

A 'room' with a 'view'

To match virtual reality to real tasks, researchers built a smoothly functioning walk-in system mostly from off-the-shelf components

The media's confidence that virtual reality has already arrived is a little premature (they seem to believe that a set of videogame goggles and gloves is all that is needed). But as recent experience at the Electronic Visualization Laboratory (EVL) of the University of Illinois at Chicago has shown, the technology still needs a good deal of work.

In pursuit of a practical virtual reality (VR) system, researchers at EVL had to develop sophisticated software applications, and real-time networks to link advanced high-speed computers with a variety of high-tech peripherals (such as sound synthesizers and location trackers). In so doing, they had to solve technical problems that had limited the usefulness of such systems.

REAL WORK. The benefits make the development worth the effort. When perfected, virtual reality systems may enhance how people work and play, contriving comfortable cybernetic environments that enliven and accelerate education and scientific modeling, in addition to devising new forms of recreation. However, VR is not child's play, as is shown by the Chicago laboratory's struggle to get the discrete parts of such systems to function flawlessly to

Greek philosopher explored the ideas of perception, reality, and illusion, using the analogy of a person facing the back of a cave alive with shadows that are his only basis for his ideas of what real objects are.

The laboratory's Cave is a new model for the design of virtual reality systems, one that offers several advantages over existing models. Unlike users of the video-arcade type of virtual reality system, Cave "dwellers" do not need to wear helmets, which would limit their view of and mobility in the real world, nor don bulky gloves and heavy electronics packs, nor be pushed about by movement-restricting platforms, to experience virtual reality. Instead, they put on a pair of lightweight "glasses" and walk into the Cave, a 27-cubic-meter room with an open side and no ceiling.

The Cave is in fact a partial cube, with the top and one vertical side missing. The other three vertical sides are 3-by-3-meter rear-projection screens facing the viewer, while the floor is a fourth screen—a front-projection one [Fig. 1]. The screens are wrapped around hard-to-detect cables that form seams at their edges. Behind each screen (or, in the case of the floor screen, above it) is a video projector that shines stereo images onto it as the viewer stands inside the cube.

The glasses trick a user's mind into

With a hand-held tracking wand, the user can, for example, strike images of bells

that one clears and closes and the other closes and clears 120 times per second. Thus each eye sees a slightly different image, so as to create the illusion of three dimensions.

Note that while the glasses alter what the wearer sees on the screens, they do not change the way he or she sees real objects (or other people) in the environment. Thus, when a computer-generated image of an object is projected into the environment, the viewer may be tempted to ask, "Is it live?" (Well, one just has to try to grasp the object; for now, at any rate, one's hand will pass through a projected image.)


The stereo glasses also have an electromagnetic transmitter attached to them that can be tracked by a location sensor. When the viewer moves, the computers that generate the images can determine that a change in viewing perspective is needed, and adjust the images appropriately.

A continuous image across several surfaces can be produced when each projector is controlled by a sophisticated computer graphics workstation. Each workstation calculates the perspective from each eye position, and sends the data to the projectors.

A hand-held tracking wand with its own location trackers lets the viewer manipulate the images projected onto the screen. This the person does by locating an imaginary position in virtual reality and interacting with it when the wand "touches" an image. For example, a user can move an object from one place to another, select various options from a menu, or strike images of bells and cause them to "ring."

An important ingredient in a virtual environment is sound. To generate the right kind of environmental sounds for

[1] The Cave consists of four screens—three 3-by-3-meter walls and the floor—onto which coordinated images are projected. The walls are actually made of a single 9-by-3-meter screen that wraps around thin cables at the room's corners. Liquid-crystal-shutter glasses grant vision briefly and alternately to each eye, allowing the viewer to see true three-dimensional images.



transfers data between these systems within the 8-ms window needed to synchronize images displayed on the Cave's four screens. It can accurately and synchronously broadcast data in real time at the high speeds necessary.

The workstations display stereo images by dividing the VGX frame buffer into two fields, each of which contains data for the image to be seen by the left or right eye; each field equals half the vertical resolution (1280 by 512 pixels). The liquid-crystal-shutter glasses (CrystalEyes from StereoGraphics Corp., San Rafael, CA) alternately allow vision in one eye at the field rate of the displays—60 Hz for each eye or 120 Hz combined. The glasses are synchronized by infrared signals generated by the CrystalEyes controller.

Without the glasses, the wall screens show double images; with them, the viewer gains a striking sense of visual depth and three-dimensional movement. Although the illusion of motion begins, albeit jerkily, in animation at approximately 10 Hz (or 10 frames per second), highly disparate areas flickered noticeably even at 30 Hz (60 Hz for both eyes). To avoid that effect, the update rate was doubled to 60 Hz (120 Hz).

Sound is also provided to enhance the illusion of depth. The use of a sound board and MIDI synthesizers in conjunction with

Flock of Birds location sensor from Ascension Technology Corp., Burlington, VT) monitors where the lead viewer's head is inside the Cave and in what direction it is turned. Information on viewer location is sent to the controlling workstation, which does the computations for image generation and sends the data to the appropriate projectors. The hand-held tracking wand also contains a tracking receiver so the system can determine where the user is holding and pointing it.

