THE JOURNAL FOR LIGHTWAVE 3D® ANIMATORS

LIGHTWAVE™ PRO

an Avid Media Group, Inc. Journal

Volume 3 Issue Number 12 • December 1995 • $8.00

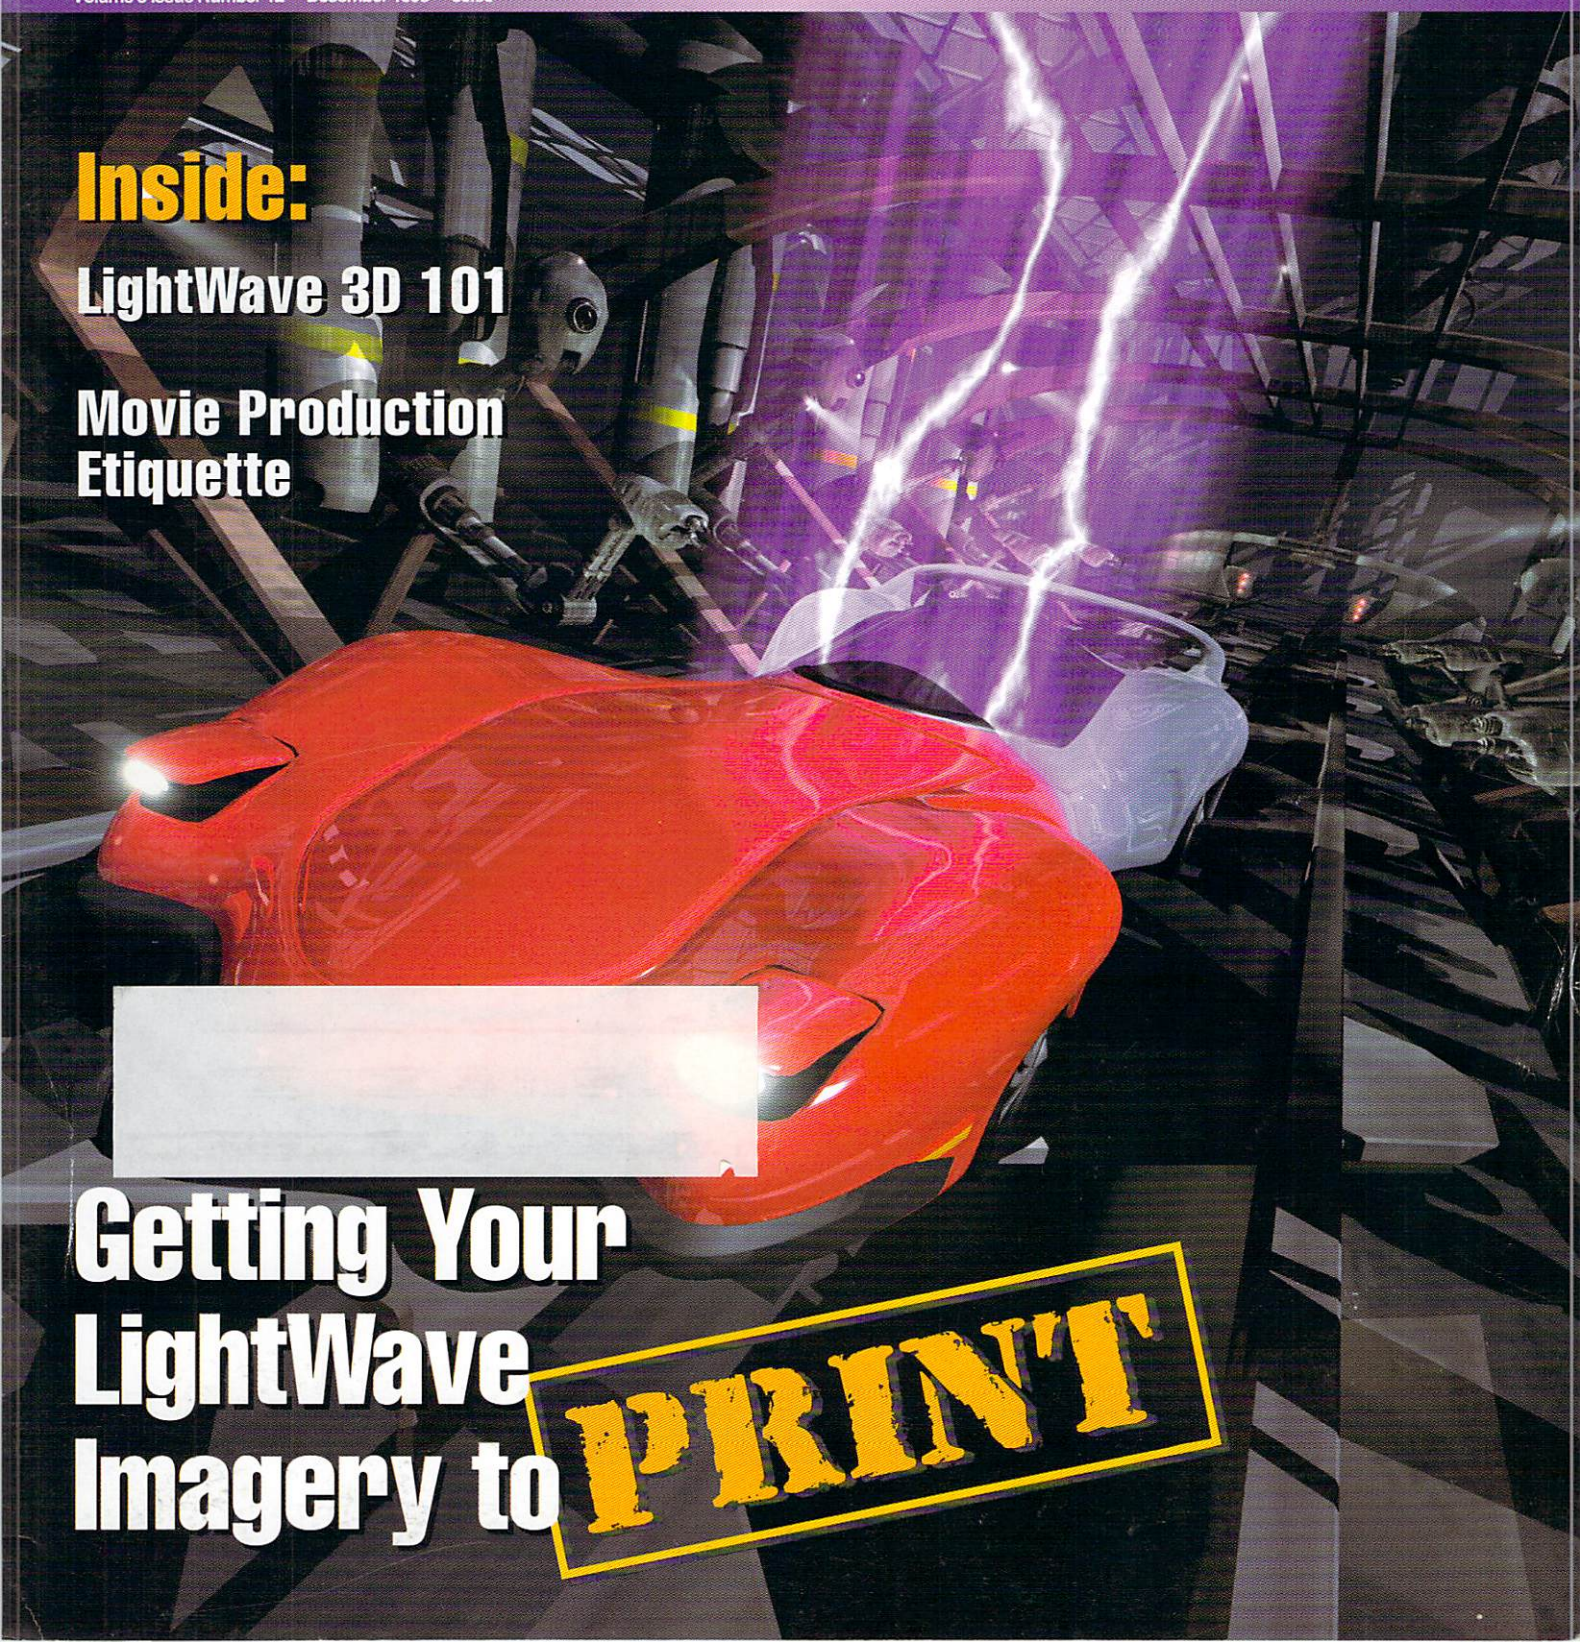

Inside:

LightWave 3D 101

Movie Production Etiquette

Getting Your LightWave Imagery to PRINT

ANIMATION STILLS

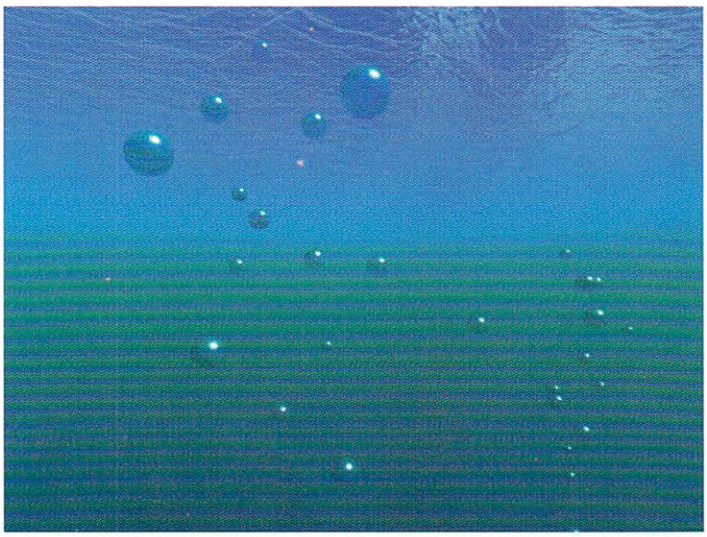

Up From the Depths

The bubbles in this image were created using Robert Hood's BUBBLES.P plug-in for Intel-based LightWave. See "lwpro@internet.online," page 4.

Cuban Missile Crisis

This image from *seaQuest 2032* is the final version of last month's *LIGHTWAVEPRO* cover. Creator Greg Teegarden decided that the shot looked better backlit, as the ship on the horizon lacked sufficient detail to hold up well for the final shot. Backlighting also added a nice specular breakup to the surface of the ocean. The surface itself has raytraced reflections to take the clouds into account. Trace reflections were not used for last month's cover because of time constraints. The final image seen here is what aired for the episode.

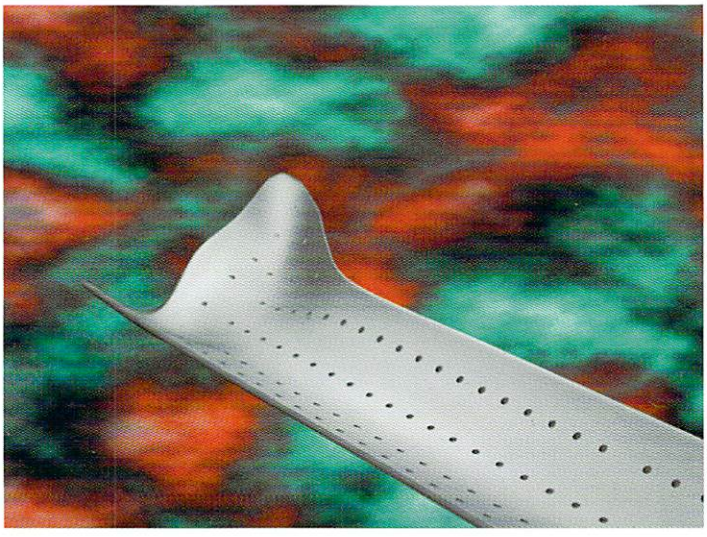

Hand Support

The result of endless patching: "Rasta Splint." This image of a hand splint was created using splines and metaform. See "LightWave 101," page 6.

EDITOR'S MESSAGE

by John Gross

It's the end of another year—boy, that went fast! Hopefully you received everything that you wanted this holiday season, including LightWave 4.0!

A whole year has gone by since the "initial" LightWave 4.0 ship date and, as of this writing, it still hasn't finally shipped in all versions. Perhaps NewTek has been a bit optimistic (to be polite) in their announcement of ship dates, but you know what? By the time you get the final shipping 4.0 in your hands, I think it will have been worth the wait. In the porting of LightWave to other platforms, much was learned, a lot of nasty bugs were discovered and fixed, and the final product is going to be that much better because of it.

