

how it works

WOULDN'T THE PRODUCT YOU'RE DESIGNING BE EASIER TO CONCEIVE IF IT APPEARED THREE-DIMENSIONAL? A NEW BREED OF STEREO VISION GLASSES AND EXOTIC DISPLAY TECHNOLOGIES MAKE YOUR 3-D DREAMS COME TRUE.

Display enhancements accept no compromises

By Brian Dipert, Technical Editor

Modern graphics hardware and software labor mightily to deliver a life-like 3-D presentation to your eyeballs and brain. Perspective control, graduated lighting and shadows, high-resolution and color-rich multitexture application, and other techniques are the means by

which today's graphics accelerators and CPUs, applications and operating systems, APIs (application-program interfaces) and device drivers, all work together in the hopes of generating a realistic cyber world. If all you're looking at are fanciful monsters and goblins, the approach works quite well. However, view a computer-display representation of an object that you know intimately and view frequently in real life, such as a human face or your backyard, and the limitations become obvious.

After all, the final display device is a solitary 2-D CRT or LCD screen. For an explanation of what's still missing, you need only to lift one finger. Close your right eye, keep your left eye open, and use that finger to block your view of an object in front of you. Now close your left eye, and open the right one. Voilà, you now have an unobscured view of the previously blocked object, and your finger has seemingly moved to the left. This simple experiment reminds us that our two eyes see the world from slightly different viewpoints and suggests that the human visual system uses this combined-image spatial perception, along with other visual cues, to mentally construct a 3-D representation of the scene.

Now take the experiment one step further. First,

place your finger very close to your face, and alternate opening and closing your two eyes. Note how far the finger appears to laterally move. Now extend your finger out in front of you as far as you can, and repeat the experiment. Even if you have short arms, the finger now hardly moves at all. This discrepancy indicates that the parallax effect, the difference in scene perception based on position, is more pronounced when you are viewing close objects and less dramatic when you observe objects at a distance. In fact, astronomers use parallax results, deriving dis-

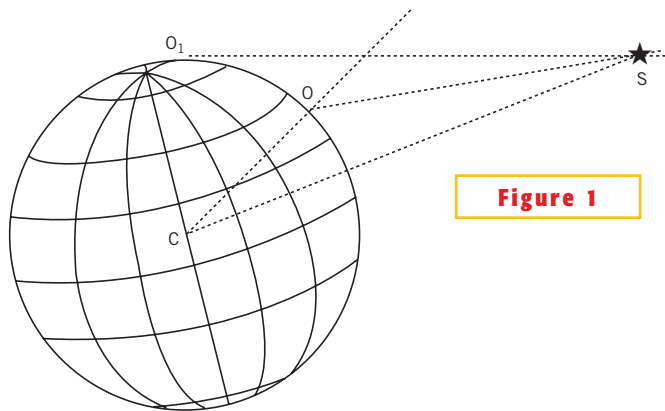



Figure 1



Perspective differences at various points on the Earth's surface combine with a little trigonometry to enable the calculation of far-off object distances.

tance from simultaneous observation of the same object at different points on Earth (Figure 1). Parallax results also depend on the distance between the two eyeballs, which varies from one person to another.

DOUBLE VISION

To create the most realistic-possible 3-D representation of a scene, a computer must generate unique image versions for each eye. Remember the old View-Master stereo viewers? Each picture on the reel (or card for the even older stereoscopes) comprised two unique images, photographed or drawn from slightly different perspectives. Bulky, expensive, and unattractive heads-up displays containing tiny LCDs located in front of each eye are the computer-age equivalents of the View-Master (Figure 2). Neither two computer monitors nor a split-screen monitor will create a stereo effect, because in both cases there's too much eyeball-to-eyeball crosstalk interference. You need to isolate each image to only one eye, but you don't want to look like a *Star Trek* extra in the process.

Hollywood has attempted to solve the image-isolation problem with color filters. Such forgettable movies as *Jaws 3-D* and *Friday the 13th: Part 3* used a single-lens projector, with each frame of the film containing two superimposed images. Objects in one image were shifted laterally to those in the other image, with the amount of shift depending on their distance from viewer. Also, one image contained an excess of one color, and the other image contained an excess of another color. When viewers wore 3-D glasses, the blue or red filter in front of each eye partially blocked transmission of one of the images. Reference 1 shows a still-image representation of this technique. Color-filter-derived 3-D technology has numerous shortcomings. The filters give the scene an unrealistic color mix, for which the brain can only partially compensate; the 3-D effect varies depending on your location in the theater. Unless you correctly position the glasses on your head at all times, you'll end up with some strange-looking results, because each eye might be able to peer out of both lenses.

