes it possible to create a crude object, alter and smooth it, then create a second object that is a proper morph target for the first object. By using multiple morph targets, objects can smoothly transform through several changes of expression. As usual with computer stuff, this is more complicated to explain than to actually do, but hang on. We will start by building a simple cartoon hand, make some morph targets, then apply what we have learned to more complicated objects.

Three Fingers and a Thumb

The first thing to do is to fire up LightWave Modeler. Go to the Options button in the objects menu. In the Polygons section, make sure the Quadrangles button is highlighted. This will ensure that all the polygons generated will be four-sided. This is important because MetaNURBS requires polygons to be four-sided in order to work properly. Later, when we freeze our MetaNURBS objects for rendering, we want to change this to triangles to reduce the possibilities of rendering errors that may happen if any of the quadrangles are non-planer.

Legend has it that most cartoon characters have three fingers because Walt Disney deemed that three fingers were plenty for a small mouse named Mickey. We can thank him for making our job 20 percent easier.

Let's start by modifying a box object.
•Create a box 1 meter square and 0.5 meters deep, with 3 segments on the X axis and 2 segments on the Y axis.
•Slightly round the edges by dragging the upper two corners down and the lower two corners up (Figure 1).
•Select the three top polygons. We will bevel these to make fingers. Hit the (b) key to bring up the bevel requester. We want to bevel without an inset, so set all values to 0 except for shift, which should be set at 0.5 m.
• Hit return and you have just created the first joint of the finger.

Because we lowered the top corners, the fingers are slightly separated. If the fingers are too spread for your taste, hit undo and tweak the angles to your lik-


Figure 1: The basic starting block for the hand.

Figure 2: The block object with the finger joints extruded using the bevel command.

ing. To create the next joint of the fingers repeat the beveling procedure. With the second joint you may want to add a slight inset to give them a gradual taper.

This would be a good point to start previewing what our final object will look like.

•Activate the 'OpenGL Smooth Shaded' option in the display panel and with no polygon selected, hit Tab.
•Hitting the (a) key will center our viewing area on the object and there should be an OpenGL display of a chubby looking fork object. Resist the tempta-

Figure 3: Extruding the thumb section.

Figure 4: Hand object with all the fingers beveled out.

tion to modify the object until after we add a thumb (Figure 2).

•Hit Tab again to deactivate the MetaNURBS function. Select the lower left side polygon and again use Bevel with the shift value at 0.25 m and an inset at 0.075 m. After it's beveled, drag and rotate the still selected polygon into a position appropriate for the heel of a thumb.

•Now select the top of the just beveled thumb heel and bevel it again. Because we messed with the perfect geometry of our object, the last thumb bevel will probably need to be adjusted back to square before you bevel it again to make the next knuckle (Figure 3).

The final section is the wrist area.

•Select the bottom most polygon and bevel it with a negative inset of about -0.25 m and a shift of 0.1 m. Again, this does not produce a perfect extrusion so in the Modify menu select Stretch and modify the still-selected polygon until it is mostly square and about 1/3 less then the diameter of the full hand.

•Bevel the bottom-most polygon again with 0.0 inset and a 0.25 m shift, and then again with a 0.25 m inset and a 0.1 m shift to taper it back in (Figure.4).

Fun Starts Here

Hit Tab to reactivate MetaNURBS. You should be rewarded with a fairly good representation of a Mickey Mouse-like hand. It's like magic, isn't it? For making small adjustments to the object while in the MetaNURBS mode, I've found that the drag tool works best. For moving full fingers, select the areas just as you would regular polygons (Figure 5).

When you get your object close to what you want, save it because at some point you will get carried away with how far you can go with MetaNURBS smoothing—trust me, it happens. We want to have the hand and fingers in a neutral position to use as a base to start with and use as a morph source.

We need to convert our MetaNURBS object into a state that can be rendered in layout. This is easily done by entering the Tools panel and hitting the freeze button. In a few moments you will have a fully polygonal object. Save this as 'HandSource.lwo' (Figure 6).

Now that we have our source object, lets make some morph targets.

•Hit the undo button to bring back the original MetaNURBS form, or clear the scene and reload the object.

•Now we'll start altering the object.

•Manipulate the glove into different positions, such as pointing, a fist and thumbs up. When you come to a shape you like, freeze it and save it (Figure 7).