There have been problems with HIIP and problems with plug-ins and problems with memory leaks (thank you, Microsoft), but all of those things are taken care of and NewTek will be providing a patch that will work with any prerelease or "final" version to the final, final version of LightWave. This patch will be provided on-line so you'll be able to update your version of LightWave without a hassle.

What will be taken care of in the patch? Functional HIIP; a workaround for the Microsoft bug that doesn't free memory (ever run out of virtual memory when rendering or using an image sequence?); more robust plug-in architecture; and, as a special bonus, Allen Hastings has greatly decreased rendering time under Windows 95 and Windows NT! Don't bother calling NewTek to find out when the patch will be ready—they will be posting information as soon as it's done.

I have a feeling that, since LightWave has been successfully ported over and there are two additional full-time programmers on the LightWave team (Arnie Cachelin and Fori Owurowa), we're going to see some major new changes in the next version of LightWave. NewTek is being very tight-lipped about the feature list and shipping date (I think they've learned their lesson), but I wouldn't

see Editor's Message, page 5

TABLE OF CONTENTS Vol. 3, No. 12

LIGHTWAVEPRO
an Avid Media Group, Inc. Journal

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

Editor ..... John Gross
Executive Editor ..... Tom Patrick McAuliffe
Managing Editor ..... Corey Cohen
Associate Editor ..... Julie Sheikman
Production Manager ..... Kristin Fladager
Art Production ..... Berimbau Design
Cover Design ..... Berimbau Design
Fulfillment Coordinator ..... Kris Nixon
Circulation Assistant ..... Debra Goldsworthy
Contributing Writers ..... Dan Ablan
..... Adam Chrystie
..... Mark Thompson
..... Glyn Williams
..... Tom Williamson
Group Publisher ..... Michael D. Kornet

Editorial Office: Avid Media Group, Inc.
1308 Orleans Drive, Sunnyvale, CA 94089
Telephone: 408-743-9250; Fax: 408-743-9251
WWW site: http://www.portal.com/~amg
E-mail John Gross at jgross@netcom.com
Are you interested in writing for *LIGHTWAVEPRO* or submitting images? If so, contact us at our office or electronically.

LIGHTWAVEPRO (Vol. 3, No. 12); (ISSN 1076-7819) is published monthly by Avid Media Group, Inc., 1308 Orleans Drive, Sunnyvale, CA 94089. A one-year subscription (12 issues) in the U.S. and its possessions is $48 (U.S.); Canada/Mexico, $60 (U.S.); overseas, $84 (U.S.). To subscribe, call toll-free 1-800-322-2843. Allow 4 to 6 weeks for first issue to arrive. Second-class postage rate paid at Sunnyvale, CA and additional mailing offices. POSTMASTER: Send address changes to *LIGHTWAVEPRO*, 1308 Orleans Drive, Sunnyvale, CA 94089.

This month's cover image was created by Aristomenis (Meni) Tsirbas, winner of the 1995 LightWave Animation Contest. It is part of a 30-second commercial for TV and film promoting Montreal's yearly auto show. The project was created with LightWave 4.0 pre-release. The car is a Metaform object using Smooth Shift for additional detail. Meni is the Creative Director at Montreal-based Trimension Inc. (514) 937-5434.

lwpro@internet.online

by Dan Ablan

Since the holiday season is just about over and some of you may not have gotten what you really wanted giftwise, this month's column may help you out. I've been surfing quite heavily lately, most nights until about 3 a.m., and it appears that there are even more and more LightWave-related sites out there. In fact, a recent report in a national newspaper stated that by June 1996, there will be 1 million more World Wide Web sites on the Internet. Yikes! You can be sure that a good percentage of those sites will be LightWave-related.

I've been spending a lot of time at ftp sites this past month, checking out some of the wonderful, free LightWave objects, images, macros and plug-ins available. I found some really great, useful objects and plug-ins. Though we cannot include them on a *LIGHTWAVEPRO* companion disk due to copyright laws, we can list the sites for you to surf yourself.

Some time ago, I found a macro for the Amiga called "Bubbles." It was a Modeler macro that generated an upward stream of bubbles, ranging in size from small to large. Once you've created some bubbles with this script (now available as a plug-in for PC LightWave), you can move the stream of bubbles in LightWave to give an underwater bubble-traveling-to-the-surface look. You can also use these for bubbles from a bubble maker, etc.

Modeler ARexx scripts BUBBLES.LWM and ABSOLUTE.LWM are now available for the Intel Windows NT platform thanks to Robert Hood. The BUBBLES.P plug-in (for Modeler) will convert all points in an object to spheres (globe or tessellated) of a radius in a range you specify. You can use the Random Points Distribution plug-in to create the points, and then this plug-in to generate the "bubbles."

ABSOLUTE.P is a nice additional plug-in that will move an object in the currently selected foreground or foregrounds to an absolute location in vector space. If you used it, you'll know that the numeric requester in Modeler will only allow you to specify relative offsets from the object's current location.

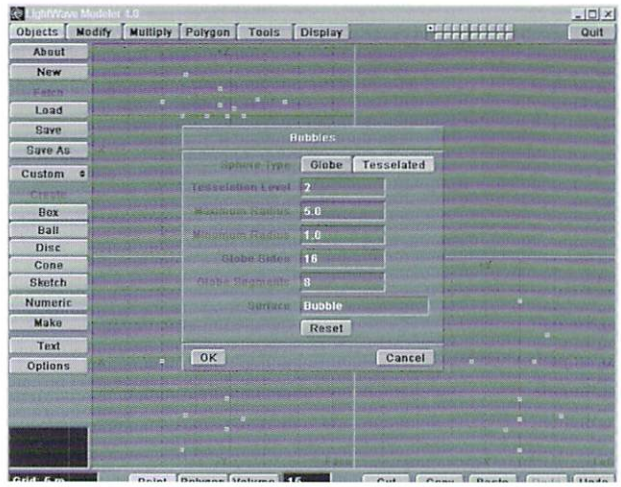

Applying the BUBBLES.P plug-in presents a requester where you can specify bubble parameters.

Both plug-ins can be found by ftp'ing to tomahawk.welch.jhu.edu(/pub/ LW_NT/incoming/arexx) or ftp.netcom.com-(/pub/bh/bhood).

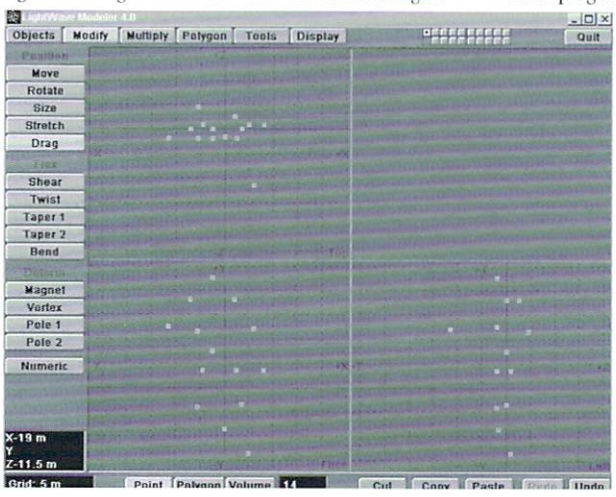

Use the Random Points Distribution macro/plug-in to generate a series of random points.

These sites contain countless other items—image maps, objects, and sometimes scene files. At the tomahawk.welch site, remember to look up directories above the ARexx directory for other available objects and scenes.

If you have not connected to an ftp site before, here's how it works: ftp stands for file transfer protocol. I ftp through a program called Chameleon, which provides me with an interface that looks somewhat like DirectoryOpus or PC File Manager. On the left side of this interface are my system directories. By selecting Connect and entering an ftp address such as the one above, plus the correct password, the right side of the screen becomes the directory of the ftp site. Simply find the directories and files you want and select copy. It's like a universal network, rather than two directories in just your own computer. As far as passwords are concerned, if there is no password listed with the posting of the address, you can assume that it's an anonymous ftp site. So, when you're ftp'ing to a site, enter in the word anonymous for your user name, and anonymous (or sometimes your complete e-mail address) for your password.

Instead of ftp, though, you can simply download from a Web site. From your WWW browsing software—Netscape, for example—enter the address http://tomahawk.welch.jhu.edu and follow the pages as instructed. It takes a little longer to download this way, but it's often easier. You can also ftp from a DOS prompt by typing a command such as ftp ftp.netcom.com.

Something else to look into at the tomahawk site is Power Tips Pro (PTP). PTP evolved from a document featuring what are commonly termed "frequently asked questions" (FAQ) that Dean Scott began com-

piling in fall 1994 as plain text. The source of its content came from the FidoNet echo, Video Toaster.

After being well-received by several LightWave users seeking answers to common questions on how to more effectively use the animation program, Ronnie Norwood came up with the idea to code Dean's LWFAQ into the AmigaGuide hypertext format (with Dean's permission, of course). The two of them became partners to revise and thoroughly edit several times over that first FAQ. A year later, PTP was born, both in AmigaGuide and Windows Help formats. Every byte of PTP in some way or another made its way through the phone lines, traveling back and forth from Dean's 486DLC/33 IBM-clone PC to Ronnie's Amiga 4000. From raw screen grabs and unedited source text to processed images and formatted text, PTP came to be entirely via electronic means. What users now have at their disposal is a fully cross-referenced guide of 100 different questions and answers on over two dozen topics for both Modeler and Layout. The information, though mostly concerning LightWave on its native Amiga, is still indispensable for new LightWave ports to the many other platforms. Check it out.

Another great site for free objects is the Avalon site at http://www.viewpoint.com/avalon.html, a link from the ViewPoint Datalabs home page. Some of the models are very nice; others need some cleanup. Not all are LightWave format, though. What you can do is download the 3DS or DXF objects, and change the format to LWO through any program that will do so on the PC, such as PolyForm or Interchange. You can also load 3DS, DXF and Wavefront objects directly into Modeler version 4.0. Or, if you have Amiga LightWave, you can directly import 3DS objects and DXF into Layout.

Moving on to other Internet places, there is one spot that shouldn't go unmentioned: Disney's *Toy Story* site at www.toystory.com. Although not directly related to LightWave, this is revolutionary filmmaking that any LightWaver should pay attention to. (Note: Keep an eye on upcoming issues of *LIGHTWAVEPRO* for a tutorial on creating your very own animated toy!)

On to the LightWave mailing lists and newsgroups for the month. I'm seeing that many of you are very hungry for additional LightWave training, information and examples of other animators' work. If you keep looking, you'll find that many LightWave users are opening WWW sites.

It's always a good idea to keep bookmarks of LightWave user sites, as they can often link you to other cool 3D sites, objects, images, scene files, and plug-ins. Much of the talk this month on the mailing lists was about animating a character in LightWave, which was partly sparked by the release of *Toy Story*, and the Blue M&M commercials (look for the article in January 1996's *VTU*), created using LightWave. Keep your eyes on *LIGHTWAVEPRO* for more on character animation techniques.

As I mentioned in the last column, it's great to see so many of you back on track and posting your LightWave questions, tips, tricks and comments, rather than complaints about software and systems.

Once again, the address for the LightWave list has changed locations. To subscribe, send an e-mail to listserve@garcia.com and indicate "subscribe LightWave." Don't forget to include your return e-mail address.

I've had numerous people ask me about surfing the net with their Amiga at the recent Avid training seminars, and through phone conversations and meetings. Granted, most WWW services center around the PC, but there is hope for the Amiga users. I-Net 225 is a professional, SANA-II compliant, TCP/IP networking software package for the Amiga. TCP/IP is an internationally accepted network protocol that exists on virtually all computer platforms, giving different systems a common means by which to exchange data. Finally the Amiga can communicate with everyone else (i.e., Mac, PC, SGI and Windows NT)!

Utilities that were written for AS225R2 should be fully compatible with I-Net 225. There are quite a number of these utilities available on BBSs and ftp sites around the world. I-Net 225 comes with a full host of clients and servers, allowing the Amiga to not only connect with remote hosts, but also to host other systems on the network. I-Net 225 is SANA-II–compliant, allowing it to operate concurrently with any other SANA-II–compliant networking package, like ENLAN-DFS. For more information, contact:

Interworks Professional Network Solutions
43191 Camino Casillas
Temecula, CA 92592-3714
Phone: (909) 699-8120
Fax: (909) 699-8279

I can't vouch for the software at this time, though I've read a few positive comments about it. Even though I use my Amiga every day, I don't use if for e-mail or Internet services, not because it's not good for such things, but because I have my PC already linked with full color.

