Getting Real

WHAT'S NEXT IN COMPUTER DISPLAYS? DEPTH AND SHADOWS BY MARK ALPERT

Nearly 20 years ago I had the dubious honor of viewing *Jaws 3-D*, one of the sequels to the infamous shark-attack movie. With my theater ticket I received a pair of cardboard 3-D glasses, with red cellophane in one lens and blue in the other. Feeling very stupid, I sat in the front row and donned the glasses only when the lights went down. I remember absolutely nothing about the movie's plot, but I recall with great clarity some of the 3-D effects. The opening sequence featured a severed fish head, which seemed to be floating gruesomely just inches from my face. I was still young enough at the time to think that this was pretty cool.

Fast-forward to 2002. I'm sitting in a darkened room at the Media Research Laboratory at New York University, staring at a device called an autostereoscopic display. The setup looks very odd: a computer monitor lies on its side, and a sheet of liquid crystal—called a parallax barrier—is positioned about three inches in front of the screen. On each side of the screen is a small camera surrounded by tiny infrared lights. When the system is turned on, I see a ghostly 3-D image floating in the space between the screen and the parallax barrier. It's a skeletal model of a human foot, with the tarsals, metatarsals and phalanges highlighted in different colors. The image slowly rotates, and it seems just as gruesomely real as that severed fish head from *Jaws 3-D*. But now I'm not wearing any red-and-blue glasses.

Standing behind me are Ken Perlin, the laboratory's director, and several colleagues who have worked on the project since 1998. Most previous attempts to create 3-D computer graphics have displayed two images on the screen and required users to wear either polarized glasses (which filter out one image for each eye) or shutter glasses (which alternately block the view from each eye). Perlin and his colleagues decided to eliminate the need for cumbersome eyewear. Their system also displays two images, one for the left eye and one for the right, but the images are interleaved on the screen. The result is a jagged mishmash of vertical strips. (Because a typical computer monitor scans horizontal lines, the system designers simply turn the screen on its side to produce the vertical effect.)

The parallax barrier turns this mishmash into a 3-D image. Think of this device as a picket fence between the viewer and the screen. The liquid-crystal sheet—in technical terms, a pi-cell—forms vertical black stripes that move rapidly from left to right. At any one moment, the user's left eye observes a striped view of one image (as if seen behind a fence) and the right eye observes a similar view of another image. The stripes prevent the left eye from seeing the image meant for the right, and vice versa. Because this fence moves so quickly—every sixtieth of a sec-

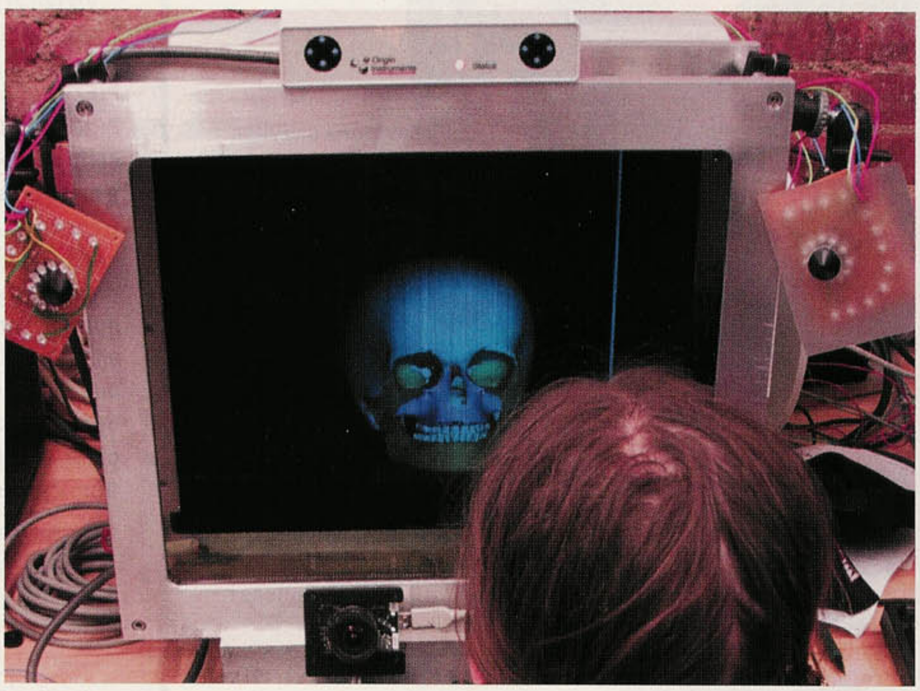

HERE'S LOOKING AT YOU: This autostereoscopic display shows a 3-D image with the help of a parallax barrier (*in front of screen*) and two infrared cameras (*on either side*) that track the user's eyes. The device was conceived at the Media Research Laboratory at New York University.