Remember, after you have done a freeze, created and saved the object that is your morph source, it is imperative that you do not add any more geometry to the object. If you add or remove any points or polygons, or even rearrange their sequence, the objects will morph together in interesting but totally useless ways.
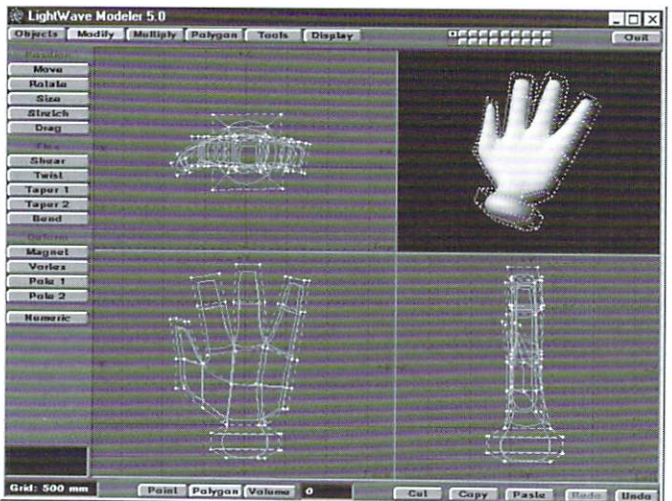

Figure 5: Hand object with fingers with the MetaNURB function turned on.

Figure 6: Hand object after Freeze function is employed to create an object that can be used in Layout.

Figure 3: The final results as it appears in Layout.

Dissolve the CutMetal.lwo and FrontPl.lwo objects to 100%, so they disappear off the screen.

• Go into the Record panel, and select Save RGB Images, give them a filename like clip. The actual file name will end up being clip001, and that's what we want.

• Now render the first frame only. The saved frame should be a black screen with a white dot. Now in the Images panel, select Load Sequence, and select the file clip001. On the Current Image bar, you should have clip (sequence). Leave everything to the default, except the Frame Offset, Change it from 0 to -1.

It's a good time right now to save your work, so save the scene, and save all objects. Now render the rest of the frames, and make sure you start the rendering at frame 2. This is a good time to take another sip of your drink, while you wait a little while for the rendering to take place.

Once all the frames are rendered, we need to crop the frame to a certain size. We'll just reuse the same scene, but with a few modifications.

• First, we will need to bring back the FrontPl.lwo object, so set the dissolve to 0%. Using the FrontPl.lwo object as reference, use the Limited Region setting in the Camera panel, to fit right over the FrontPl.lwo object. In other words, the only area you want to render is that occupied by the plane object. For mine, I brought the camera as close as possible to the object. Then I slightly adjusted the Limited region setting by keeping the key 'l' pressed and dragging the sides of the Limited Region box.

• Make sure that the Frame Offset (Image panel) is set to 0, (it used to be -1) and that all the objects are now dissolved to 100% or cleared from the scene. We don't need then anymore. Set up a new name in the Record panel (Use "cut") then just click the Render button, and re-render the entire animation, starting at frame 1.

This will give us the clip map which will be mapped on the FrontPl.lwo.

• Clear your scene, and reload FrontPl.lwo and CutMetal.lwo. In the Images panel, click the Load Sequence button, and select the cut001 file. Now go back to the Objects panel, and select the FrontPl.lwo object, and click the Clip Map option. For Texture Image, select cut (sequence), and hit the Automatic Sizing button. Also make sure that your Texture Axis is the Z axis.

This just about does it. The cut will be uncovered over time, which is exactly

what we wanted. All we need to do now is add a lens flare, and some sparks, and we're done.

A Few Touchups

To add a little realism to the scene, we need to jazz it up a little. We'll add some sparks flying from the point where the torch is cutting, and a lens flare or two, with some sort of flicker for the light that would be produced.

• In Modeler, in the Tools/Custom button, select the RandPoints plug-in. Set the number of points to 20, and select sphere and fall-off. If you don't have LightWave 4.0 and up, you won't be able to use this plug-in to create random points. Just plop down 20 points in whatever fashion seems random to you.

• Now scale everything down by a factor of 0.15. That will bring the points closer together.

• Select Points2Polys from the Tools/Custom button, then extrude these single sided polys by -1.0 m in the Z axis, in one segment. We need to make it look like the lines are generated from a point. So select the set of 20 points that lie around the origin, and scale them by a factor of 0.01.

• Go through this process of making these spark-looking objects as many times as you want. I did this two more times for a total of three spark objects.

The final result looks somewhat repetitive, but with more objects and a bit of patience, you can get some decent results by mixing the sparks up somewhat randomly. Also, remember that the more points you create for each spark object, the rendering time increases considerably! So keep the point count down.

• Now change the surface of all the sparks to something like 'Sparks surf'. Then save all the objects to Sparks01.lwo, Sparks02.lwo and so on

It's time to put everything together in Layout, and we must not forget to surface the sparks we just created.

• Load up the FrontPl.lwo and CutMetal.lwo. objects, along with all the sparks objects you just created in Modeler. Go into the Surface panel, and select the surface 'Sparks surf'. Plug in these settings.:

Surface color	255, 215, 100
Luminosity	00%
Transparency	0%
Transparency Texture	
Texture Type	Grid
Texture Falloff	Z = 100%
Texture Center	Z = -1
Texture Value	100%
Line Thickness	1.0

This should give the sparks a nice look. Now let's put everything together.