If you do get connected with your Amiga, you'll soon find out what all the talk is about. Between E-mail lists, newsgroups and cool WWW sites, I probably spend close to two hours a day on-line. It's almost as bad as being in Vegas. Remember to e-mail me with your web sites, questions and comments. This column is for you, so keep in touch.

And, as always, a few cool sites to link into when you get a chance:

http://zebu.uoregon.edu/galaxy.html
http://www.odci.gov/cia/publications/95fact/index.html
http://www.acs.ryerson.ca/~aonsen/
http://www.electriciti.com/~dwhite/

See you next month!

LWP

Dan Ablan is president of AGA Digital Studios in Chicago, Ill., using LightWave for animation and Internet graphics. Reach him at dma@mcs.com. Also, check out his latest web page at http://www.mcs.com/~dma/home.html.

Editor's Message

continued from page 3

be surprised to see something before you expect it (just as long as you don't expect it now!).

Speaking of Windows 95, I had the opportunity to perform the *LIGHTWAVEPRO* benchmarks under this platform. Here's the results for a 100 MHz Pentium with 32 MB RAM and a comparison to the same machine running Windows for Workgroups 3.11 and Windows NT:

Scene	Time (Win 95)	Time (WFW 3.11)	Time (Win NT)
DOF.lws	9m 35s (575s)	10m 47s (647s)	9m 17s (557s)
Raytrace.lws	1h16m 15s (4575s)	2h 3m 53s (7433s)	1h 14m 18s (4458s)
Textures.lws	3m 44s (224s)	3m 57s (237s)	2m 43s (163s)
ZBufSort.lws	9m 37s (577s)	9m 6s (546s)	6m 32s (392s)

So what does this tell us? Some scenes are much faster under NT (Textures and ZBufSort), but the others are pretty close. With the final patched version, we should see times that are faster straight across the board, bringing them much closer to each other. That's good news for Windows 95 users, but I still highly recommend getting Windows NT if you use LightWave professionally. The multitasking and memory handling is much better under NT, and if you do crash, you won't take your whole system down!

I predict that 1996 is going to be the year of the plug-in. I feel that we're going to see an explosion in the number of commercial plug-ins available for Layout and Modeler. The few that I've seen so far have been pretty impressive and I know there're many more cool plug-ins in the works.

What's in store for *LIGHTWAVEPRO* this coming year? Hopefully, you'll be seeing many new changes. We plan on having new features and departments in *LIGHTWAVEPRO*, maintaining the quality you've come to expect.

John Gross
Editor

LWP

LightWave 101:
Give Me a Hand!

by Adam Chrystie

In the October issue, we made a simple mountain using splines. This time we will be making a hand splint using splines and a bit of metaform. Please save your work often and use multiple file names. Never delete the spline cages once they are done—you can always improve this object once you have finished the tutorial.

The first step is to analyze our intended object. First, we must decide how we are going to model with our splines. Along which axis will we begin modeling: X, Y or Z? Another important decision is to determine how many key cross section splines you will need for the object. Key cross section splines are used to add detail and more pronounced curvatures in the object. I compare them to keyframes in the time line of our animations. As we build our object we will discover the need to create additional key cross sections.

Starting Out

The hand splint starts at the forearm and extends until the initial joints of the fingers. Another portion of the splint extends further to cradle the thumb. I know it's hard to visualize the object, so take a look at the screen shots to get an idea of what you're building. I had five key cross section splines that were modeled along the object's Z-axis. The initial curve consisted of three points. I decided to plot the curves first and then go back and move the points to make a more defined shape. See Figure 1 to get an idea of the shape I used for the beginning of the arm splint. Follow these steps to begin the splint:

- Create the curve in Figure 1.
- Select the **Copy** (c) button to store a copy of the curve in LightWave's memory.
- Go to layer 2 and use the **Paste** (v) button.
- Move the curve in layer 2 to the next key spline cross section (Figure 2).
- Continue to use this method of copying and pasting curves until you have something that looks like Figure 3. When you do, cut all the curves and paste them back into layer 1. *[Editor's note: A quick way to achieve this is to make all of the layers active and choose Cut, then Paste. All curves will be pasted into the first active layer.—JG]*

Now we need to make the other defining boundaries of our arm splint object. We must make curves for the right and left sides and the bottom of the object.

- Select all left points of the curves starting with the first curve and ending with the last curve. Do not skip a point. Select them in order! If you make a mistake, push the (/) key or click in a gray area on the bottom of the Modeler screen to clear all of the points.
- Push Ctrl–p or **Make** in the **Tools** panel to build a curve using these points.
- Do the same for all the right-hand points.
- Select all the center points of the curves in order and make a curve using these points.

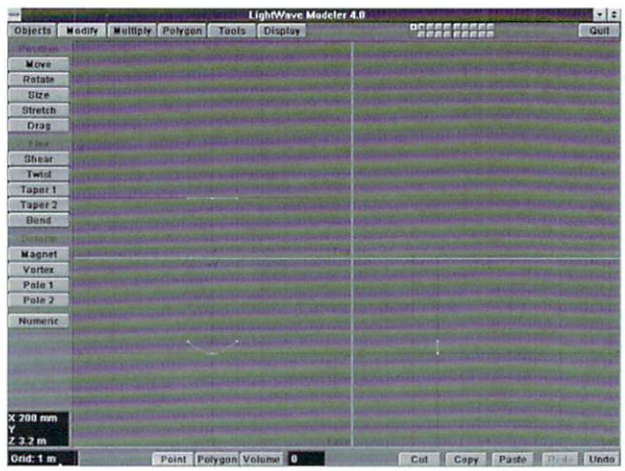

Figure 1

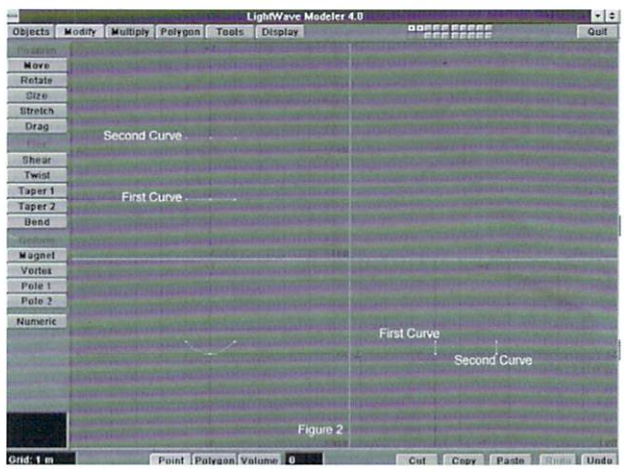

Figure 2

Adding Detail

Your screen should look similar to Figure 4. You might be noticing that this hardly looks like a hand splint! How are we going to define the hand area with only three points? By adding more points to the end of the splint object. More points will allow the mobility needed to shape the thumb region and the hand region. Let's get to it!

- Select the last curve.
- Select **Add Point** (**Polygon** panel) and add points to the end polygon. My splint had a total of six points. Try to add the points equidistant from each other (Figure 5).
- **Move** or **Drag** (**Modify** panel) each individual point so it mimics Figure 5.

Great! We're getting there! What would happen if you tried to patch the curves right now? Go on, try it. After patching, you'll notice that the splint looks like a smooth, curled piece of paper. It doesn't have the dramatic depressions that form a valley for the thumb to sit in. Overall, it is a very boring object. Let's take care of the thumb detail. We'll need to make a curve that will travel down the middle of the thumb. This curve will later be moved negatively along the Y-axis

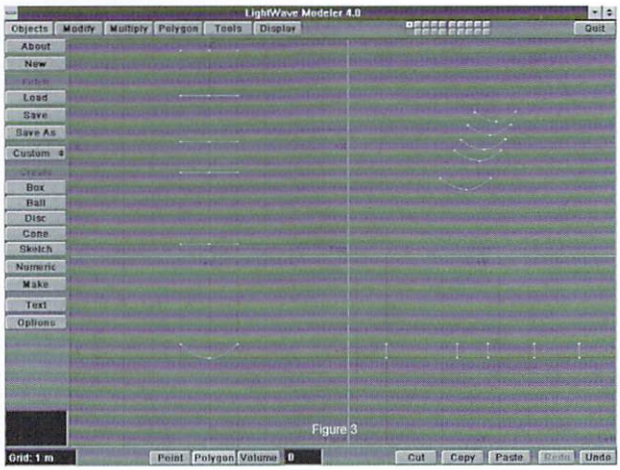

Figure 3

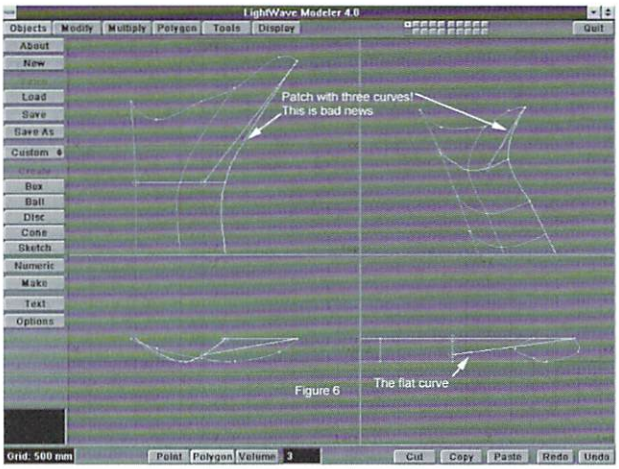

Figure 6

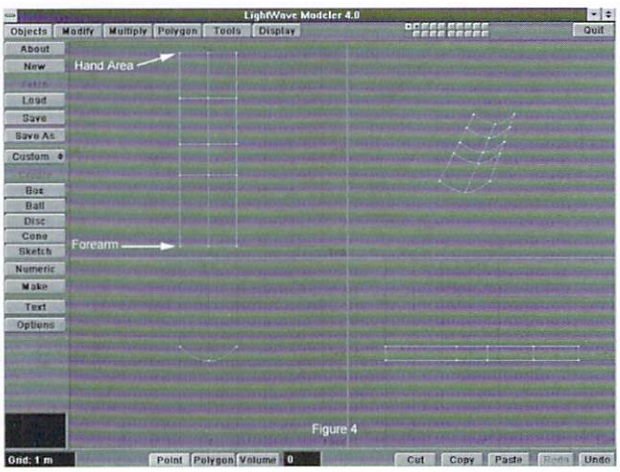

Figure 4

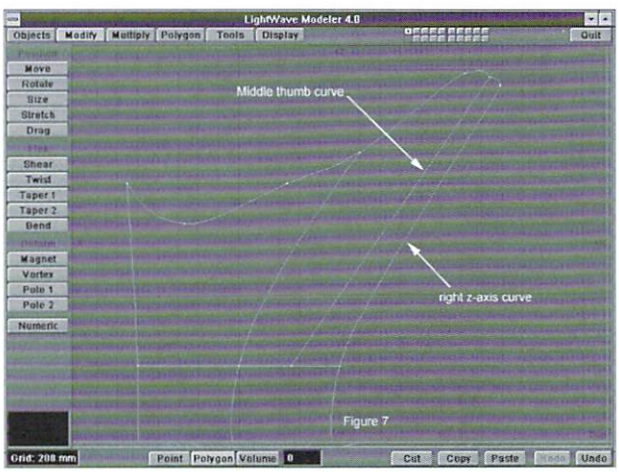

Figure 7

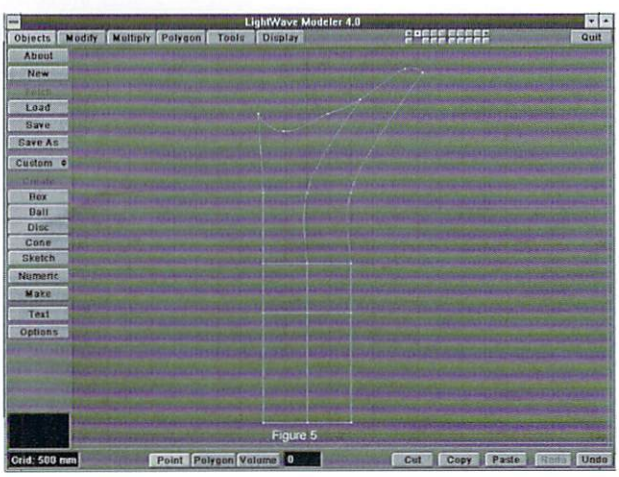

Figure 5

to create the needed depression in our object.

- Select the second-to-last cross-section curve.
- Use **Add Point** (**Polygon** panel) and add a new point between the center and right point of the curve.
- Select the point created in the above step and select the rightmost point of the last cross section curve. Make a curve out of these points (see Figure 6).