Some Imax theaters have put a modern twist on the old red-and-blue-glasses idea—a variation that, as you'll soon see, computer-peripheral manufacturers have also embraced. In this setup, the Imax projector contains two sets of optics, two reels of film, and other replicated equipment that matches the stereo nature of the camera that the filmmakers used when they originally filmed the scenes. Each lens, with a uniquely oriented polarizing filter ahead of it, casts a slightly different image on the screen. The audience wears similarly oriented polarized glasses; therefore, each lens of the glasses blocks transmission of one of the two images. The polar-

izer approach doesn't distort the images' color as does the color-filter technique. However, polarization darkens the presentation, and the effect again varies depending on where you sit in the theater, how you place the glasses on your head, and how well the lens-to-lens spacing matches the distance between your eyes. Polarized glasses are also more expensive than those constructed with simple color filters. The cost of both the glasses and the dual-duty projector has prevented this 3-D technique from breaking into the movie mainstream.

Recall that at greater than 30 fps (frames per second) or so, the eye and brain can no longer discern individual displayed frames, creating the illusion of real-life motion. The fact that a computer display is progressive-scan, not interlaced as with a television, further increases the required frame-per-second figure, as does the fact that fast-paced interactive computer games require low latency between the user's actions and the display reactions. Suffice it to say, then, that many 3-D-graphics users demand at least 50 fps from the graphics subsystem that they use, and they're often willing to trade off image quality to achieve their performance objective.



Figure 2

Heads-up displays may be functional, but they sure aren't fashionable (courtesy i-O Display Systems).

WE DON'T NEED NO STINKING GLASSES

The frames-per-second performance threshold is a key parameter that computer-peripheral manufacturers must consider as they come up with lower cost alternatives to heads-up displays for stereo-image viewing. The most common technique they employ involves glasses with LCD "shutters" that alternately darken and become transparent many times per second. In synchronization with the shutters, often coordinated via the horizontal- or vertical-synchronization signals or DDC (display-data-channel) signal transitions, the computer graphics subsystem draws an image or portion of an image that targets one eye.

Early LCD shutters were capable of only very slow




Figure 3

Lenticular displays use subpixel rendering and microlenses to present a stereo image to the user, with no glasses required.

“open” and “shut” speeds, which caused flickering and ghosting as images leaked through to unintended eyes. Having the LCDs close to the eyes helps to minimize flicker, because the shutters are outside of the eyes’ focal length, but the eyes and brain can still detect slow operation. Until recently, you could not set monitors to high vertical-refresh rates (near or, ideally, greater than 100 Hz) without damaging them, unless you also set your display to a low resolution. And graphics accelerators did not have the performance capability to render anything other than simple, color-deficient and low-resolution scenes at triple- or even high double-digit frame rates.

As a result, the graphics hardware and software would alternate between images targeting the right and left eyes on a horizontal line-by-line basis within each frame. The advantage of this approach is that if, for example, you want to project a 50-fps image to each eye, you have to achieve an only 50-fps display-frame rate. One key disadvantage of this approach, though, is that the effective vertical resolution each eye sees is only half that which the monitor and graphics subsystem delivers, and the eye and brain must (imperfectly) interpolate across the in-between-line blanks. If you set the display to 800×600 (SVGA), for example, each eye would see only 300 horizontal lines of resolution, with gaps between them.

Today’s high-end graphics hardware easily hits high double-digit frame-per-second speeds, and modern CRT monitors tolerate you setting them to triple-digit refresh rates. In response, the latest generation of shutter glasses alternate between the eyes on a frame-by-frame basis, with each eye seeing half of the total displayed frame rate but at full resolution. Compared with the earlier example, a 100-fps refresh setting now means that each eye sees a 50-fps refresh rate but at the full 800×600 monitor resolution. The glasses’ LCD shutters now operate only at a frame-by-frame toggle rate, not the earlier line-by-line rate. LCD monitors still use line-by-line im-

age alternation; their screens do not yet achieve the refresh rates possible with CRTs.