• Add a null object to the scene. Select the Null object and then select Motion Graph and load the motion Curve1.mot (the one you saved earlier). We'll use the null as the parent of all the sparks and the lens flares.

• After loading the motion, adjust the Z position for all the key frames to -0.1. This will offset the null a little bit from the actual surface of the metal object. Next, parent all the spark objects to the null which will now move along the desired path.

You will now need to animate these sparks.

• At frame 0, size all the sparks to 0. Two frames later, scale the first spark to 1.0 and create a key for it. Now go in the Motion Graph panel, and for End Behavior, choose Repeat. This will cycle through this cycle endlessly.

• For the second spark object at frame 0, set its size to 0.5. At frame 1 set its size to 1.0 and at frame 2 set its size to 0. Set its End Behavior to Repeat.

• For the third object, set its size to 1.0 at frame 0, size of 0 at frame 1 and 0.5 at frame 2. Remember to select Repeat for End Behavior. If you have more that 3 spark objects, you'll have to a bit of choreographing to get objects the right size at the right moment, but it can be done (Figure 4).

Now for the lens flare.

• Go into the Lights panel, and add as many flares you want. I only used one, but more flares can give a better illusion of randomness. Go into the Flare Options panel, and click the envelope button for the flare intensity. Here you can set down points in a random order, to give the illusion of the light getting brighter and dimmer as time goes on. Remember that the light won't vary very much. Just keep it around 80%, and vary it from 95% to 75% or so. This will create a shimmering light.

Another way to get the random envelope is to use the ARexx script Random Envelope Generator in Modeler on an Amiga. I don't know if this exists for any other platform as a plug-in, so do it manually, as mentioned a bit further down. This ARexx script will ask you for ranges, how long the animation is, and so on. Be creative, but remember what I've told

you: You want a shimmering light, not a light that is at 0% one frame and ramps up the next to 100%. But don't let me stop you from experimenting! I'm only here to point you in the right direction and to give you some pointers. The rest is up to you.

Well, we are done, and you might want to take the last few sips of your beverage before you toss it. Render your animation, and you should have a half-way decent looking torch cutting. This is by no means an end product. Like I said earlier, be creative with it! And let me know if you find other off the wall uses for this tutorial. Good Luck!

LWP

Adrian Onsen can be reached at either aonsen@easynet.on.ca at his home account, or at aonsen@acs.ryerson.ca at his school account, where he studies engineering.

Tip of the Month

Make the Most of ScreamerNet

Did you ever run into a situation in which you have a few small scenes to render, but will be out of the office for a long time? At times like this, a batch render would sure be great. If you're on an Amiga, you can use the Batch Render feature of Power Macros. Or, whatever system you're on, you can use ScreamerNet.

And did you know that ScreamerNet decreases rendering time? You can set up ScreamerNet on a Pentium system, for example, to render more than one scene at a time and you can fully maximize your CPU by running ScreamerNet twice. That' right, twice! on one machine!

Here's what you do from a Windows/WindowsNT setup:

1. You need to create a Working Directory for ScreamerNet. Create a folder somewhere on your drive, and name it SN (for example).

2. From the program manager, select file/new. Choose a new program item.

3. In the Description entry, type in ScreamerNet.

4. In the Command Line entry, type C:\NEWTEK\LWSN -2 C:\SN\JOB1 C:\SN\ACK1

5. Now, enter the working directory, such as C:\SN.

6. Lastly, create a shortcut key, if you like. This will create an icon that you can click and run to initialize ScreamerNet. You can repeat this process but when creating the second icon, enter JOB2 and ACK2.

Try this out. Double click the new icon you've just created. Next, enter Layout and open the ScreamerNet panel (SN). Choose ScreamerNetII as the Net Rendering Method. Toward the bottom middle of the screen, select Add Scene To List. Select a scene to render. Select another. ScreamerNet will check for available CPUs and inform you how many were found.

If you created two ScreamerNet icons in the program manager and ran both of them, LightWave will find two CPU's to render with (or so it thinks), even on one machine. What this basically will do is fully maximize your CPU. When one frame is being loaded,and saved, another is rendering. This shaves a little time off of each frame's render.

By the way, there's a few things you should be aware of before using ScreamerNet for

rendering. Make sure that your sceneis ready to go before you save it. Meaning, set up the Record Panel and where the frames should be saved. Set the resolution and anti-aliasing, and save the scene. Saving the scene with the Record settings enabled allows ScreamerNet to know where to save the frames. Lastly, you don't need to have the scene loaded in Layout to use ScreamerNet. It's better if you don't have it loaded as it will take up more RAM. Simply open Layout, and run SN.