Problems and Solutions!

We've run into two problems with our model. Look in the Left view of Modeler. Notice that our new spline is a straight line. We will need to add another point somewhere along the curve. The other problem we have is that we've altered our model so there's a three-sided patch to the right of our "straight curve." See Figure 6 for details. Let's take the straight spline first:

- Switch to **Point** select mode (space bar) to make sure that no points are selected. Deselect any selected points.
- Switch to **Polygon** select mode and select the rightmost curve (the curve that runs down the entire Z-axis of the splint). If you selected a small curve, you've selected the wrong curve.
- Use Add Points (**Polygon** panel) to add a new point just below the point that was used for the creation of the previous straight spline.
- **Move** (**Modify**) the point above our new point so that it forms a smoother outline resembling a thumb. Don't worry if the curve isn't perfectly smooth. We will fix this later.

Your model should be similar to Figure 7. We have now added new points and rearranged the positions of our points while also retaining almost the original shape. We still have not converted the three-curved patch to a four-curved patch. Look at your monitor. You'll see that the curve in the middle of the thumb and the outer right-hand curve form two sides of a possible patch box. Check out the tip of the thumb and the second-to-last spline cross section. These two curves form the forward and rear sections of the patch to make a four-curved patch, which is what we want. We need to split the curve at the thumb's tip from the right-hand curve and to split the tip of the left side of the thumb from the last cross section. Here's how!

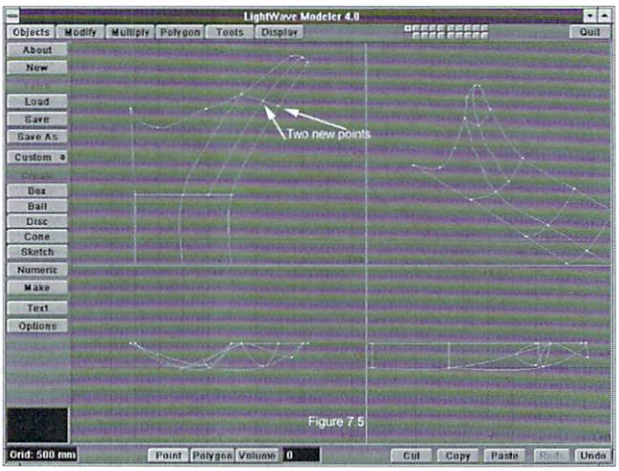

Figure 8

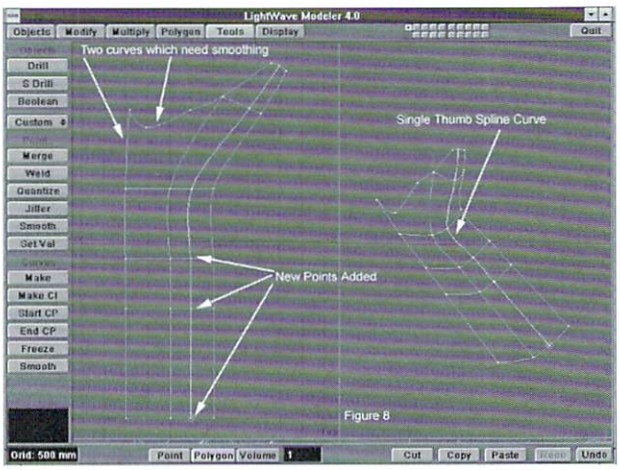

Figure 9

- Select the right Z-axis curve.
- Switch to **Point** select mode.
- Select the second-to-last point of the right-hand curve.
- Use the **Split** (**Polygon**) command to split the right-hand curve into two separate curves. This will form another small curve at the top of the thumb region.
- Switch to **Polygon** select mode. The two curves should be selected. Select the Right-hand curve and the new thumb tip curve if they are not already.
- Use **Smooth** curves (**Tools**) to make the two curves blend into each other. Be sure to use the smooth command that is in the **Curves** area of the Tools panel, not the **Points** area!

You need to perform the same steps to the left side of the thumb. The steps are exactly the same except that they are performed on the opposite side of the thumb.

Now it's time to fix the curve that will control the depth for the thumb region. Remember, this "straight curve" has only two points. After adding the middle point to the straight curve, we'll add another point to the right Z-axis curve. The new points will be used to create another cross section curve in the thumb region. (Later, you can alter the positions of these points to give the thumb more definition.) Look at Figure 8 and create a curve that is similar to mine or follow the steps outlined below:

- Change to **Point** select mode. Deselect any points.
- Change to **Polygon** select mode. (Remember, the space bar changes modes.)
- Select the middle thumb curve and right Z-axis curve.
- Use **Add Point** (**Polygon** panel) to create a new point on the middle thumb curve and right Z-axis curve. Look at Figure 8 for the point placements.
- Switch to **Point** select mode and deselect all points.
- Reselect the points of our soon-to-be curve in order from left to right or right to left.
- Make the curve (Ctrl–p).
- **Move** (**Modify**) the point's Y-axis positions. The middle point of the new cross section curve should be lower than the two outer points.

Take a look at the new detail that we've created in the thumb region. Notice how most of the patches neatly form quadrangles that have two points of a different patch in common. Almost all our splines neatly match each other in this fashion (or at least they should).

Can you find the area that does not conform to the above rule? Follow our thumb curve along the Z-axis. Somewhere the geometry should cease to form perfectly aligned quadrangle spline sections. If this is not corrected, heinous, model-threatening rendering errors will occur. Assume we use patch values of 5 for perpendicular subdivisions and 5 for parallel subdivisions without fixing the alignment error. You can see in Figure 8's wireframe preview window, that we already have two small patch regions that lie directly above one large patch region. The junction between these two curves is where the error will be produced.

Notice how the two smaller regions fit perfectly along the width of the larger patch section? As a result, after we patch the large section, we will have five polygons along the brim adjacent to the two smaller patch sections. But, after patching the two smaller sections we'll have a total of 10 polygons, which will be adjacent to the five polygons of the larger patch section. This really appears ugly on Modeler's screen, and believe me, it looks nasty when rendered in Layout. Since the points are shared, you'll get ghastly seams.

Well, because you're smart, you'll think to yourself: "What if I play with the subdivision settings to make the large patch section have 10 polygons? Won't that fix the problem?" Yes, it will, assuming that you select your curves in the correct order and create a number of faces in each small patch area that equals the number of faces along the adjoining edge of the large patch area. (See sidebar.)

There are two other solutions that will rectify the problem. A more advanced approach "stitches" the two dissimilar path sections together. (See Ken Stranahan's article in *LIGHTWAVEPRO* Volume 1, Issue 2, December 1993.) The second, easier method calls for extending the thumb curve down the entire Z-axis by adding new points to the preceding splines and then creating a new curve using these points. We will use this technique here. Finally, we will need to **Merge** (**Polygon** panel) the two Z-axis thumb splines together. After doing this you should have something that looks like Figure 9.

We're almost done using splines. Hang in there! There are a few modifications we should make before patching our model. Look at Figure 9. Can you find two curves that should be smoothed? Look at the left Z-axis curve and follow it down its length until you come to the last point. This point should also be the first point of our last cross section curve. Select both of these curves and use the **Smooth** curves (**Tools**) function on them.

Now we need to move some points around so our object has some curvatures along the Y direction. Basically, the perimeter points of the entire model should have larger Y values than the inner points. Check the points that were added to create the detail splines. Make sure these points are positioned correctly on the Y-axis. You might have to lower Y values of some inner points in the thumb region. The inner points should be positioned so they maintain the arch profile that we began our model with. Finally, look at the points defining the last cross-section curve in the model. The perimeter points of the thumb region initially should have Y values that match the other perimeter point's Y values. As we travel towards the left side of the cast, try to decrease the Y values gradually, forming a slight depression. The leftmost and rightmost points should have a Y value that is near the Y values of the other perimeter points.

Once you get things looking smooth, **Save** the curves as a separate object!

Patch Time

Yes, it is now time to use the lovely **Patch** (**Multiply** or Ctrl–f) command! Let's begin by analyzing the right-hand splines. To make life simple, I used

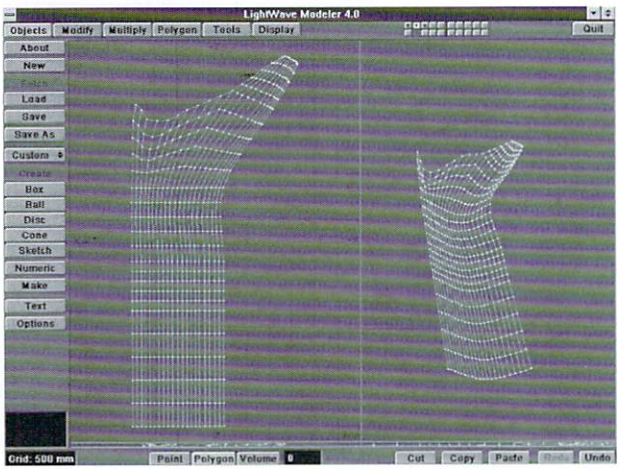

Figure 10

a value of 5 in each of the Patch requesters. I also set the Patch method to **Length**. All of these items are set after you select four connecting curves and issue the **Patch** command. I used values of 5 for two reasons: If we use equal patch values in conjunction with four-sided patch regions, then we do not need to concern ourselves with the order of curve selection since each patch is four-sided and equal in polygon resolution. (See October's introductory spline article, "LightWave 101.") Also, a polygon value of 5 gives us enough definition because there are many patch areas on the right side that do not span a great distance.

Do you notice a difference between the splines of the right side and the splines of the left side? There are more spline patch regions on the right side of the model than on the left side. So? Well, when you try to patch one of the left side regions, the resulting polygon mesh will not be permitted to follow the curve profiles accurately because the area is just too large to be well-defined by only five sets of polygons. Compare the patch regions of each side. How do their patch densities differ? Do they differ along the X-axis or along the Z-axis? In my opinion, they have contrasting X-axis values. The right side of the model's X-axis has two patch regions while the left side only has one patch region along the X-axis.

There are two simple solutions. We could add another long spline Z-axis spline to the left side, or we could alter the mesh's resolution when we issue the **Patch** (**Multiply**) command. I chose to increase the mesh's X polygon count by playing with the values of the Patch requester. Now is a good time to review the laws by which LightWave interacts with the parallel and perpendicular values in conjunction with curve selection order. Basically, I increased one field to 10 and kept the other at 5. Use the curve selection order that produces 10 polygons along the X-axis and five polygons along the Z-axis. Patch all the regions if you have not already.

Clean the Messy Mesh

Phew, are your hands tired! What I would give for an autopatcher plug-in! Let's clean up our messy mesh.

- Use the **Copy** button and **Paste** the mesh into layer 2.
- Make layer 2 active and make sure you're in **Polygon** select mode.
- Push (w) to bring up the polygon statistics screen and select all curves.
- Push (z) to erase the curves from the model.
- Use **Merge** points (**Tools** panel or m).
- Select **Align** (**Polygon**) to force all the polygon normals to point in the same direction. Make sure they are facing the way you want.

Oh So Shallow

The mesh has now been cleansed. Figure 10 displays my version of the object.

- **Copy** (c) and **Paste** (v) the mesh into layer 3.

A two-dimensional model is like dating a shallow person. We need to create some depth!

- **Extrude** (**Multiply**) the object slightly using one segment. I extruded the object so that you can see a noticeable depth when it's viewed in Modeler. Don't overdo it. Hey, how about some rounded edges?
- Click on the **Subdivide** (**Polygon** panel or Shift–d) button and select **Metaform** from the subpanel. I metaformed my object twice.
- The final step is to seek and destroy non-planar polygons by using **Polygon** Stats (w) and selecting the (+) button next to **Non-planar** polygons. Triple (t) these polygons.

Hey, it's time to save your object and render it out in Layout! While in Layout, you might have to boost the **Smoothing Angle** (**Surfaces** panel) to a higher value. I used 114 degrees.

Don't limit your splining experiments to this project. Alter these spline cages, patch them, and make morph targets for character animations. After metaforming, you might want to drill some holes in the object by using **Boolean** Subtract (Shift–b).

Have fun splining until we meet again.

LWP

Adam hopes that those who push for peace in the Middle East won't succumb to the forces of extremism, terrorism and political games. Peacemakers are the bravest of warriors, for they often face invisible enemies. Adam can be reached at adamcbry@cats.ucsc.edu.