MAKING IT WORK. For the multiple projectors to function in stereo, the vertical-scan portion of their video signals had to be synchronized. The 120-Hz screen update rate produced by the Silicon Graphics hardware was incompatible with the projector's image-synchronizing, or genlock, input. To compensate, the Electronic Visualization Laboratory built a simple stereo sync processor (the only hardware employed in the system that was not available off the shelf). Using a signal from the CrystalEyes controller, the custom processor selects every other vertical sync signal emitted by the source VGX and so creates a 60-Hz signal that the slave VGX genlock inputs can handle.

A ScramNet optical-fiber network was employed because Ethernet would not have provided the real-time speeds fundamental to maintaining the virtual reality illusion.

ScramNet replicates this data and transmits it simultaneously to all other parts of the Cave network. A data filter gets rid of redundant transmissions and unchanged data, in effect increasing bandwidth.

One of the largest practical challenges in implementing the Cave was making the seam between floor and wall screens "disappear" into the application. When the viewer looks at a seam between screens where the corners have not been properly displayed, the effect is jarring; the objects look bent and so the illusion breaks down.

ScramNet sets up an interlock through which the workstations can share display data quickly enough and keep screen images enough in sync for the corners to present a continuous illusion of depth. Also, the 3-mm-diameter cable joining the screens physically at their edges minimizes the physical gap between them.

Another early problem was a side effect of the high, 120-Hz video rate. The long decay time of the green phosphors used in the video projectors caused the green channel of the video to persist longer than the others. Because the display alternated from eye to eye before the green component of the images faded, viewers would see double images. Moreover, all images appeared in a hazy green light.

A quick fix was to add a blue

20/110, which is four times better than current head-mounted displays and on a par with binocular omni-oriented monitors (Booms).

The Cave is relatively immune to errors due to head rotation and nodding. With head-mounted displays, the projection plane must move as the viewer's head moves to look left or right, up or down. Therefore, in head-mounted displays, the viewer must wait for the computer to adjust the monitors accordingly after turning, a lag that may cause a loss of equilibrium and even nausea. In the Cave, the projection plane is stable; a glance to the left inside the Cave shows the image already on-screen.

Also, the Cave is a less intrusive interface. A head-mounted display can be uncomfortable and disorienting, because the viewer is cut off from the real visual world. Someone wearing it is blind to the real-world risks of, say, walking into walls or tripping over the cables needed to connect the headset back to the computer. A viewer in the Cave can move around at will, because the outside world is still visible through the liquid-crystal glasses. Better still, because he or she can still see the outside world, the Cave frees the viewer to collaborate more fully.

More than one user will, however, have to interact in the same environment if virtual reality is to become an effective visualization tool in fields such as research and education. While the head-mounted display and the Boom interfaces do allow multiple users to share an environment, it is only at the high cost of duplicating the interface hardware. The participants must themselves be represented with virtual selves in each other's simulated environments, which adds to a system's location-tracking and computational workloads. In other words, in order for me to see you in a head-mounted display environment, the computer must calculate in virtual terms physically in relation to me and render your image accordingly on my headset monitors.

In the Cave, two or more users may look around as they wish. They need not be represented virtually, because they represent themselves physically. At present, though, the screen images are shown from the perspective of only one person, the so-called lead viewer, who is being tracked.

CONTINUING CHALLENGES. Viewer location tracking remains a challenge. A lag between changes in viewer position and the updating of the screen was sometimes noticeable. While less troublesome than the lag experienced with head-mounted displays, it still shows up if the viewer moves too quickly. The lab is looking into extrapolation techniques to better predict user motion and reduce interactive delays to sensor input.

More advanced technology may be capable of maintaining several perspectives separately on-screen, so that each viewer enjoys an individual perspective. However, this technological step is not trivial; with

20th-century Cave paintings

Although virtual reality is basking in its popularity as an entertainment medium, researchers in the field are working with other scientists on several practical applications for this technology. Realistic examples of what the Cave can do are evident from descriptions and photographs of applications in mathematics, biochemistry, and neurosurgery that have already been developed.

Fractal Explorer. In the growing field of chaos science, fractals are mathematical representations of repetitive elements of compound events or objects. The Fractal Explorer, created by Randy Hudson, Alan Verlo, and Louis Kauffman of the University of Illinois at Chicago (UIC), enables participants to explore three-dimensional fractals and their chaotic attractors by moving around them, changing their shapes and colors, and displaying them with different graphical primitives (for example, triangles or squares). Viewers may wave a wand to rotate the fractal and view other sides of it, or to radically revise the fractal by manipulating the attractors it was built on.