Here's a comparison of a Layout render versus a ScreamerNet render on a single Pentium 100 machine with Windows NT.

Texture Examples scene
Medium Res, Low AA at 8
Layout Render   3.43
SN Render   2.59

Another example is a simple logo animation that took 2.03, but only 1.10 per frame with ScreamerNet. Over the course of an entire animation, the time savings add up. Running two copies of ScreamerNet will cut even a few more seconds off each frame. However, if you have a hefty scene and limited RAM, you may want to only run one ScreamerNet.

Now that you can render things faster and set up a batch render, you'll have plenty of time to do other things, or sleep.

by Dan Ablan

LIGHTWAVE**PRO** 13

The Surfacing Jungle
Tips and Techniques for Polishing Up Your Images

by David Ebner

A lot of 3D animators out there seem to have their strength in one area or another. Yet every aspect—lighting, modeling, surfacing and animation—is important. Any weakness in just one of these areas can bring the whole piece down. I've seen a lot of tutorials aimed at modeling and even more dealing with animation. This one is strictly designed for obtaining more desirable surfaces and offers suggestions and a look at some of the new possibilities offered by LightWave 5.0. We'll also examine a few new plug-ins that expand the LightWave's capabilities.

Approaching Your Surfaces

I like to work on surfaces under a fresh scene aside from that containing animation movement. Placing a couple of colorless lights under the general conditions and lighting ratios of which the subject will be placed, I start examining the model and surface breakdowns. By working in this surfacing scene, I can also make key frames, like a scratch pad, up close on certain surfaces and with differing orientations. Your objects should be surfaced so they can be successfully placed under different lighting conditions to affect the tone, illumination and mood while the object itself maintains balance.

A lot of people are finding that because they have surfaced their model during or after animating and lighting, they end up having to resurface or repaint maps for different shots. Your car shouldn't get a new paint job depending on the scene and neither should your surfacing. True, a little tweaking from scene to scene may be advantageous for experienced surfacers, but generally try surfacing models this way.

Image mapping

I see people still trying to map an entire surface of an object with a single image. Break that surface down into smaller, geographic regions. You can still use the same surface characteristics such as diffusion and specularity, but by breaking up the surface into regions of differing axes perpendicularly, and by using related images for these surfaces you'll greatly reduce the smearing and stretching of images along your object. In the example pictured, the body of this creature was divided into several surfaces based on whether or not I would map downwards (Y axis), sideways (X axis) or front on (Z axis). Once these regions were assigned with their appropriate surface names, reference images from the Modeler screen were captured (screen grabbed) and images were painted. Capturing the reference image along the same view that is being mapped enables you to perfectly line up the texture and details with the geometry and features.

With some time invested in making sure the edges of the image maps don't contrast in color or value against the neighboring images, you can achieve very precisely detailed surfaces making up what appears to be one uniform surface. Even with the introduction of 3D paint programs I still have not experienced the precise and predictable results as with the methods just described. You simply can not bypass the work it takes when it comes to achieving believable image mapping.


Figure 1: The panther's head was surfaced separately from the body.

Diffusion/Luminosity/Lighting

The source of light can be thought of as an umbrella under which many surfaces are present. The relationships of brightness between these surfaces are primarily defined by the diffusion and luminosity against other surfaces.

First, it is important to note that within the color values or color mapping there also is the definition of value or tone, meaning lightness and darkness. With three-dimensional shapes, it is usually desirable to have no shading introduced by the color map. The distinction of differing colors, therefore, is to allow the lighting given in the scene to properly shade the surface. Upon the introduction of light, the model takes on a new look as it is being rendered. The shading and addition of shadows give the model further form.

Upon default, the Diffuse Level is full strength allowing the full effect of the lighting given, while Luminosity is set to off at zero. To affect the tone of a surface beyond that of the color map, enter a new value or introduce a texture defining the diffusion. This diffusion surfacing control is independent of light levels creating specularity highlights or shadows.

The introduction of luminosity further intensifies the brightness of the whole surface. This illumination, however, is not derived from the source of light but is obtained evenly by its own value or applied in relation to an image or procedural map. This self-illumination reduces the effect of light shading and of shadows and may completely overcome them. Since luminosity is independent of light, surfaces which acquire their brightness from this control do not adapt well in different light. This control should be reserved for those surfaces truly generating their own light or for simulating certain lighting effects.