Movie Etiquette

How Not to Make Enemies on Set

by Tom Williamson

In the fast-paced world of moviemaking there are unwritten rules. Rules that need to be learned and heeded if you want to go anywhere in this business. With the growing success of LightWave 3D, there are more and more feature films and television shows are showcasing LightWave's tremendous power and flexibility. Facilities like Foundation Imaging and Amblin Imaging have opened the floodgates, and small graphics companies are springing up all over the place (mine included). When work began on our latest feature, several of Computer Café's employees joined me on set, most for the first time. As the evening wore on, I found myself answering a myriad of questions on terminology and etiquette, based on my 10 years in this business. I then decided to write this article.

Be Professional

This is the most important rule. When you consider the delicate balance of multiple departments, tons of equipment, actors, locations, personalities, technical problems and script revisions, it's amazing that movies get made at all! There is also a huge dollar amount connected to each single element that makes up a film crew. Filmmaking is incredibly expensive. Probably the one thing that makes it all work is specialization and professionalism. When a good crew is assembled, from the director to the craft services (the people who provide the snacks and drinks), it's amazing. The whole thing works like a well-oiled machine. That's why it's very important to stay professional. Your actions on set are a direct reflection of your company (and, in some cases, your field). I'm not saying you have to be a stuffed shirt, but when it's time to work put on your serious face. With over 100 people on the average film crew, you will meet an assortment of personalities, and this industry breeds eccentrics. But even that flamboyant, purple-haired makeup artist knows that when "final touches" (the call to make all last-minute adjustments to hair, makeup and wardrobe) is called, it's time to stop doing Ethel Merman impressions and get to work. Every person there is a highly skilled craftsman. They deserve your respect and courtesy.

Stay out of the Way

Depending on where you are shooting, things can get very cramped on set. When all the pieces come together for a shot there usually isn't much room left. There's the cameras, dollies, tracks, cables, lights, lightstands, c-stands, flags and a plethora of other equipment. Then there's the director, the first AD (assistant director), the second AD, the second second (another second AD, really!), the DP (director of photography), the camera operator (sometimes the DP), first AC (assistant camera), second AC and third AC. And don't forget the key grip, the grips (the people who handle all the equipment it takes to make a movie), electricians, the best boy (an electrician in training, traditionally nicknamed "Sparky"), sound, boom, hair, makeup, special effects, actors, stuntmen and a few dozen other people. When the shot's getting close, it's very important to stand clear. It's a fair bet that if someone's running in yelling "hot points ("I'm carrying something pointed"), watch your back," they're doing something more important than you at that particular moment. It really seems to come down to that: all the jobs are important, just not all at once. When there's lighting going on, the gaffer (the guy in charge of getting lights to the places the DP wants) rules the set. When there's an elaborate or dangerous special effect coming up, the FX guys rule the set. It's the same with most of the departments, and everyone else knows to get out of the way and stay quiet to let the department of the moment work. By the way, no matter who "rules" the set, the director is always top dog.

Keep Your Eyes and Ears Open

I would say that 90 percent of filmmaking is waiting. There is usually an incredible amount of "down time" to kill. When you're on a set, you'll notice a lot of people just standing around, including yourself. You'll also see a lot of walkmans, gameboys, books and other various ways of passing time. The busiest people are the camera crew and grips—they seem to always have something to do. Most people on set will be needed at one time or another, and the last thing you want to do is make everybody wait while they look for you. It's important to stay on top of what's going on. Most key people carry walkie-talkies (or just "walkies"), and by keeping your ears open you can keep pretty informed on what's happening. When you hear the words "picture's up," it means shooting is about to begin, and if you're needed, you had better be there. Other key phrases: "Rolling"—cameras are rolling (usually preceded by "roll sound"); "Cut"—cameras have stopped rolling, "Lock it up"—don't move; and "Fire in the hole!"—a loud sound is coming. And about a dozen other phrases get screamed and then echoed by the ADs and the PAs (production assistants). When you're on a sound stage it's important to listen for the bell, which will ring to signal a shot is about to happen. If you happen to be outside the stage, you'll see a "whoopee light" spinning around above the door. Don't go in! Even though most sound stages have very quiet double doors, you do not want to be the reason a shot was botched. If you just pay attention—as hard as it gets when you've been sitting for two hours—you'll stay out of trouble.

Show Some Respect

This shouldn't be too hard to follow. The thing to remember is that what you're doing is a privilege. If you're working on a movie set it's probably because you've worked your ass off and paid your dues. (Or you knew somebody—funny, eh?) The thing to remember is that everyone else there did, too. They wouldn't be there if they didn't want to be. It's very hard work, with long days and incredible deadlines. Most people on the set are there because they have a real love for the business, and they did what they had to do to get there. Treat everyone with respect, from the producer to "transpo" (the guys who drive the equipment around and shuttle people). When you're watching the crew in action you'll probably hear a lot of "Yes, sir" and "No, sir," and "please" and "thank you." These people are not kissing butt. They genuinely respect one another.

Stay Friendly

When people of different professions come together there is usually a lot of curiosity about the work

other people do. I try to be as open and informative as I can. This really helps me when it comes time for me to rule the set. When I come running in yelling "Nobody move, I need a clean plate" (a shot with no actors in it, usually for rub-throughs or split-screens), everyone pretty much knows what I mean and gives me what I want. If you walk around brooding like the silent artist type, no one will want to talk to you. They will probably be grumbling under their breath when you're telling them what to do later. Being friendly also helps you make connections for future projects, and gives you someone to talk with to pass the time. It's the nice people you remember.

Know the Hierarchy

No one likes to have someone go over their head, and it's important to know the chain of command to avoid this. For our business, visual effects, there are several people we need to keep in communication with. The main person is usually the effects supervisor for the production. He or she is responsible for all the shops doing work on the film. Most questions and comments should be directed to the effects supervisor, who then can pass them to the proper person. If a dialog is opened by the effects supervisor and another department head for you, you may then communicate directly. Usually the effects supervisor is pretty busy and will make a point of introducing you to various department heads, allowing you to communicate freely. The last thing you want to do is go to the DP and say "It's way too bright" before you've been introduced. It's just not cool. Some effects supervisors like to keep a thumb on all involved shops; fortunately, they're the exception.

If you are going to talk about one of your shots with someone you don't know—the camera operator, for instance—be sure he or she is not busy and that you introduce yourself. If you have a specific problem, you should go to the department that can help. For lighting problems, talk with the DP or gaffer. If you need a "stinger" (an extension cord), talk with electric. If you need an "apple box" (a wooden all-purpose box, available in full, half, quarter and "pancake" sizes), talk to the grips. All creative decisions should be run by the effects supervisor or the director.

Usually, a film crew becomes very close very quickly. It's almost like a family. If you can, try to be there for the first couple of days while this "bonding" is happening, so you won't feel like such an outsider. If you only show up in the last week of production, or only for the shots for which you're needed, expect to feel like a red-headed stepchild. Do not confuse this with loitering. If you're not needed, you shouldn't be there, and you're probably not welcome. And you should always ask before bringing a visitor, especially on a bigger production. Secrecy is a very real and favored practice in this industry. Finally, try not to bug the director, who's got enough headaches. Try the second AD, then the first AD, before going to the head honcho.

This is by no means a complete list, but it covers the basics. Just using common sense will get you through almost any situation that should arise. If a problem should crop up, nip it in the bud. The quicker you smooth things over, the better. Keep these little tips in mind on your first few productions and you'll step on fewer toes. Remember, you're only as good as your last job.

LWP

Tom Williamson is the effects supervisor for Santa Maria, Calif.–based Computer Café. He's an avid fan of special effects, GI Joe and semprini. He also makes a mean chicken cordon bleu.

Tip of the Month:

How is the Glow Effect calculated?

The Glow Effect found in LightWave 4.0 is calculated based upon the color of the surface with **Glow Effect** applied and the amount of glow area to non-glow area within the **Glow Radius**. Every pixel in a surface that is using Glow Effect will contain an area of glow surrounding it. It helps to imagine (or render) the following circumstance:

Picture a pixel on the edge of a plane that has **Glow Effect** enabled. Assume that the **Glow Intensity** is set to 100% and **Glow Radius** is set to 10 pixels. Figure 1 shows the Glow Radius surrounding the pixel. Notice that there is an equal amount of glow to non-glow area within the Glow Radius. In this case, a pixel located in the non-glow area, but directly adjacent to the glowing pixel, will be colored a value equal to 100 percent of the pixel color. Additionally, it will fade to nothing 10 pixels away from the glowing pixel.

In the case where there is an unequal amount of glow to non-glow area (Figure 2), the glow color will take on an appropriate value equal to the amount of non-glow area. In the case shown in Figure 2, because there is 50 percent more of the non-glow area within the **Glow Radius**, the color of a non-glow pixel adjacent to a glowing pixel is reduced by that amount. The opposite is true, of course, as in the case of an inside pixel at the point of the letter V.

—John Gross

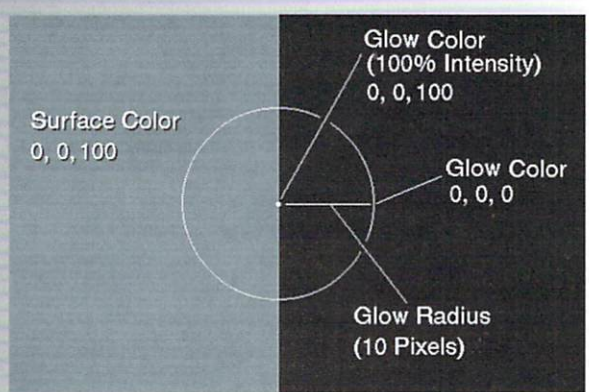

Figure 1

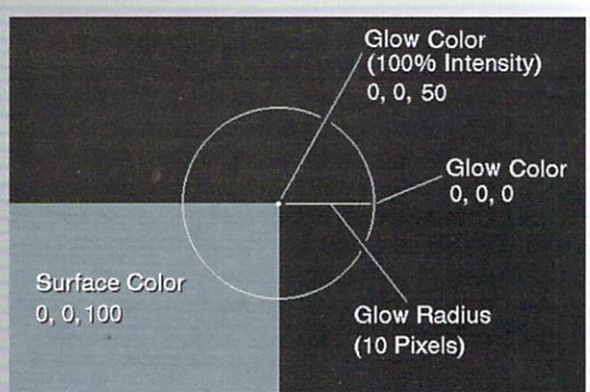

Figure 2

Cleaning House

Performance-Increasing Tips for Windows NT and LightWave

by Glyn Williams

Picture this: There are just seven weeks till the deadline. In that time, you have about 20 shots to storyboard, models to produce and texture, and scenes to build, animate and render. Losing sleep? Working long, long hours? Is this your life?

If so, you may have opted to save time and run LightWave on a more powerful system—and that probably means an Intel- or an Alpha-based platform. (The term "Intel platform" is a nice way of saying "PC-compatible." PC seems to send a shiver down many a spine.)

LightWave for Intel is nice enough to allow you the choice of Windows 3.1, Windows 95 or Windows NT. While Windows 3.1 is inexpensive and Windows 95 offers a bunch of new features, NT is by far the best choice for LightWave users because it is simply much faster at throwing around the vast lumps of memory that LightWave requires.

If you are using an Alpha-based machine, NT 3.51 is your only option. NT is not a version of Windows, but a completely new, 32-bit operating system produced by Microsoft to capture the workstation market. NT is designed with features like true preemptive multitasking and broad, uncluttered access to a real 32-bit memory space. And yes, while AmigaDOS did this happily on a 512K Amiga, NT finds it a bit of a squeeze on a 16 MB Pentium. To provide backward compatibility on top of the NT kernel, the programmers added a look-alike GUI and a facility to run old 16-bit Windows applications.

It's time for a simile, I think. If AmigaDOS was a rodent—fast, agile and light on its feet—then NT would be a huge, rumbling battle tank: massive, complex, unstoppable and pretty heavy on fuel. Because NT is so complex, you'll want to squeeze a little extra performance out of it using a range of methods. That added boost can be a real boon—a three-minute-per-frame speed-up over 1,000 frames can make the difference between meeting a deadline and missing it. Here's a collection of tips for improving the performance of LightWave under Windows NT.

Find the Bottleneck

There's always one somewhere. With a program like LightWave, the bottleneck is either going to be the processor or memory. Examining your system while running a rendering task is helpful. If the disk light is always on, the hard disk is thrashing, the system monitor says there are hundreds of page faults per second and the free-memory indicator says there is less than 3 MB of free RAM in the system, then chances are the bottleneck is memory.