LCD shutter glasses are lighter and smaller than heads-up displays, but they’re still somewhat odd-looking. If you don’t want to put full-blown LCD shutter glasses on your head, you have other options. Place ahead of the monitor a screen with an electrically adjustable polarization orientation that coordinates with the monitor’s horizontal- or vertical-sync signal and put on Imax-like polarizing glasses. Or, dispense entirely with the glasses. Lenticular displays incorporate a sheet of cylindrical lenses on top of an LCD divided into subpixels. By selectively controlling the subpixels and their grouping, the system can project images to each eye for a stereo 3-D effect (**Figure 3**). Keep in mind, though, that the subpixel rendering used here is subject to many of the same trade-offs present with ClearType, CoolType, and other subpixel-font-rendering technologies (**Reference 2**).

WHERE’S YOUR HEAD?

Now that you know the various options available to display stereo images, back up and consider how the graphics hardware and software create the images in the first place. First, you need to tell the graphics subsystem the distance between your eyes and the size and aspect ratio of your monitor. Then, you need to tell it how far away each of the pixels is that comprises each displayed 3-D object. The Z-buffer, a per-pixel memory array containing non-linear depth values of 0 to 1, most commonly stores this information. The graphics subsystem uses the Z-buffer to determine, when two objects overlap, which one it should render in front of the oth-




Figure 4

Depth-dependent lateral shifting of both eyes’ images is the least computation-intensive approach to 3-D stereo; more elaborate but accurate techniques completely render each image from the corresponding eye’s unique vantage point (courtesy Elsa).

er. If a program is incompatible with a stereo-vision display, it may be because the program bypasses the Z-buffer and controlling API and writes pixel data directly to the frame buffer.

Having determined eye, monitor, and object-depth data, the graphics subsystem can now calculate the parallax effect and generate appropriate laterally shifted images (Figure 4). For greatest accuracy, though, simple lateral shifting isn't enough. Consider, for example, a situation in which you might be looking straight at one of the faces of a cube. If the cube is smaller than the projected space between your eyes, your right eye might be able to also see just a bit of the cube face on the right side, while the left eye can see "around the corner" to the other side. Images ren-

dered with all perspective effects comprehended for each eye are more computationally challenging than simple depth-dependent horizontal shifts of objects in an otherwise-identical image, because they require complete re-renderings of the entire scene beginning at the earliest polygon-based stages. However, they most closely approach real life.

Observer position is the final piece of data that is valuable to graphics hardware and software when generating stereo images. Is every computer user's head always exactly centered in front of the monitor and at a fixed distance from the screen? Or, far more likely, are these coordinates user-specific and do they, in fact, vary over time even with the same user? Head location and orientation combine with

TRUTH IN ADVERTISING

Just a few years ago, 3-D stereo glasses cost hundreds or thousands of dollars, were bulky and heavy, and generated so much flickering and ghosting that they gave users headaches after only a brief period of use. So it was with no lack of skepticism that I approached an evaluation of Elsa's 3-D Revelator glasses. A wired version that supports simultaneous operation by as many as four users cost \$55 (from Elsa's online store; less elsewhere) and an infinite-user wireless configuration cost \$75. The price was right, but did Elsa make unacceptable trade-offs to achieve it? Fortunately, my pre-review cynicism was unnecessary; these glasses exceeded even my most optimistic predictions.

I tried out the glasses with two graphics boards: Elsa's Erazor X², based on Nvidia's GeForce 256 (installed in a Pentium III 133/533 system using Intel's VC820 motherboard), and Creative Labs' 3-D Blaster, based on Nvidia's TNT2 Ultra (installed in a Pentium III 100/600 system using Intel's SE440BX-2 motherboard). Originally, Elsa's glasses worked only with the vendor's own 3dfx-, Nvidia-, and S3-based graphics boards, but in recent months

Elsa has added support for other companies' Nvidia-derived products. Driver installation was uneventful with both systems, and Elsa includes a utility that lets you toggle back and forth (with corresponding system reboot) between its drivers and other ones you've installed and intended for your board. Unlike previously developed stereo glasses from other manufacturers, the 3-D Revelators required no installation of any special application-specific patches.

The visual quality of games such as Acclaim Entertainment's Forsaken, Criterion Studio's Redline Racer, and Rage Software's Incoming has always impressed me. Elsa's stereo glasses generate an enhanced 3-D effect that takes graphics to a whole new level of realism and, frankly, blew me away. Missiles create a much greater sense of urgency when they seem to be coming right at you. Racing a motorcycle through a desert canyon is far more realistic when the canyon walls seem to be hurtling by you. And the ability to stretch the background of a room or an outdoor panorama back behind the monitor glass dramatically increased the depth perception.