The exchanging of diffusion value into that of luminosity can be a quick way of shadow strength reduction. An example being a backdrop object in which you

do not desire the shadow of a logo to be so dark. With some luminosity/diffusion balancing you can achieve this without lessening the light, though the shading effects also will be diminished. Before breaking the rules it is important to understand these relationships and characteristics

Additive Surfaces

This switch in the Surfaces panel is often overlooked and sometimes not taken advantage of due to the lack of understanding. When this mode is invoked the entire surface is treated differently and for an entirely new purpose. This tool gives you the ability to gently or intensely lighten anything behind these surfaces with defined shapes, procedures or images. A good example of its use is in cloud additions in which the dark tones always marry with the back drop. Also I accomplish light rays, subtle surface glimmers, artificial specular highlights and secondary reflections with it. Only values above 0, 0, 0 with this kind of surface have any effect as the tone also determines the strength. A transparency map therefore is not needed in most cases since the dark regions of the surface are transparent by default. The overall strength of this surface can be controlled with the diffusion or luminosity values, though usually I tend to use these surfaces as luminous as opposed to diffuse.

Reflection Mapping in Variable Lighting

If your model is going to be placed in different lighting setups you'll have to remember that a reflective surface will acquire some of its brilliance (if not all) from the reflectivity. This is particularly important to note if using an image to define the reflection as opposed to the environment. The image used must "know" or be modified for the lighting conditions. A sequence of images or surface morphing may be necessary to get realistic looking reflections while animating through changing light. Without compensating for conditions, reflections tend to be way too bright in darker lighting.

Multi-Mapping with 5.0

Beginning LightWave surfacers often keep forgetting that some of the algorithmic textures use the surface color or value in the main part of the surfaces panel with the effect. Veins for instance adds its effect to the color value. This concept is important to grasp because starting with 5.0 you can have multitudes of textures and maps and it is important to know how they are ordered and combined.

LightWave's image mapping capability has now expanded allowing you the ability to stack or blend image maps and procedural textures together in each of the texture areas. This is achieved with a new button labeled "Add New Texture" located alongside "Previous" and "Next" buttons for navigation. The maximum strength of a texture (whether image based or procedural) is controlled with a texture opacity field.

Alongside each image map you now have the option of adding an adjacent image defining the strength of the image map. The tonal values of an image applied as a texture alpha defines the strength (transparency) of the map. This way you can apply color, diffusion and other factors in just the areas you want. A good example of this can be related while surfacing a planet's atmosphere. Some colorful gases are desired within your already defined cloud work. Without building an additional set of geometry with more transparency calculations factored in, you can now lay in another image map with a high contrast alpha image map to achieve the desired effect. This way you avoid incurring painful rendering times and do not have to alter the original image map. If the new gases are too strong, simply lower their texture value.

Also with 5.0 you can further limit the extent of your image map by checking off the width tiling on the texture height and/or texture width. By doing so your image will only affect the area within the numeric sizes centered, of course, around the texture center. It is now possible to have a color map such as photo of cardboard for the whole box, and the use a second image sized down and re-centered for the label all in one surface.

Figure 2: It's a good practice to screen capture reference images from Modeler

Remember that, as soon as you have an image map used with 100 percent Opacity, no other textures previous to that one will show their effect (assuming that you are not using Texture Falloff) as the last 100 percent Opacity image map fully replaces any previous surface values. You can test this by trying to use multiple transparency maps, each with 100 percent Opacity. Only the last will show its effect, while "hiding" the others.

With three different limiters you have possibilities never before so easily achieved in combining several images for surfacing.

With procedural maps there already is an inherent definition of texture strength though you still have the overall texture opacity level which can be used to dissolve out a texture. Random bumps can be made even more irregular by combining two textures of different coverage, just as additional variances of underwater and veins can easily be obtained.

The ordering of textures is important to keep in mind. The last texture to be added can be thought of as the top layer of application or the final coat of paint. The transparency of this texture or areas of non application can be thought of as worn away paint revealing the next layer. The bottom of the stack is the surface color or value in the main part of the surfaces panel.

Though the combination of textures so far is limited to priority and strength weighted in a normal application, I am looking forward to one day maybe seeing some other methods such as additive and subtractive.

Surfacing Tools in Modeler

Also exciting is the surfacing tools offered directly in Modeler. You can now characterize your surface's color, diffusion level, specularity and glossiness, double-sided on/off, smoothness and smoothness angle all in Modeler. These are limited in that you can not apply texture maps; that is to say until you use the "Power Texture" plug-in. With this new plug-in which looks much like Fori's Power View for OpenGL viewing before OpenGL modeling and Layout views became a reality, you can now interactively apply texture to your surface, seeing the results as you position and scale your image.