NEW YORK UNIVERSITY CENTER FOR ADVANCED TECHNOLOGY ond the parallax barrier cycles through three positions—the user doesn't perceive the stripes. Instead each eye sees a full image, which is oriented according to the eye's position. That is where the infrared lights come in; the cameras track the positions of the user's eyes by measuring the reflection of the light off the user's retinas. Only one user at a time can see the 3-D image.

I tested the limits of the autostereoscopic display by moving my head in a variety of ways, jutting it up and down and back and forth. The system adjusted the image of the skeletal foot to match my movements, allowing me to view the thing from practically any angle. Occasionally I glimpsed some distracting black gaps in the image, perhaps caused by stray reflections of the infrared beams. The skeletal foot also appeared a bit dim because the moving stripes of the parallax barrier block some of the screen's light. Perlin assured me that this problem could be fixed by employing a brighter screen to compensate.

As I gazed at the 3-D foot, it occurred to me that surgeons might find this kind of display quite useful. But the first applications of the technology are more likely to be entertainment-oriented: N.Y.U.'s Media Research Lab is conferring with a software company about developing 3-D displays for computer game systems. (As if those games aren't grisly enough already!) It also occurred to me that the purveyors of Internet pornography might have their own reasons for promoting such displays. The members of Perlin's team wisely declined to comment on this prospect.

At Columbia University, just a few miles north of N.Y.U., researchers at the Columbia Automated Vision Environment (CAVE) laboratory are working on another display technology that promises to enhance the realism of computer graphics. Called the Lighting Sensitive Display, this device can judge the direction of the exterior light striking a computer screen and then adjust the image on the screen so that it has the appropriate

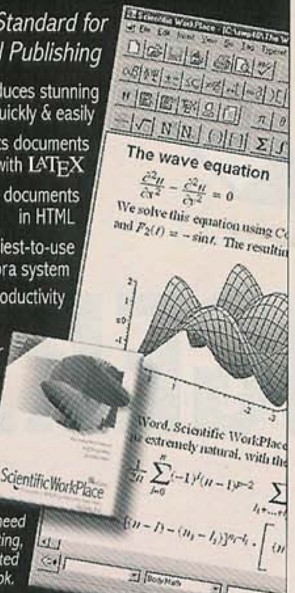

TECHNICALITIES

LIGHTING SENSITIVE DISPLAY developed at Columbia University changes the shadows and reflections in an image to match the exterior light. Note how the pepper mill's shadow shifts from right to left.

COLUMBIA AUTOMATED VISION ENVIRONMENT

shadows and reflections. Shree Nayar, the head CAVE man, led me into another darkened room in which a flat-panel display hung on the wall like a painting. On the screen was a still life containing a bowl of grapes, a decanter of wine, a pepper mill and a dead fish on a cutting board. (Why do dead fish keep floating to the surface of this story?) Embedded in the screen's wooden frame was a hemispherical camera with a wide-angle lens that could detect light sources anywhere in the room.

Nayar handed me a gooseneck lamp and urged me to wave it in front of the display. As I moved the lamp, the image of the still life changed: the shadow of the pepper mill moved from left to right, glints of light appeared and disappeared on the grapes, and the decanter sparkled. The effects were almost instantaneous—the screen can change at 15 frames per second. Although the resolution was not quite good enough to show the tiny reflections off the scales of the fish, it was nonetheless an impressive show.

The story of how Nayar and his colleagues developed the Lighting Sensitive Display is equally impressive. At first they considered ray tracing, the technique used by computer animators to cast shadows on their cartoon landscapes, but they soon realized that this method was too computationally intensive to change the image quickly enough. Instead they decided to photograph an object under hundreds of different lighting conditions and store all these images in the system's memory. (Photographing the still life took so long that Nayar says the fish began to stink.) When the embedded camera detects a light source at a certain angle, the screen quickly calls up from memory the image with the appropriate shadows and reflections.

Because this approach involved a huge amount of data, Nayar's team developed an ingenious way to compress the information: divide the image into many parts using a Cartesian grid. Moving the exterior light doesn't change every part of an image; many parts don't change at all, and other parts change very little. So there is no need for the display system to store more than one or two images for each of the less changeable parts. Using this compression method, Nayar's group was able to reduce the system's memory requirement from an unwieldy four gigabytes to a more manageable 10 megabytes.

Nayar says the technology could have some educational uses. His group has already created a high-quality rendering of Michelangelo's *David* that can be viewed with the Lighting Sensitive Display. Nayar himself is an amateur painter, which perhaps explains why his fish-and-fruit image resembles a still life by one of the 17th-century Dutch masters. "The Dutch were obsessed with getting the fruit right," he says. "To show the translucency of the grapes, they used layers of oil and put the highlights beneath the surface." Maybe the typical computer user doesn't require quite the same photorealism on his or her screen, but it's nice to see that graphics designers are at least aiming high. SA