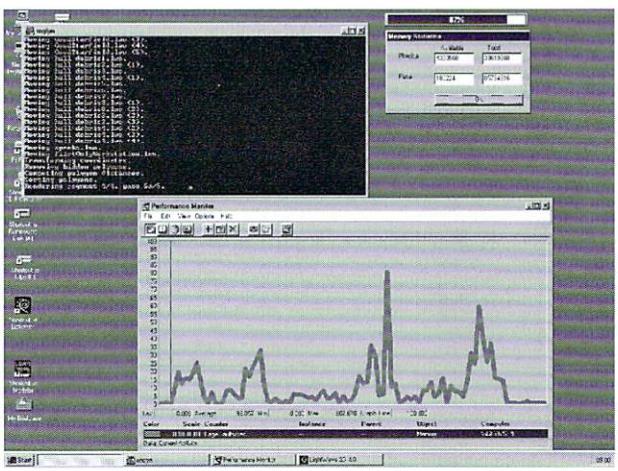

Figure 1: Here, with the Performance Monitor tool, the memory/page-faults-per-second index shows a rendering session very low on memory.

When NT runs out of real memory, it starts to use virtual memory (a swapfile on the hard disk). Though a little paging to disk is acceptable, if more memory is required, the system will start to thrash, the machine will access the hard disk constantly and performance will plummet. Remember, a silent PC is a happy PC. Or, if you prefer, the drive light on the front of the PC is a reminder to get in your car and drive to the computer store and buy some more RAM chips.

If the bottleneck is not RAM, the only way to improve performance is to get a faster processor. LightWave is written in 32-bit code and makes heavy use of the math coprocessor, meaning it really benefits from the architecture of the Pentium processor. A P90 will run LightWave about four times faster than a 486 DX2-66. Better still, a DEC Alpha 21164-based machine may achieve six times the performance of a P90!

The overhead for NT itself is a modest 12 MB. If you are running LightWave on a 16 MB system, only the tiniest of models can be accommodated before the system has to start using virtual memory extensively. For most professional work, 32 MB is really the required minimum.

Reduce LightWave's Memory Usage

LightWave and similar programs eat RAM for breakfast. The software itself uses amazingly little memory, but loaded objects and images consume a great deal. The good news is that you can reduce the memory consumption of a scene by taking the following steps:

Reducing the Segment Memory Setting

- LightWave uses a big chunk of memory to buffer the image being rendered. For large images, this can be a significant amount of RAM. The **Segment Memory** setting (**Camera** panel) allows you to specify how much RAM LightWave allocates to this buffer. A typical NTSC video image requires a Segment Memory setting of approximately 8.7 MB to hold the entire image. By reducing this setting to 4.5 MB, LightWave will render the image in two half-screen segments—a little slower than usual, but now you'll have an extra 4.5 MB of RAM at your disposal. If the system was desperately out of RAM, this extra 4.5 MB would boost performance. With a bit of experimentation, you should be able to balance these conflicting requirements. *[Editor's note: To determine the size of the image buffer needed (in bytes), multiply your output width by output height by 24.—JG]*

Cutting Out Unwanted Geometry

- Objects that are never seen still use precious RAM, so cut 'em out or cut 'em down. Consider hidden or obscured features—objects behind the camera—or complex objects in the far distance, where a simpler object would do.

Reducing Image Memory Usage

- Image files use a lot of memory. LightWave 4.0 shipped with HIIP image converters, allowing the program to load JPEGs and other file formats. Unfortunately, HIIP loaders pad out images to 24-bit, regardless of the number of bitplanes they contain. I suggest you avoid using these loaders completely. To do this, rename the IFF.EXE and TGA.EXE files to OldIFF.EXE and OldTGA.EXE. (LightWave's internal loaders will then take over, which won't pad out the images to 24-bit for smaller bitplanes.) LightWave has internal loaders for IFF and 24-bit TGA files. *[Editor's note: NewTek is aware of various HIIP problems and is working on a patch to resolve them.—JG]*
- Use 8-bit IFF images wherever possible. Whenever you are using an image map solely for a surface value other than Surface Color, you can use an 8-bit image because only the luminance value is taken into account. This can save megabytes for modest-size maps. And don't use a 2K-pixel-by-2K-pixel map when a 1K-pixel-by-1K-pixel map will do. Make a habit of browsing the image list and looking at memory usage. You may be surprised. If using a clip map, you can get by with two-color—black-and-white—images, as only two values are used for clip maps.

Maximizing NT's Available Memory

There are a number of tricks you can try:

- Close down any programs not in use—don't have Photoshop running in the background. Don't use wallpaper, NT's OpenGL 3D screensavers, or any screensavers for that matter. (Even NT's clever circular clock is a waste of perfectly good RAM.)
- Disable unwanted services. For example, if you don't have a printer attached to your machine, disable the spooler service from the control panel, which will save more than half a MB of real memory. You'll need to log on as an administrator to do this.
- I've heard rumors of other ways to reduce NT's voracious RAM requirements: Stop the automatic loading of the WOW subsystem (which involves altering the registry), allowing NT to run old 16-bit Windows apps. Or rename OS2.EXE, which apparently allows character mode programs for the OS/2 operating system to operate. (I cannot say whether there are any real benefits to these procedures.)
- Disable BIOS shadowing features. NT would rather have the RAM. (I suspect this will only liberate a few K.)

I'm not going to go into depth about how to perform these procedures. If you are interested, there are a number of ways of finding out how, including NT manuals, Internet newsgroups and Microsoft.

Improving Virtual Memory Performance

NT uses virtual memory, to some extent, all the time. When RAM is used up, the system swaps pages of real memory to disk into a swapfile called PAGEFILE.SYS. In a memory-bound application, improving NT's usage of virtual memory can make quite a difference.

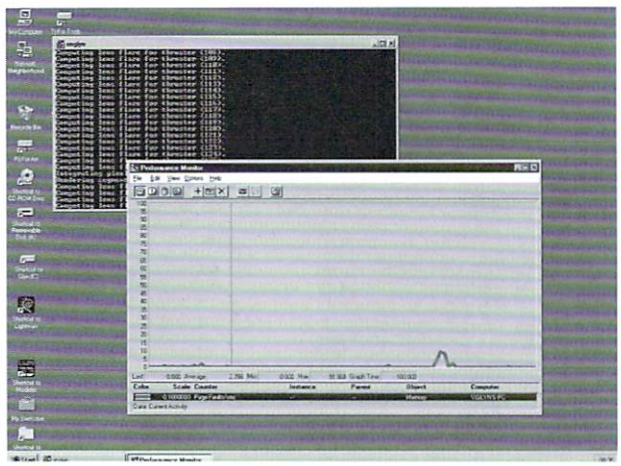

Figure 2: In this rendering session, memory consumption has been reduced, resulting in less paging and faster rendering times.

- Using the Control panel/system, set the minimum amount of swapfile so it's at least your actual amount of RAM+12 MB. You'll need to be logged on as an administrator to do this. Having to grow the swapfile amount past its initial setting slows down NT a lot and you won't need to do this if you have increased the amount of RAM in your system since installing NT.
- If you have multiple disk drives, place swapfiles on each of them. (Note: Never have a swapfile on a compressed area.)
- If you are using NTFS and are feeling adventurous, use striped volumes and put your swapfile on them. Striped volumes share data across more than one disk, so if each of two disk drives works at 4 MB per second, you can read and write data to a striped volume at 8 MB per second. Be smart about backing up if you are choosing this method: If you lose one disk, you'll lose all the data!
- Don't fragment swapfiles all over the disk. If you can make a contiguous swapfile, you'll be helping yourself out. If you are using a FAT disk and have access to Windows 95, use this method: (1)Boot up Windows 95. (2) Delete Pagefile.sys (95 doesn't use it). (3) Run Defrag. Ignore what it says about fragmentation and get it to unify free space. (Don't use a DOS de-fragger—you'll lose your nice long filenames.) (4) Reboot NT. NT will remake a large, contiguous swapfile onto unused disk space.

Avoiding the Memory Leak

LightWave and NT allocate and free memory differently. When LightWave loads a big scene file, it uses vast tracts of memory (virtual and otherwise). When you press Clear Scene, this RAM is not recovered until you either exit LightWave or load another scene.

When doing a large rendering operation, even on a single machine, use LightWave in ScreamerNet mode. This renders significantly faster than using Layout itself to do the rendering. If you work like we do, the last thing you do at night is set up ScreamerNet with a nighttime batch of rendering work. If the machine running the ScreamerNet control panel is also going to be used for rendering, press Clear Scene and close down LightWave completely to free up the RAM. Run the ScreamerNet node (LWSN.EXE) and then run LightWave once more to get the ScreamerNet control panel. *[Editor's note: A bug in the Microsoft compiler used with LightWave and Modeler causes memory to be "tagged" as freed rather than actually freed. This situation can cause problems when loading images sequences and loading and clearing large scenes. The same patch that NewTek is working on for the HIIP problems will contain a workaround for these Microsoft problems.—JG]*

Finally, here are some cool features of NT you may not have heard about:

- The NewShell for NT is a replacement GUI for NT that is almost identical to Windows 95. This much-improved interface has several benefits for LightWave users, but beware: it does require more memory. You can download it from many on-line services or at ftp.microsoft.com.
- If you are using an NTFS volume, you can elect to employ file compression. NT will apply a fast, reliable and transparent compression to any directories you wish, which is particularly useful for volumes containing 24-bit movie sequences and large texture maps.
- If you do manage to get home, why stop working? With NT's remote-access server (RAS), you can log onto your office server machine exactly as if you were in the office, albeit with a pretty sluggish network link.

LWP

After a career of designing and programming games, Glyn Williams is now co-owner and director of Particle Systems, a startup games development house based in Sheffield, England. The company is making heavy use of LightWave in the production of its current title, due for release in 1996.

LightWave and Print Technology

by Mark Thompson

For just about every 3D animator, there comes a time when there is a desire or a need to get a print of an animation frame. A glance through the back pages of *VTU* reveals that there are a number of businesses equipped to provide this service. Send them your 752x480 LightWave-generated frames, and they will happily deliver photographic slides or prints without a breath of effort on your part. As a convenient and portable means to show off your work, these are quite adequate. However, upon critical inspection, you will note that the prints appear fuzzy, lacking the crispness you normally associate with photographs or magazine ads. Besides that, the colors frequently look washed out or possibly shifted in hue. While there are some bargains to be had in the area of hardcopy output—the Fargo Primera printers being a good example—for the most part, you get what you pay for. And when your company's image is on the line, it can be a grave mistake to show anything but your best in this highly image-motivated industry. But producing this level of quality requires a bit more effort on your part, and it helps tremendously to learn a few basics about the print trade before diving in.

There are two routes to take in producing high-quality hardcopy output: photo output or offset printing. From LightWave's perspective, they are essentially the same. The difference comes in on the back end when preparing the work for delivery to the service bureau. The photographic process can typically be more cost-effective for one or two stills, but it's completely impractical if you plan on producing a large number for advertising purposes. This is where the economics of printing really shines. Since this is likely the area of interest for most animation companies, it is where I will focus the most attention. Besides, the work required for photo output is basically a subset of that for print.

Before I continue, I must preface the following with the announcement that I am, admittedly, no expert in the printing business. The information here is what I have gleaned from going through this process myself many times, while trying to learn what is going on along the way. Once you have a few basics under your belt, you too can be walking and talking like a page layout professional, and hopefully make some darn nice hard copy. With that disclaimer aside, let's get to printing with LightWave.

You Say Tomato

Much of the confusion that exists between animators and the people in the print industry stems from differences in terminology, even though the fundamental processes involved are quite similar. In print, the resolution of an image is typically expressed in terms of what "line screen" is being used. Line screen is also known as lines per inch, or lpi. The other unit of measure is dots per inch, or dpi. Strange as it may seem, the dpi resolution of a print is completely independent of the lpi resolution. That's because the lpi describes the resolution of the image the printer is resolving, while the dpi is the physical dot density of the print device. So why would the two values be independent from one another? Because a combination of multiple dots is used to form each pixel of the image. And the density of dots used can vary across different print jobs. This deserves some further explanation.