Some More 5.0 Surfacing Additions

LightWave 5.0 is a major advancement, especially in the area of interactivity. This of course is exemplified with the dynamic OpenGL layout, modeling and previewing modes. Another rad tool with 5.0 is the image viewer accessible through the generic plug ins interface. This viewer not only displays images loaded but also a frame from an image sequence. Along side "Polygon Edges" in the Objects panel is "Cel Look Edges" which gives a color and weight definable outline that follows the contours of surfaces. This is very effective when balancing luminosity against your diffusion, achieving more of a two-dimensional ren-

From the Net
A Steaming Hot Cup of Coffee!

by Dan Ablan

They told me when I was young to "enjoy" it now. But of course, like most kids, you can't wait until you're older. The time frame from the start of the school year to the end seemed like an eternity. Now, we find out that the adults were right - time really does go by fast. So fast in fact, that it's time to start thinking about the holidays again. When you think about fall and winter, you often think of hot chocolate (sometimes with a little kick to it), or hot coffee. Of course, if you're not from an area of the country that can get to -20 degrees during the winter, your thoughts of warm liquids may not be as strong. In either case, you may want to animate such a setting. Ironically, three people have e-mailed me asking how you would go about making a steaming cup of coffee, or hot chocolate as the case may be. Well, if you've ever played with fire (in LightWave 3D, that is) you can easily create steam.

Somthing's Brewin'

When you think of winter, hot chocolate, snow, you think of thick full mugs, rather than dainty tea cups. Your first step to creating this masterpiece is to build the mug from which the steam will rise. I'll be using MetaNURBS in LightWave 5.0, but if you are using LightWave 3.0 through 4.0, you can follow the mug tutorial in the original LightWave/Toaster manual, or use a Boolean Subtract operation on an elongated disc. The mug is just a reference for this, and can be changed later. But, for 5.0 users, follow along.

1. From the Objects panel in Modeler, select Box, and press n to call up the numeric requester. Press the reset button to create a mug that is 500 mm in size, and click OK. Don't worry about modeling to scale right now, the importance of this tutorial is the steam.
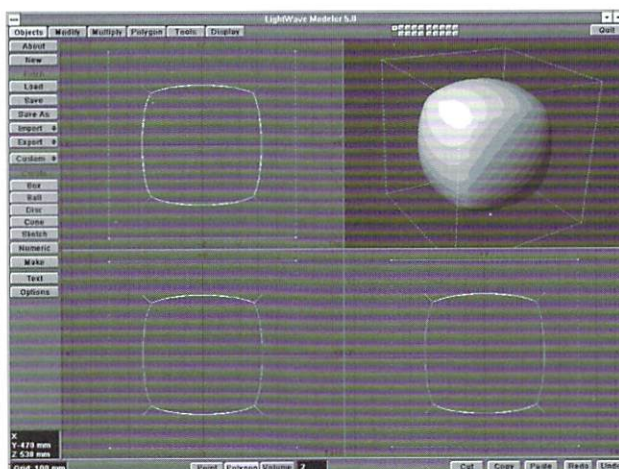
2. Once the box is made, press the TAB key to

Figure 1:Once the box is made, select the bottom and top polygons.


Figure 2: The bottom and top polygons are Beveled with an Inset sharpening out the rounded edges.

turn on MetaNURBS (Ctrl i works too).

3. Select the Polygon selection tool from the bottom of the screen. Then, click on the bottom side of the NURB cage. Once you do that, you should see that there are three polygons selected in the selection information box next to the Volume tools.

4. Click the side edge of the NURB cage to deselect the middle polygons leaving only the bottom

and top polygons selected, as in Figure 1.

5. With the bottom polygon selected, press b to call up the Bevel tools. Enter an Inset value of 10 mm and leave all other values set to 0. Click OK, and you'll see that you now have a not so ball-like, but rather, the beginning of a mug, as in Figure 2.

6. Deselect the bottom polygon, leaving just the top polygon selected. Again hit the b key for the Bevel tools. Enter an Inset value of 40 mm and click OK. Press b once more, and enter another Inset value of 40 mm. This will create the rim of the mug. Now, Beveling again with all values set at zero other than Shift, set a Shift value of -850 mm. This takes the selected polygon and pulls it down into the mug.

7. Lastly, Bevel the bottom inside polygon, which should still be selected, and set an Inset value of 20 mm, leaving all other settings at zero. Press Ctrl d to Freeze the NURB cage turning it into polygons. The last thing you need to do is create the handle. Save the mug as it is. Then, in a new layer, do the following:

1. Make a box with these values:

	Low	High	Segments
X	500 mm	620 mm	1
Y	-400 mm	400 mm	12
Z	-50 mm	50 mm	1

2. Click OK to make the box. Press the TAB key to turn on MetaNURBS, and then select the Bend tool from the Modify panel. Make sure the mug is selected in a background layer for reference. From the Top view, click and drag the segmented box over to the left roughly 180 degrees, as in Figure 3.

4. Now select the Rotate tool (hit y on the keyboard) and either rotate it 90 degrees on the Z axis.

Figure 3: A handle begins to take shape for the mug.

