

Seeing Triple

Anticipated for decades, machines are finally displaying real objects in three true dimensions

Inventors have struggled for years to create displays that can conjure vivid three-dimensional images that users can manipulate and interact with. Chemists could exploit such marvels to design new drug molecules. Oil and gas explorers could see exactly where to aim their drills. Surgeons could pass probes or radiation beams through

the collated slices of diagnostic data produced by magnetic resonance imaging (MRI) and computed tomography (CT) machines to test procedures before performing an operation. But shortcomings such as a flickering image, a narrow angle of view or the need to wear special glasses have bedeviled the devices.

Two companies have recently mixed their own technology with off-the-shelf components, including the Digital Light Processor (DLP) chip from Texas Instruments, to create interactive systems called 3-D volumetric displays that overcome these limitations. The two firms'

products are just now transitioning from the laboratory to commercial models.

Spinning Algorithms

WAIT A MINUTE. Aren't holograms three-dimensional and viewable without funny glasses? Yes, but they are recorded once as a final image and thus do not allow interactivity. Engineers have also knit together cubes and spinning arrays of light-emitting diodes to give a full-bodied view, but the resolution is coarse, restricted by the connections among the diodes. Other contenders appear to be 3-D but really are not; the Heliodyisplay

By
Stuart F. Brown



ANGLE OF ATTACK: Data from a CT scan is projected in 3-D in PerspectaRad, revealing a brain tumor core [yellow diamond, center] and possible paths for radiation treatment [green rays].

from IO2 Technology in San Francisco projects floating images onto a vertical plane of fine mist suspended above the instrument that seem to possess depth, but the illusion comes from an absence of depth cues, not actual imagery in the third dimension. Users who want to load 3-D medical data into a machine one day and 3-D military scenes the next and then turn, prod or alter them while on view can exploit two inventions worthy of the term “volumetric”: Perspecta and DepthCube.

Perspecta, developed by Actuality Systems in Bedford, Mass., might best be described as a crystal ball for looking inside objects. A transparent polycarbonate dome houses a flat, disk-shaped screen 10 inches in diameter that rotates on a vertical spindle at 900 revolutions per minute. The system takes data generated

by a CT, MRI or PET (positron-emission tomography) scanner and mathematically divides the information into 198 radially disposed segments, like an apple thinly sliced about its core. Held in a frame-buffer memory, the data slices are fed to three DLP chips. DLPs are arrays of hundreds of thousands of tiny mirrors, each of which can be individually tilted by onboard circuitry. They form the heart of projection televisions and new slide projectors, as well as digital movie projectors that may replace film reels in theaters. In Perspecta, each DLP is assigned a color and projects its light through a prism onto the rapidly spinning screen, creating a 3-D apparition.

A lot of mathematical heavy lifting was needed to make Perspecta work. “It took us three or four years to invent the algorithms that let us slice the image

data,” says Gregg E. Favalora, Actuality’s chief technology officer. “For instance, we have a patent on how you draw a straight line on a rotating screen—because it is not obvious what dot to pick on that grid as it spins around.”

Perspecta creates a glowing, semi-transparent image. Every volumetric pixel (voxel) that seems to be in a certain point in space really is there, but it becomes visible only when the screen sweeps through that point—hence the fast rotation rate. The screen is made from a plastic that looks like taut tissue paper and is 50 percent reflective and 50 percent transmissive, allowing the imagery to be seen from any angle by observers gathered around it. The unit provides both vertical and horizontal parallax; if a viewer moves his or her head up and down or left and right, background objects that were previously obscured by foreground objects will come into view, as they are perceived in the real world.

A user can also wield a penlike mouse to zoom into or out of the image, rotate and flip it, or change colors. This feature has only recently been made possible by rapid advances in computer graphics. “In our first demo in 2002,” Favalora recalls, “it took 45 minutes of processing time just to perform a click-and-drag. Now we get a video card off the shelf for a few hundred dollars, and it computes the problem just like that.”

It was also a while before Favalora and his colleagues realized that perfecting the core technologies would not be enough to build a business; they had to identify an initial market and develop a ready-to-use system tailored to it. That niche turned out to be radiation therapy

THE AUTHOR

STUART F. BROWN was formerly a staff writer at *Popular Science* and *Fortune*. He calls his now freelance beat “the man-made world,” which includes aerospace, transportation and biotechnology.

for cancer tumors. Doctors need to carefully plan the paths along which they aim radiation beams, trying to maximize the killing effect where the rays converge on a tumor while minimizing damage to nearby healthy tissue. Because oncologists must work with 2-D slices of scanner data, planning the beam paths for a treatment can take several hours. Actuality developed its PerspectaRad system as an add-on to existing radiation therapy equipment manufactured by Philips Medical Systems.

PerspectaRad sports the 3-D display plus software that connects the device to the Philips systems. When a doctor pushes a button, an image of the CT data for, say, a brain tumor appears in 3-D. Another button adds the radiation pathways chosen by a dosimetrist, who plans the treatment. The physician can

see exactly where the beams will strike the tumor, which healthy brain tissues they will pass through, and the dose cloud—the volume of tissue that will be affected by the radiation. This imagery helps doctors adjust the beams to improve treatment or reduce damage. The first PerspectaRad systems cost about \$90,000. Greater production could lower the price to \$65,000, according to Favalora, but the displays are unlikely to reach consumer markets.


Nevertheless, treatment stands to gain. James Chu, head of the medical physics department at Rush University Medical Center in Chicago, recently studied 12 patients with brain tumors for whom treatment plans were developed using both PerspectaRad and conventional methods. The plans were reviewed by doctors who were unaware of

which method had been used. The protocols developed with PerspectaRad turned out better in six cases, equivalent in four cases, and worse in two cases. In one patient, PerspectaRad made it clear how to reduce incidental damage to the optic nerve. Calling the results “interesting,” Chu is planning a larger study that will include patients with tumors in other parts of the body. “When working just with CT data,” he says, “you have to look at individual slices and somehow integrate them all in your head to get a 3-D picture. With Perspecta, you see the 3-D picture directly.”