Much like a CRT screen, print media produces its continuous range of colors by combining primary colors. A CRT uses red, green and blue light in what is known as the additive color process. The colors add together to create the spectrum, with full intensity on all three yielding white. In print, a subtractive color process is used. Cyan, magenta and yellow (the complements of red, green and blue) are subtracted from the white print medium. Both methods rely upon the human eye to integrate the individual dots of primary color into a smooth spectrum of millions of shades. On a CRT, a color is created by changing the relative intensity of three point light sources (RGB) typically grouped together in a triad configuration. However, in offset printing, there is no method for varying the intensity of each dot. Instead, very high dot densities are used, and the percentage of dots in a particular area determines the intensity (once again relying on the eye to integrate them). Intensity can also be modulated by varying the dot size. So while a CRT could conceivably only require three dots per image pixel, an offset print would use many more. Consider that the typical monitor has around 200 dots per inch (three dots for each RGB triad). Compare that to 2,500 dots per inch on a commercial printer. A Linotronic 330 printer can be driven to deliver from 635 to 3,386 dpi. And, though this number is independent of the desired lines per inch, higher dpi values obviously do a better job of resolving higher lpi images. The main thing to remember is that lpi is associated with the image itself while dpi is a function of the device reproducing the image. With that aside, let's get back to line screen and how it relates to your LightWave rendering resolution.

A Pixel By Any Other Name

Animators don't describe the resolution of their images in terms of lines per inch—they specify pixels. So, to bring things into perspective, the pixel per inch resolution of an image is equal to double the line screen. This makes sense when you consider that you need two pixels (one on and one off) to define a line. Graphics programs sometimes refer to image resolutions in terms of dots per inch, but what they really mean is pixels per inch, or ppi. A 300 ppi image (common for color print media today) has a line screen of 150. However, all these resolution descriptions have a "per inch" tacked onto the units. In order to relate your LightWave renderings to these resolutions, you must also associate a physical size to the image. Although it may be painfully obvious, ppi = pixels/inches. Therefore, a 752x480 LightWave rendering printed at 300 ppi would be 2.5 inches x 1.6 inches on a page. Clearly, LightWave's standard video resolution is pretty small by print standards. Of course you don't have to render at the full resolution set by the line screen. But if you don't, your image can suffer from jagged artifacts and/or fuzzy detail.

So let's consider a practical example. A standard magazine requires an image 8.375 inches x 10.875 inches. The magazine publisher informs you that he will be using a Linotronic 330 to print and that he needs you to deliver color negatives at 133 line screen and 2,540 dpi. 133 line screen used to be the most common resolution for color offset printing and is

still used a great deal for magazines today. This means you must render at 266 ppi for maximum image sharpness (remember, double the line screen). To accomplish this, choose **Custom Size** in LightWave's **Camera** panel and enter 2228x2893 (8.375 x 266 = 2227.75 and 10.875 x 266 = 2892.75). Also, whenever you are outputting for print, remember to use a square pixel aspect ratio. You will notice in the Camera panel that the number of segments required to render this image has probably jumped up to a rather large number. I typically set my **Segment Memory** to 8.67 MB so my video-resolution stuff will render in one segment. At this high output resolution, an 8.67 MB segment size results in 19 segments. For speed purposes, you want to render in as few segments as possible within the limits of physical memory. Note that under NT, virtual memory will allow you to specify an outrageously large segment memory size. However, when you vastly exceed the available physical memory, the speed penalty resulting from disk paging far outweighs the advantage of rendering in a single segment. So keep the segment size within physical limits. You can judge this by seeing if your drive is doing lots of thrashing for each frame or segment.

Once the image is rendered (and saved), you could stop there and pass it on to your print service, but you're not really done. There are two reasons to finish the job before handing it off. First of all, you usually pay an extra premium to have the service bureau convert and rasterize your print. This could come to as much as $50 or $100 for a single page. But beyond financial reasons, you would be shocked at just how few print shops have any clue about digital imagery, and I mean *really* shocked. Time and time again I have told printers that I have a magazine-sized, 133 linefeed CMYK TIFF file, and that I need color negatives and a proof. This is the standard language of a print shop, and yet only two out of 10 shops I contacted had any idea how to do this. And of the two that thought they could handle it, one had to be instructed every step of the way. One large printer/publisher I spoke with had a special person to work with Photoshop, Quark, PageMaker, etc. This company, by far, had the least idea of what to do. I can only hope that you have better luck than I did. So what is involved? Not much, actually. I'll give a quick summary of the steps in Photoshop.

Off and Running

Start by using the Print Manager to add a driver for the printer your output will be going to. For the example above, you would select the Linotronic 330 driver. Next, load the image into Photoshop, mode convert it to CMYK color, and image size it to two times your linefeed (266 pixels/inch in this case) while constraining the file size. Image Size should now report the exact width and height your print will be. Go into Page Setup and select your printer. The other settings will require some information from your print shop, but it's typical to select registration and crop marks, negative image and emulsion side down for color-separated film. Finally, go into the Print requester, select the appropriate print quality (2,540 dpi in our case), choose CMYK printing and enable Print to File. Also, to keep file size down, you will want to select Binary Encoding. That is essentially it. You should now be the proud owner of a multi-megabyte, color-separated PostScript file ready for printing. There are other details, such as assuring proper color calibration, but they go beyond the scope of this article.

The final challenge is to get your data to the print shop. Bernoulli drives seem to be the most widely accepted means of data transport. It's a good idea to save a CMYK TIFF to the drive if there is room. This way, if there is a mistake in your Postscript output, and your print shop people aren't completely brain damaged, they can properly re-create the Postscript output themselves.

A quick note about print color space. Earlier, I referred to CMY and how it is the complement of RGB color. So where does the K come from? K stands for black, and there are two reasons it is added to CMY. Consider that black is created in print by combining all three inks at full density. This can cause an excess of ink, resulting in bleeding and unnecessarily long drying times. Adding a black ink reduces the need for the other three. Second, CMY by themselves do not produce a truly dark black. Black ink obviously helps here as well. The two different processes for removing CMY and replacing them with black are called Under Color Removal (UCR) and Gray Component Replacement (GCR). Which method to use depends on the kind of paper that your work will be printed on.

It Can't Be That Easy!

Now back to LightWave. I have described how to properly set your camera resolution for printed output, but the simplicity of this task is only the result of recent 4.0 changes. If you are working with an older version of the software, the task is a bit more complex. Pre-4.0 LightWave did not have a setting for custom resolutions with non-video aspect ratios. (You can input the proper sizes, but try rendering out the image!) Instead you must choose a resolution that has a 4:3 dimensional ratio.

Let's return to our example above. Our 8.375-inch x 10.875-inch print requires a LightWave camera resolution of 2228x2893. However, that is nowhere near a 4:3 image aspect. First, you will want to rotate your camera 90 degrees so that you render 10.875 inches x 8.375 inches. This will maximize your use of the layout window. We need to cover a 2893x2228 area with a 4:3 dimension. If (Xresolution x 3)/(Yresolution x 4) is less than 1, set the camera resolution to (4 x Yresolution/3) for X and Yresolution for Y. Otherwise, use Xresolution for X and (3 x Xresolution/4) for Y. In our case, that means a custom resolution of 2971x2228. Then frame your image so that either top to bottom or left to right fills the camera view. In our case, we have 97 percent of the image used left to right. Render as described before, with a square pixel aspect and an appropriate segment memory size. Take the resulting image into your favorite image processor, crop it to 2893x2228 and rotate 90 degrees to get you back to 8.375 inches x 10.875 inches. The image is now ready for CMYK conversion and the rest of the steps described above.

That's all there is to it. If you are having color negatives made for a magazine ad or another type of mass reproduction, it is a good idea to get at least one proof made. The proof is a hardcopy positive that shows what the final print will look like based on your negatives. Not only will it be needed by whoever prints reproductions from your negatives, but its also handy for checking how accurately your monitor is calibrated to your print/image-processing application.

You are now fully armed with the information you need to start creating your own magazine ads, full-color brochures, press kits or whatever. Note that if you ever want to do something for the cover of *LIGHTWAVEPRO*, you only have to render the RGB image in LightWave and leave the rest for Avid Media Group. *[Editor's note: If you measure the image part of this cover you'll see that you need to build an image that is 8.5 inches square (8.5 x 266 = 2261). So render out at 2261x2261!—JG]*

LWP

Mark Thompson is director of animation and special effects for Fusion Films, Inc., and producer of the LightWave in.focus training tape series. Send questions or comments to Fusion Films, Inc., 51 Derry Street, Merrimack, NH 03054, or e-mail Mark at mark@fusion.mv.com.

Basic steps to printing bliss:

(1) Determine the physical size of the image you wish to print (in inches).
(2) Talk to the print shop and determine what line screen will be used.
(3) Set your LightWave render resolution to 2 x line screen x width by 2 x line screen x height.
(4) Select Square Pixels for the aspect ratio.
(5) Render and save to selected file format.
(6) Load image into Photoshop or equivalent and convert to CMYK.
(7) Set image ppi to 2 x line screen without altering the physical size.
(8) Get print details from print shop/publisher (printer used, dpi, negatives, emulsion up/down, etc.).
(9) Configure print driver and output to a PostScript file.

SIDEBAR

Curve Selection Order and Matching Polygon Edges

If you do have to "dead-end" a curve into the side of another, as we did in "LightWave 101" on page 6, carefully choosing the proper number of perpendicular and parallel segments after selecting your curves in the proper order can give you smooth edges between adjoining patch areas.

- Figure 1 shows the Z curve in the thumb area of the splint as it would appear if it was not extended along the length of the splint. Select the curves in the order shown to patch the curves and select **Patch** (Ctrl–f).
- The Patch requester's Perpendicular and Parallel segment values are determined by the last curve selected (number 4 in this case). Enter 5 for **Perpendicular** and 10 for **Parallel.** In order to make sure that the segments are produced in relation to the knots (or points) along the curves, see that **Knots** is selected for both segments.

Figure 3 shows the first patch section created. Notice that there are five rows of polygon segments perpendicular to the last curve selected and 10 rows parallel to the last curve.

- Deselect all curves and then patch the area to the right of the last patch section by selecting the appropriate curves in the same order. (It doesn't matter if you accidentally select some faces along with your curves—they are ignored in the patch operation.)
- Deselect all curves and patch the large adjoining area below the two sections, using 10 for both **Perpendicular** and **Parallel** values. Figure 4 shows all three patch sections.
- Perform an **Automatic Merge Points** operation (m) to eliminate all of the duplicate points along the shared edges and eliminate "seams" in the rendered image.

One problem that you may encounter while performing a patch operation of this type is acquiring sections in your object with dissimilar curves. In the case of this splint, not running the thumb curve down the length of the splint created polygons like those in Figure 4. The problem is that we have lost the continuous defined shape by dead-ending the thumb curve. You can, however, run the curve down the length of the splint as outlined in the article and still follow the above procedures to make free-flowing shapes. Just because a curve is in an object doesn't mean that you have to use it to patch—it can be there just to help define the surrounding curves. Figure 5 shows the results of patching this way. Notice the smoother-flowing polygons.

What have we learned from this? If possible, try to continue your curves throughout an object to define the overall shape, and patch the proper ones to generate as even a patch area as possible across the surface. If you do have to dead-end a curve, here's a good way to get around it.

—John Gross

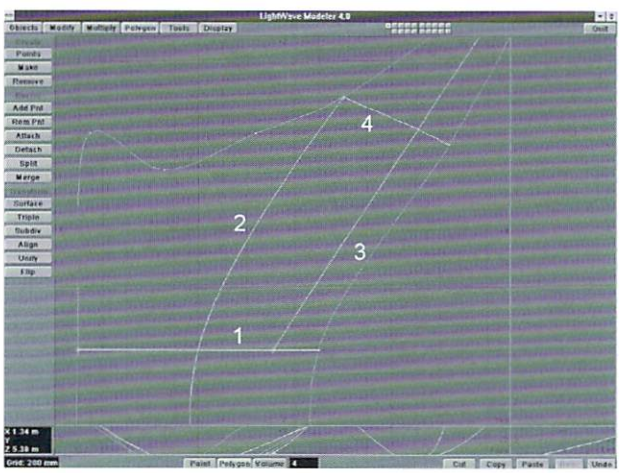

Figure 1: The four curves chosen in the proper order.

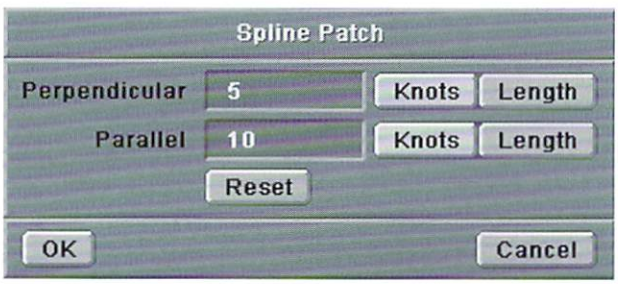

Figure 2: The Patch requester.