Figure 4: Points are added and a closed curve is made.

Position the handle up towards the right side of the mug from the face view. Once in place, press Ctrl d to Freeze the NURB cage into a polygonal mesh. Before you Freeze the NURB cage, you can drag and pull points to create a cool looking handle, or a handle with grooves.

5. With the mug still in a background layer, and the handle lined up, press Shift b for Boolean operations and choose Union. In a moment, your mug will have a handle.

Making Steam

Now you need to make the steam. The object you need to build is even easier than the mug. In a new layer, select the Disc tool.

1. Press (n) to call up it's numeric requester,

Figure 5: The Lathe requester for making the Steam object.

and enter these values:

Sides	16
Segments 1	
Bottom 0	
Top	0
Axis	Z
Center	

X	Y	Z
25 mm	800 mm	0 m
Radii		
X	Y	Z
325 mm	300 mm	0 m

2. Click OK to make the disc. Now press k to remove or "kill" the polygons leaving only the points. Select the points that make up the disc from roughly bottom center to top right, or 6 p.m. to 2 a.m. if this were a clock. Once just those polygons are selected, press the z key to delete them. You should be left with four points.

3. Now add roughly six more points evenly spaced (about one grid square) up from the last point in the face view. Make the last point you add on the center axis, but at the bottom near where you started. Once all of the points are in place, select them in order, and press Ctrl p to make a closed curve, just like the one in Figure 4. Use this figure as a reference for point placement. This doesn't have to be exact, but basically similar to the figure.

4. Select the Lathe tool from the Multiply menu, and Lathe the curve around the Y axis. Figure 5 shows the Lathe requester.

5. Click OK and press Enter to Lathe the curve.

6. Press q and in the Change Surface requester, name the object Steamand save the object.

## It's Steaming Up

Load both the Steam and the Mug object into Layout. Parent the Steam to the mug. This isn't a very important step, but helps if you decide to move For the mug, make it a simple whitish gray, and perhaps had a bit of fractal noise to dirty it up a bit. For the Steam, use the following:

Surface Color	120, 120, 120		
Luminosity	100%		
Diffuse Level	0%		
Specular Level	0%		
Transparency	100%		
Transparency Texture			
Texture Type	Fractal Noise		
Texture Opacity		100%	
Texture Size			
X .06	Y .25	Z .06	
Texture Falloff			
X 18.0	Y 60.0	Z 18.0	
Texture Velocity	Y .05		
Texture Value	0%		
Frequencies	1		
Contrast	1.0		
Edge Transparency Transparent			
Smoothing On	Double Sided On		

You can add a Displacement Map to the Steam object to give it a wiggle. The Texture Velocity will make the steam move upwards, but the Displacement Map can sort of shift the steam a bit from left to right. The final cup is added to a kitchen counter with the Sunday paper (see color pages).
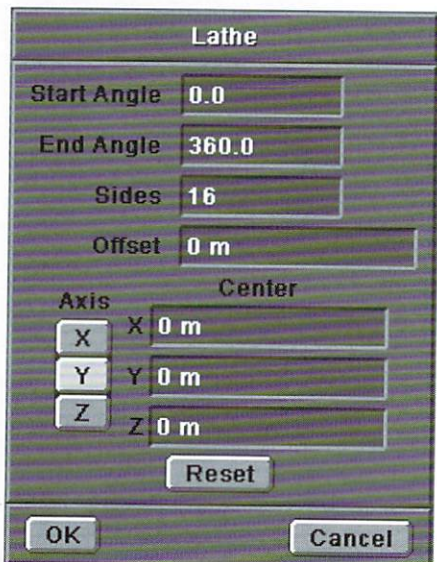
Jungle

dering effect. Also try this with Allen Hasting's Cel-Look shader.

Pertaining to the overall quality of output, Anti-aliasing has been improved upon in different ways. There is now a numeric control over the anti-aliasing strength of image maps as well as the addition of new anti-aliasing rendering options which offer greater results.

New Plug-ins

It is very exciting seeing LightWave and third-party tools develop with more and more diversity and capability. I'm impressed with this growth and the fact that the ease of use and organization of the program and features have been maintained. One very exciting new add-on being offered now is the film grain plug-in by Steven Worley. Originally developed for and commissioned by our facility, Steve's grain plug-in is perfect for aiding in the integration of 3D into feature film work. Steve has achieved the realistic appearance and behavior of film grain and has programmed this with user-adjustable parameters for surface application, selected surface application or complete image application. He has also developed an algorithm for the removal of film grain from shot footage. His attention to detail goes to the level of grain strength tapering based upon the film print transparency to grain ratio of actual printed film.