A new way of exploring attractor space was also developed in the Cave. It involves anchoring a length of virtual ribbon to the end of the wand. The wand can draw the ribbon anywhere in the Cave's space while the attractor tries to draw the ribbon in by the other end. The effect is somewhat like hanging onto a long strip of cloth to keep it from being sucked into a vacuum cleaner.

The Fractal Explorer allows users to comfortably manipulate visual images that are described by very complex numbers. A less extreme example of similarly useful mathematical manipulations would be in business: viewers could explore and manipulate the data in a vast database, adjusting bar graphs or pie charts to perform "what if" analyses with the wand.

Brain structure modeling. As computer workstations become more powerful and affordable, computer modeling becomes a more attractive way for scientists to witness and understand what goes on at microscopic scales. An example is modeling what goes on in the brain at a neural level: a simulation of the inferior olive and the cerebellum was created by Jim Bower, Upinder Bhalla, Maurice Lee, and Erik DeSchutter of the California Institute of Technology, and Jason Leigh of UIC.

At a networked demonstration during Siggraph '92 in Chicago, the Caltech researchers modeled a cerebellar Purkinje neuron (a large, round nerve cell) using their Intel Touchstone Delta supercomputer in California. At the same time, in Chicago, a Sun Microsystems Sparcstation was used to create a model of the inferior olivary neuron (a cell from nervous tissue located alongside the medulla oblongata, which is responsible for automatic actions such as breathing). These two neuron models were linked using a T3 line over NSFnet and the data generated during their interaction was displayed in the Cave.

With the use of the hand-held wand, participants could stimulate the olivary neuron; this in turn triggered an action potential that was transmitted to the Purkinje cell. This model not only demonstrated ways in which researchers could comfortably see and learn what goes on at a microscopic level, but also how the Cave (and other virtual reality paradigms) might work in connecting viewers thousands of kilometers apart.

Mapping cognitive function. During surgical procedures involving the frontal or parietal cortex, brain surgeons must be able to identify and avoid critical motor, sensory, and speech areas. Because these functions are never organized in the same way from one individual to the next, researchers are trying to find ways to easily map the location of these functions.

A way of doing this was demonstrated at Siggraph. Created by Leo Towle, Robert Grzeszczuk, and Martin Ryan of the University of Chicago and Steve Cohen of UIC, it takes the areas of the brain associated with cognitive processes—such as language, attention, and memory—and maps them onto three-dimensional images of the brain surface created from magnetic-resonance imaging (MRI) scans.

The voltage potentials evoked in the brain of a patient to whom an image of a checkerboard is shown are measured by 21 electrodes over a 10-ms interval. The evoked-potential range is mapped onto a color scale. The colors, when transferred to the MRI image of the brain, indicate the level of activity in that portion of the brain. Further, each electrode is assigned the sound of a different musical instrument, whose pitch represents the intensity of activity. The music produced, with its complex instrumental and rhythmic components, lets medical researchers analyze events using the common human ability to recognize complex sound patterns.

Virtual reality models in brain research are to the advantage of both patients and science. An anatomically correct three-dimensional image, one that surgeons and researchers can walk around and poke, does not disturb a live subject. It also allows a more quantitative evaluation of how cognitive processes are organized within the human brain.


—T.A.D., D.J.S., and C.C.N.

each additional viewer, the display update rate increases, and more brightness is required from the projector.


Despite these difficulties, the Cave has drawn long lines of eager would-be visitors at industry shows since it was first demonstrated at the showcase event at Siggraph '92, held in Chicago. The applications demonstrated have included practical and educational programs, such as molecular

biology models, superconductor design, fractal mathematics, weather mapping, and environmental management. Some programs were simply entertaining. For example, some visitors had the opportunity to play music with the hand-held wand by ringing bells floating around them.


Other useful applications of the Cave model of virtual reality include scientific visualization, medicine, architecture, and art.



The Fractal Explorer turns complex math concepts into 3D objects.



The interaction of the Purkinje and inferior olivary neurons are examined by means of supercomputers and virtual reality.



Cognitive functions are mapped onto an image of the brain taken from magnetic-resonance imaging systems.

In 1994, the Electronic Visualization Laboratory plans to return to Siggraph with an event called Vroom (Virtual Reality Room). The intent is to display the latest in head-mounted displays, binocular omni-oriented monitor, and Cave technologies. The lab is currently asking others for applications to display in the Cave.

At the same time, EVL is also working on the creation of a "road-show" Cave,

which would be easier to transport. While this one-of-a-kind working model costs about US \$600 000 in all, it could eventually become the prototype for home- or business-based Caves that might cost no more than a "home theater" does today. ♦

ABOUT THE AUTHORS. Thomas A. DeFanti is director of the Software Technologies Research Center and codirector of the Electronic Visualization

Laboratory at the University of Illinois at Chicago. His current interests are interactive systems and computer graphics languages.

Daniel J. Sandin is codirector of the EVL and a professor at the university's School of Art and Design. His current interests include interactive computer graphics and rendering of defraction.

Carolina Cruz-Neira is a Ph.D student at the EVL. Her interests include interactive computer graphics and virtual reality.