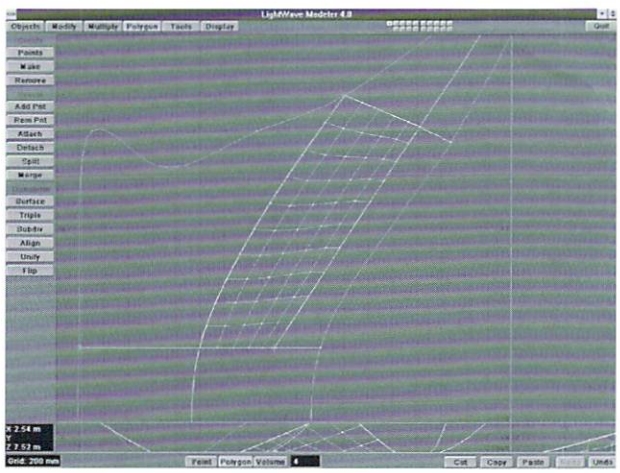

Figure 3: Perpendicular and Parallel polygon segments relative to the last curve selected.

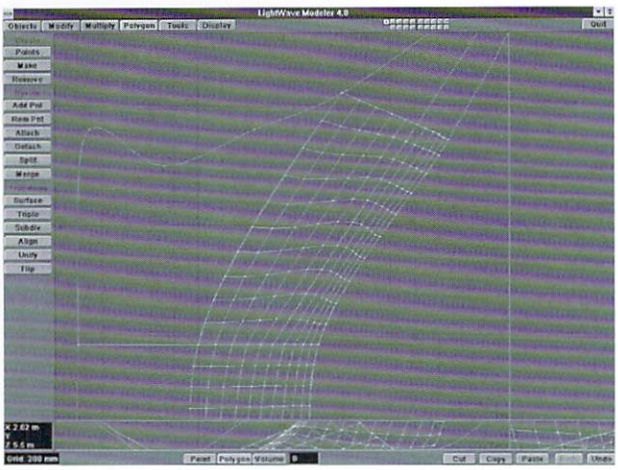

Figure 4: All sections patched. Notice that the polygons share points along all adjacent edges.

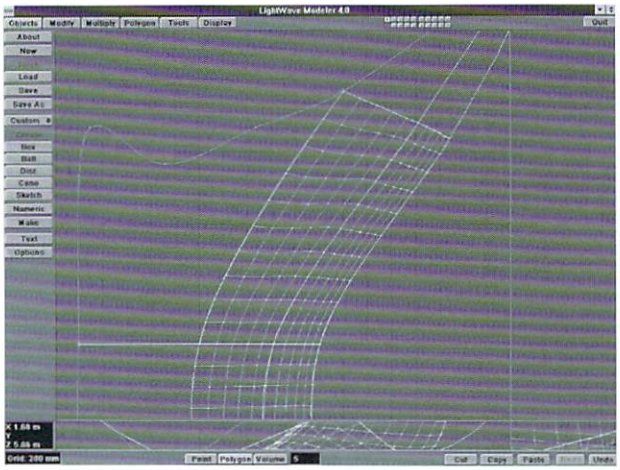

Figure 5: A better way to apply the same patching orders. Notice the entire patched section flows much nicer than in Figure 4.


THE BEST OF
LIGHTWAVEPRO
THE BEST OF LIGHTWAVEPRO COMPILATION BOOK
■ A collection of over 100 articles previously published in LIGHTWAVEPRO newsletter.
■ Written by the world's top LightWave 3D experts, including animators from Amblin Imaging, Area 51, Foundation Imaging and other top animation houses.
■ Includes a CD-ROM with sample models, surfaces and LIGHTWAVEPRO disk programs.
ORDER NOW
ONLY
$39.95!
(California residents also add $3.10 sales tax)
CALL 800-322-2843
or 408-743-9250 for
MC/VISA orders
or make checks payable to:
Avid Media Group, Inc.
ATTN: LWPRO Compilation
1308 Orleans Drive
Sunnyvale, CA 94089
COMPILATION BOOK
NOTE: U.S. customers add $3.00 for shipping cost. All others add $7.00 for shipping cost.

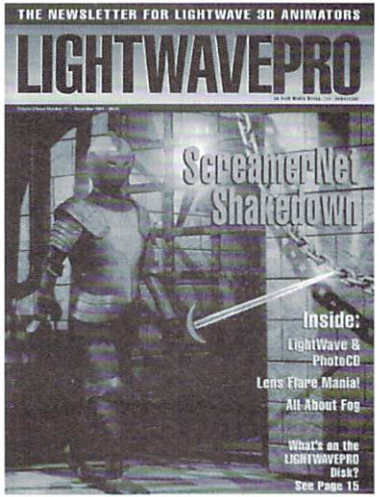

TO ORDER WITH YOUR VISA OR MASTERCARD CALL TOLL FREE!

1(800) 322-2843

LIGHTWAVEPRO

THE NEWSLETTER FOR LIGHTWAVE 3D® ANIMATORS

Make a Smart Investment

Subscribe Today!

Canadian/Mexico: Add $US12
Overseas: Add $US36
Allow 6-8 weeks for delivery.

Make checks payable to LIGHTWAVEPRO.
Prepayment required on all overseas orders.

YES! Enter my one-year (12 issues) subscription to LIGHTWAVEPRO at the Special Introductory Rate of only $48—that's 50% off the cover price!

Name_____

Address_____ Apt. #_____

City_____ State _____ Zip _____

☐ Payment Enclosed ☐ Bill Me

BFYLW

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 2263 SUNNYVALE,CA

POSTAGE WILL BE PAID BY THE ADDRESSEE

LIGHTWAVEPRO

273 North Mathilda Avenue
Sunnyvale, CA 94086-9313

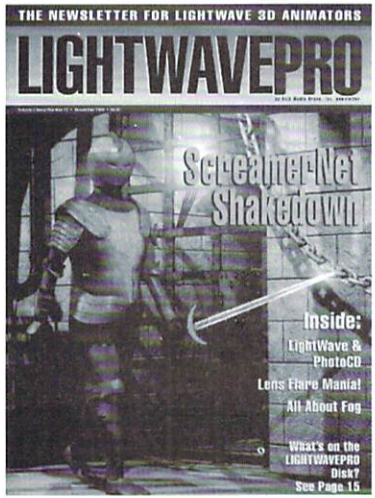

TO ORDER WITH YOUR VISA OR MASTERCARD CALL TOLL FREE!

1(800) 322-2843

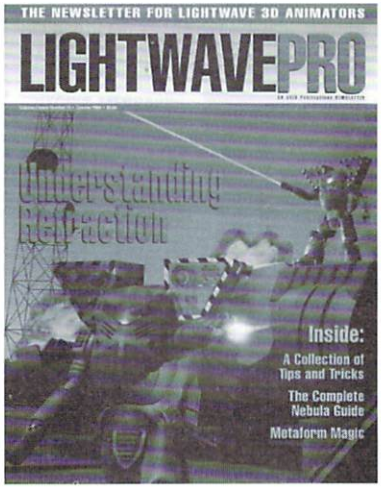

TO ORDER WITH YOUR VISA OR MASTERCARD CALL TOLL FREE!

1(800) 322-2843

LIGHTWAVEPRO

THE NEWSLETTER FOR LIGHTWAVE 3D® ANIMATORS

Make a Smart Investment

Subscribe Today!

Canadian/Mexico: Add $US12
Overseas: Add $US36
Allow 6-8 weeks for delivery.

Make checks payable to LIGHTWAVEPRO.
Prepayment required on all overseas orders.

YES! Enter my one-year (12 issues) subscription to LIGHTWAVEPRO at the Special Introductory Rate of only $48—that's 50% off the cover price!

Name_____

Address_____ Apt. #_____

City_____ State _____ Zip _____

☐ Payment Enclosed ☐ Bill Me

BFYLW

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 2263 SUNNYVALE,CA

POSTAGE WILL BE PAID BY THE ADDRESSEE

LIGHTWAVEPRO

273 North Mathilda Avenue
Sunnyvale, CA 94086-9313

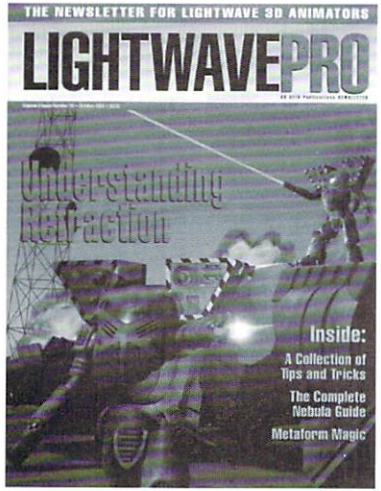

TO ORDER WITH YOUR VISA OR MASTERCARD CALL TOLL FREE!

1(800) 322-2843

Services:
3D Animation
Custom Modeling
Graphics
Animation Packages
T.V. News Graphics • ID's • Opens

Industrial Animation
Forensic Animation
Product Animation

BETACAM SP • 3/4"
S-VHS • VHS OUTPUT AVAILABLE

Hardware:
DEC ALPHA 21164 (266 & 300 MHz)
DEC ALPHA 21064 (275 MHz)
MIPS R4600 (133 MHz)
PENTIUM PROCESSOR
NETWORK SOLUTIONS
MICROSOFT SOLUTIONS

DESKSTATION
TECHNOLOGY
RAPTOR
RENDERING
FARM

Digital Solutions inc.

TWO-TIME
EMMY AWARD
WINNING
ANIMATOR
ON STAFF!!

CALL FOR THE RAPTOR DEALER NEAREST YOU!

Raptor 3 is a trademark of DeskStation Technology
Pentium is a trademark of Intel Corporation
Alpha is a trademark of Digital Equipment Corporation

Digital Solutions, Inc.

2700 East Sunset Road, Suite C20 • Las Vegas, Nevada 89120
Voice **702/261-9686** • FAX 702/261-9856

RAPID TURN AROUND!

LIGHT ROM 3®

for Windows/NT/Amiga

3 CD ROM SET $49.95

You have read about the resources available to artists through the Internet and BBS's, but you lack the time and finances to access them. LIGHT-ROM, published every 6 months, brings to you these resources along with exclusive contributions for an incredible value.

THE REVIEWS ARE IN

"The range of objects is very impressive. LIGHT-ROM is worth every penny."
Jason Holbourn • Amiga Format — January 1995

"A must have for Lightwave users"
Graeme Sandiford • Amiga Format — August 1995

"This CD is a must have for Lightwave users."
R. Shamms Mortier • Video Toaster User — August 1995

"This has got to be my favorite CD of all time."
Graeme Sandiford • Amiga Shopper — April 1995

"This CD ROM is an excellent value."
David Taylor • Amiga Shopper — August 1995

"The nicest collection I've seen for Lightwave models."
Erik Flom • Video Toaster User — October 1995

ROM #1

★ Over 6500 Lightwave objects divided into categories including Anatomy, Aviation, Botany, Buildings, Furniture, FX, Holidays, Household, Logos, Music, Ships, Space, Sports, Tools and Vehicles.

★ Showcase directory features contributions from Lightwave artists around the world including Alan Chan, Jamie Cope, Erik Flom, Tim Nign and Dean Scott.

★ Every Lightwave object is represented in thumbnail renderings (IFF, TARGA & JPEG) for easy previews.

ROM #2

★ Collection of 3D objects in other formats including **Imagine** (175 megs), **3D Studio** (100 megs) and **Sculpt 3D** (30 megs).

★ 700 Textures in JPEG format complete with thumbnails.

★ Public Domain graphic programs and demos for the PC and Amiga.

★ Collection of 3D landscapes in Lightwave, 3D Studio and Imagine formats.

★ Video Toaster directory with wipes and CG fonts.

ROM #3

★ Bonus "DEM ROM" includes over 1000 DEMs (digital elevation maps) for use with VistaPro, World Construction Set and Scenery Animator on any platform.

★ Use DEM ROM (along with any of these programs) to create realistic backgrounds or flights through your favorite scenery to incorporate into Lightwave or other 3D programs as background sequences.

★ All DEMs include thumbnail renderings of their topographical maps.

LIGHT-ROM 3 includes 3 CD ROMs with over 1.9 gigabytes of data and over 17,000 files!

Write for details on how to contribute to the next issue of Light-Rom.

3 CD ROMs for only $49.95 plus shipping.
Shipping and handling is $4.95 for U.S. & Canada – $6.95 International.

GRAPHIC DETAIL

GRAPHIC DETAIL INC. • 4556 SOUTH 3RD ST. • LOUISVILLE, KY 40214 • USA
VOICE / FAX – 502-363-2986 • E-MAIL – michael@iglou.com

MasterCard

★ ORDERS ONLY – 1·800·265·4041 ★

VISA