Frame-rate impacts were

unnoticeable in most cases.

Most times, the games were still playable at their original quality, color-depth, and resolution settings, though with the older TNT2 Ultra board, I found it necessary to sometimes back off to less performance-strapping settings to eliminate display stutter. The GeForce 256 and the system containing it took on the added processing burden with no perceptible strain.

My 21-in. monitor, capable of greater-than-100-Hz vertical-refresh settings, enabled me to operate the glasses with no perceptible flicker. I experienced no eyestrain even after several hours of use, probably due to both careful setup (I even used a ruler to measure the spacing between my eyeballs) and the large monitor.

The glasses, as expected, did not have perfect, 100% light transmission with open "shutters," so I raised the monitor's brightness setting. They also didn't fully block light transmission when the shutters were "closed." As evidence of this fact, I could detect slight ghosting, manifesting as a solid image in the middle with faint images on either side of it.

The lack of OpenGL support disappointed me, both because I

wanted to try Quake III and Unreal Tournament and because digital-content-creation (CAD, 3-D modeling, and others) software capability also would have been cool. Limited program-specific OpenGL support currently exists for Elsa-branded cards, although Elsa's recent enhanced relationship with Nvidia promises to improve this situation (Reference A). I didn't try SciTech's OpenGL wrapper for DirectX, but I've heard that it lets you play OpenGL-based games and doesn't significantly cripple performance.

At press time, Elsa's drivers were based on Nvidia's 5.31 driver version, many months removed from Nvidia's state of the art. But considering the price of the glasses and the plethora of DirectX-based applications out there, the 3D Revelator glasses are a great value. Check 'em out. Also, take a look at the 3-D still images and video clips sold by i-O Display Systems, whose LCD technology many other glasses vendors, such as Elsa, also employ.

REFERENCES

A. Dipert, Brian, "Thinning ranks don't diminish graphics pace," *EDN*, Oct 12, 2000, pg 28.

all the other previously mentioned variables to enable the computer to present to the user the most realistic stereo scene representation possible.

To get a feel for the additional information that a stereo-display system needs versus a standard single-monitor setup, look at the custom graphics drivers that Elsa developed for its 3-D Revelator glasses (Figure 5). A one-time configuration utility enables you to set the monitor-refresh rate and fine-tune the

drivers to your eye-to-eye spacing, monitor size and aspect ratio, and eye-to-monitor distance. You can enable and disable the stereo effect both through the control-panel window and via a keyboard hot-key sequence and adjust the clipping window that determines which portion of the scene is displayed.

Elsa's drivers also support the optional Dyna-Z feature, which dynamically alters the stereo effect depending on the range of Z (depth) values present

FOR MORE INFORMATION...

For more information on products such as those discussed in this article, go to our information-request page at www.rscanners.ims.ca/ednmag/. When you contact any of the following manufacturers directly, please let them know you read about their products in *EDN*.