Another impressive surface plug-in is included with version 5.0, the surface blur. With this plug-in you can now soften surfaces individually. At SIGGRAPH, I was able to take a look at some very impressive additions such as Particle Storm, a new and improved Mesh-Paint, Vertilectric, Wavefilter and WavenetPRO along with a peek at a few in development. These plug-ins are all excellent, adding additional tools and ease of use to LightWave. I'm sure they will be explored even more in future articles to come in LIGHTWAVEPRO!

LWP

David Ebner is an animator and president of Computer Cafe Inc. in Santa Maria, Calif., who enjoys good cinema and good food. Just like all of you. I hope.

MetaNURBs

Because the surface attributes of the morph targets are never seen, you can change the names of any polygon you wish. In this case we can name the different parts of the glove to make selecting the polygons easier for later on when quarters start getting cramped and it is difficult to sort out which polygon belongs to which finger.

Once you have a number of targets, we can jump into Layout and see if this really works

Stay on Target

Enter layout and set up a typical object morph.

• Load the source object and two target objects. In the objects panel have the current object be "HandSource.lwo," select MTSE for multiple morph targets (Multiple Target, Single Envelope), and put "HandTargetA.lwo" in the metamorph target box.

• Select "HandTargetA.lwo" as the current object and choose "HandTargetB.lwo" as the meta-morph target.

• Go back to "HandSource.lwo" as the current object and click the (E) button to create an envelope. When the envelope is between 0% and 100%, the object will morph between the source object and the first target. When the envelope is between 100% and 200%, the object will morph from the first target to the second.

• Do a preview to make sure every thing is working right. To make a compliment of the left hand we have created, make a clone of the source object and under the stretch command numerically enter a value of -1 on the Z axis. This gives you a mirror image of the original.

This technique saves on geometry but because we're basically turning the object inside out, we have to turn on Double-sided in the surfaces panel to have it render properly. It will clone with the exact same morph envelope, so you probably want to alter that to loosen up the animation (Figure 8).

About Character?

The technique described above for gloves will work exactly the same for faces. Rather than walking you through the creation of a head and face, I suggest you refer to the LightWave User Guide and "Advanced Tutorial #7," which describes the creation of a cartoonish character object. By following the rules laid out above, you can easily create several morph targets of the object. It is amazing how easy it is to change an expression by simply dragging a point or two. Grabbing the right points at the edge of the mouth, for instance, quickly produces a groan-like expression. It is great fun to see how many expressions you can cycle through by hitting the undo and redo buttons.

MetaNURBs is so easy and intuitive, once you work through one tutorial, you will have learned enough to keep you busy for a long time.

LWP

*Patrik Beck is a cool animator guy who can write. Reach him by e-mail at zippie@execpc.com.*

# BUSINESS REPLY MAIL

FIRST-CLASS MAIL   PERMIT NO. 1286   BOULDER, CO

POSTAGE WILL BE PAID BY ADDRESSEE

LightwavePro

P.O. Box 54240
Boulder, CO 80323-4240

Burn Baby, Burn!

Sparks fly as the flame from an acetylene torch cuts through a metal plate in Adrian Onsen's "Cutting Edge" tutorial.
*Copyright©1996 Adrian Onsen*

The Panther

Successful surfacing takes proper planning and attention to detail. See David Ebner's tutorial for more information.
Copyright©1996 David Ebner

See for yourself what the future holds for digital video and 3D animation!

Video Toaster User Expo '96

- Sun Microsystem's **Java™**
- Intergraph's **RenderGL™**
- Digital Semiconductor's **FX32!™**
- Apple's **QuickDraw 3D™**
- Amiga's **6000T**
- Intel's **MMX™ Technologies**

- **Exhibits:** Over 30 vendors showing Video Toaster, Flyer, and LightWave 3D compatible products.
- **Specialty Sessions:** Digital Technologies on the Internet, 3D in Print, Artists in Action and the Tools They Use.
- **NewTek University:** featuring training on LightWave 3D, Video Toaster and Flyer techniques.

Seminars November 5-8 Exhibits November 7-8 Universal City Hilton

---

PRE-REGISTERED EXHIBIT PASS ONLY $20:

- Exhibits
- Keynotes
- Grand prize drawings
- Strategy Sessions
- Specialty Sessions
- LightWave Theater
- Technology Forums: Sun Microsystems, Intel, Digital Semiconductor, Apple, Intergraph and Amiga!

NEWTEK UNIVERSITY

Learn how you can make the most of your powerful equipment.

More than 24 Video Toaster, Flyer, and LightWave 3D seminars to choose from!

Call

1-800-643-3976

for more information

Grand Prize Give Away!

Loaded 500 MHz Raptor ReFlex WorkStation
Compliments of DeskStation Technology

PLUS, additional drawings from Intergraph, Viewpoint Datalabs and more!