Chu is also excited by Perspecta’s ability to display moving images of internal body parts. Because internal organs and tissue move as the heart beats and the lungs fill and empty, it is very useful to be able to discern the axis of motion of a tumor. With this information, a doctor could direct a lower-energy radiation beam along the motion axis, rather than a more intense beam across it, lessening collateral damage. Chu says Perspecta could also allow more precise implanting of radioactive “seeds” in the prostate gland to treat cancer there, by allowing a physician to better compensate for motion of tissue that occurs when the needle that delivers the seeds is inserted.

PERSPECTA: CRYSTAL BALL

A transparent dome, translucent screen and optics all rotate at 900 rpm to create a 3-D image. A computer sends graphics data to electronics below the pedestal, which instruct three light-processor chips to focus light from an arc lamp through a projector lens. It reflects a beam up through the spinning shaft and across relay mirrors onto the screen. A second, larger dome (not shown) encases the spinning parts for safety.




Voxels in Glass Plates

DEPTHCUBE, the other interactive volumetric display, was developed by LightSpace Technologies in Norwalk, Conn. It is a rear-projection monitor shaped somewhat like a computer terminal, with a face that measures 16 by 12 inches. The “screen” is four inches thick and made of 20 vertical, transparent glass plates sandwiched together, yet the imagery looks about 12 inches deep. Collaborators standing randomly in front of the screen will each see objects with the corresponding perspective. Internal structures appear and disappear as a viewer’s angle changes. The system would be handy for, say, a team of product engineers studying how parts drafted on a computer-aided design system would—or would not—fit together.

When LightSpace president Alan Sullivan built his first prototype eight

DEPTHCUBE: THICK-SCREEN IMAGING

An arc lamp shines light through optics and a prism, breaking it into blue, red and green beams that are reflected by light-processor chips through a projector toward relay mirrors. The mirrors direct the beams onto one of 20 screens. Each screen consists of two glass plates that contain conducting layers and a liquid-crystal mixture that scatters light. By illuminating the four-inch-deep set of screens in succession, the system creates a 3-D image that appears 12 inches deep. The sample image of half a human head allows particular structures to be seen [sinuses are bright yellow; cartilage is orange], whereas others can be hidden.



years ago, he managed to coax a trio of Texas Instruments DLP chips into projecting depth-related imagery onto the 20 plates, which are separated by thin air gaps. In DepthCube, each DLP contains 786,432 mirrors that tile an area similar to that of a fingernail.

Sullivan still needed a convenient way to generate depth information, and he was elated when he realized that an affordable, commercial 3-D graphics card could suffice. Graphics cards use a color buffer—a morsel of memory—to assign the appropriate color to every pixel on a two-dimensional screen. But the cards also have a hidden component, called the depth buffer, which describes the depth of every pixel. In a normal application, the depth buffer is largely untapped, because only the frontmost layer of a pixel must be defined to create a 2-D picture. So a place for Sullivan's depth information, he muses, "was in there for

free." The information drives the 20 plates, known as liquid-crystal scattering shutters, which can rapidly change from a transparent state to a scattering state. That trait allows a plate to let pixels pass through to other plates as needed yet also enables it to display a pixel. At any moment, all the plates are blank except for one, but the processors project coordinated image slices 50 times a second onto every plate, creating full depth, height and width.

The prototype DepthCube conveyed three-dimensionality to a viewer nicely but only within the four-inch depth of the screen; items in an image appeared almost like flat scenery elements on a

theater stage, standing in front of and behind one another. That was when Sullivan, who formerly studied ultrahigh-energy lasers at Lawrence Livermore National Laboratory, had a smart attack that won him a patent. It occurred to him that the so-called antialiasing algorithms used to smooth jagged edges in 2-D images could also be applied to smoothing the transitions between the DepthCube's 20 planes. This innovation makes the display's 15.3 million physical voxels look like a whopping 465 million virtual voxels. "We produce 31 subplanes between the physical planes, so the perceived resolution is much higher," Sullivan explains. As a result, to the human brain, the images can appear to be as much as 12 inches deep.

The image data that are fed to the chips can come from nearly any 3-D software that runs the OpenGL application programming interface, a common protocol used by computer-aided design and engineering programs such as Catia or ProEngineer. LightSpace has sold a handful of DepthCubes to research institutions, including the U.S. Air Force Research Laboratory and Hokkaido University in Japan, for about \$50,000 apiece. Sullivan acknowledges that the market is limited at this price but says he can see the path to a product that would cost about \$5,000. "There's nothing in our architecture that's different from what's in a rear-projection TV, except for the liquid-crystal shutters," he says, "and those could be produced quite cheaply in volume."

The products developed by these two young companies are garnering respect from boffins in the 3-D world, and more applications will follow. Optical scientist Steve Hines, owner of HinesLab in Glendale, Calif., states that "both these groups are doing extremely hard things and have pulled them off." The natural places to sell the technology, he adds, will be where the money is: "medicine, the military and the movies." ■

MORE TO EXPLORE

Volumetric 3D Displays and Application Infrastructure. Gregg E. Favalora in *Computer*, Vol. 38, No. 8, pages 37–44; August 2005.

A Method for the Real-Time Construction of a Full Parallax Light Field. K. Tanaka and S. Aoki in *Stereoscopic Displays and Virtual Reality Systems XIII*. Edited by A. J. Woods et al. *Proceedings of the SPIE*, Vol. 6055, Article 605516; January 30, 2006.