3-D STEREO GLASSES

AnotherWorld

+82-31-337-2468
www.anotherworld.to
Enter No. 396

Artificial Parallax Electronics

www.apec.com.tw
Enter No. 397

ASUSTek

1-510-739-3777
www.asus.com
Enter No. 398

Elsa

1-408-961-4600
www.elsa.com
Enter No. 399

i-Art

+886-2-23952883
www.i-art.com.tw
Enter No. 400

Illico/i-O Display Systems

1-415-323-8270
www.illico.com
www.i-glasses.com
Enter No. 401

Metabyte/Wicked3D

1-510-494-9700
www.metabyte.com
www.wicked3d.com
Enter No. 402

NuVision

1-503-614-9000
www.nuvision3d.com
Enter No. 403

StereoGraphics

1-415-459-4500
www.stereographics.com
Enter No. 404

Tetratel

1-780-448-1019
www.tetratel.com
Enter No. 405

3DTV

+33-2-4011-6299
www.3dmagic.com
Enter No. 406

VRex

1-914-345-8877
www.vrex.com
Enter No. 407

VR Standard

1-562-921-2000
www.vrstandard.com
Enter No. 408

Woobo

+82-2-5466-006
www.woobo.com
Enter No. 409

3-D STEREO DISPLAYS

Dimension Technologies

1-716-436-3530
www.dti3d.com
Enter No. 410

Dresden 3D Display

+49-351-463-0
www.inf.tu-dresden.de/D4D
Enter No. 411

Elsa

1-408-961-4600
www.elsa.com
Enter No. 412

Heinrich-Hertz Institute

+030-31-002-0
www.hhi.de
Enter No. 413

NuVision

1-503-614-9000
www.nuvision3d.com
Enter No. 414

Philips

1-212-536-0500
www.philips.com
Enter No. 415

RealityVision/Richmond Holographic Studios

+44 (0)171 739 9700
www.brunel.ac.uk/depts/mes/Research/Groups/vr/vrsig97/proceed/008/hasdpape.htm
www.cuni.cz/~pavlik/skola/3d/autostereoscopy/rhs.htm
Enter No. 416

Sanyo

1-619-661-6835
www.sanyo.com
Enter No. 417

StereoGraphics

1-415-459-4500
www.stereographics.com
Enter No. 418

3D Experience

01279-817219
www.3dexperience.co.uk
Enter No. 419

VRex

1-914-345-8877
www.vrex.com
Enter No. 420

OTHER COMPANIES MENTIONED

Acclaim Entertainment

www.acclaim.com

Creative Labs

www.3dblaster.com

Criterion Studio

www.csl.com

Intel

www.intel.com

Nvidia

www.nvidia.com

Rage Software

www.rage.com

SciTech's

www.scitech.com

S3-based

www.s3.com

3dfx

www.3dfx.com

OTHER 3-D STEREO RESOURCES

Google

www.google.com/search?q=stereo+3d

Stereoscopic 3D Virtual Reality

www.stereo3d.com

SUPER INFO NUMBER

For more information on the products available from all of the vendors listed in this box, enter No. 421 at www.rscanners.ims.ca/ednmag/.

in a given scene. In a shallow room, for example, the stereo effect might be exaggerated; in a deep room, it may be more subdued. When you enable the stereo display, additional hot-key settings let you adjust the minimum and maximum Z values that trigger a parallax adjustment. This feature is useful when, for example, you don't want a 2-D object, such as an airplane-instrument panel, to project in front of the monitor. Other hot keys enable you to shift the entire scene more in front of or behind the monitor screen, narrow or broaden the lateral shift, reduce or increase the nonlinear Z-buffer mapping factor, or save the current settings in an application-specific profile.

Incorporating stereo-display support obviously requires source-code access and extensive modifications to the OEM drivers that the graphics accelerator vendors develop for their chips.

This reality puts quite a software-development

burden on the 3-D-glasses or display manufacturer, sometimes with undesirable consequences.

The drivers might only communicate with the DirectX API, for example, not with OpenGL or Glide, unless you also use a performance-sapping API "wrapper" with limited features. They might work with only one vendor's chips; a one- or few-family subset of those chips; or, worse yet, only one vendor's boards based on those chips.

Support for other nonuniversal graphics-board features, such as video I/O or TV-tuner functions, may disappear when you install a stereo-display-aware driver. If the stereo-enabled drivers lag by several generations the standard driver software that the silicon vendor supplies, you'll have to accept compromises in stability, performance, quality, and features. And the additional calculations that the graphics subsystem needs to create each image, coupled with the fact that it's now creating two stereo images for every one that users view the "normal" way, can cause the frame rate to plummet. But don't let these disadvantages prevent you from at least trying out a 3-D stereo monitor or pair of glasses. Match them with the right graphics card and software, and the results will amaze you. □



Figure 5

A highly customizable software driver ensures that each user has the best-possible stereo viewing experience (courtesy Elsa).

You can reach
Technical Editor
Brian Dipert at
1-916-454-5242,
fax 1-916-454-5101,
e-mail bdipert@pacbell.net.



REFERENCES

1. Dipert, Brian, "Balancing in three dimensions," *EDN*, April 27, 2000, pg 54.
2. Dipert, Brian, "Display technology's results are compelling, but legacy is unclear," *EDN*, Oct 26, 2000, pg 63.
3. Dipert, Brian, "Getting glitzy with graphics for embedded systems," *EDN*, April 1, 1999, pg 79.
4. Wilson, Andrew, "Autostereoscopic system delivers 3-D images using cylindrical lenses," *Vision Systems Design*, May, 2000, pg 45.