

The Cave

Audio Visual Experience Automatic Virtual Environment

**Carolina Cruz-Neira
Daniel J. Sandin
Thomas A. DeFanti
Robert V. Kenyon
John C. Hart**

The CAVE is a new virtual reality interface. In its abstract design, it consists of a room whose walls, ceiling and floor surround a viewer with projected images. Its design overcomes many of the problems encountered by other virtual reality systems and can be constructed from currently available technology. Suspension of disbelief and viewer-centered perspective, are often used to describe such systems.

Suspension of Disbelief: This term arose from film criticism and is defined as the ability to give in to a simulation—to ignore its medium. The early attempts of the entertainment industry to achieve better suspension of disbelief laid the foundations for current virtual reality research. Suspension of disbelief is a fundamental part of the effective use of a virtual reality interface. Until we can ignore the interface and concentrate on the application, virtual reality will remain a novel experience instead of a serious visualization tool.

Viewer-Centered Perspective: The perspective simulation of common visualization systems dates back to the Renaissance, and is based in a mythical camera positioned along an axis extended perpendicular from the center of the screen. Viewer-centered perspective simulates the perspective view from the location of the viewer. To maintain correct perspective, a sensor that continuously reports the viewer's position to the simulation is commonly used. Without this perspective, the viewer becomes less a part of the environment, and a full suspension of disbelief becomes increasingly difficult.

Paradigms

Research in virtual reality began in 1965, when Ivan E. Sutherland proposed the "Ultimate Display," which would completely override the user's senses, totally immersing the user in the computer simulation [16]. Modern virtual reality research has split into four distinct directions, based primarily on differences in display devices.

Cathode Ray Tube (CRT): This is the simple monitor and is the most basic visual paradigm for virtual reality, though other kinds of monitors are also used. The most




Figure 1.
Figure 3.

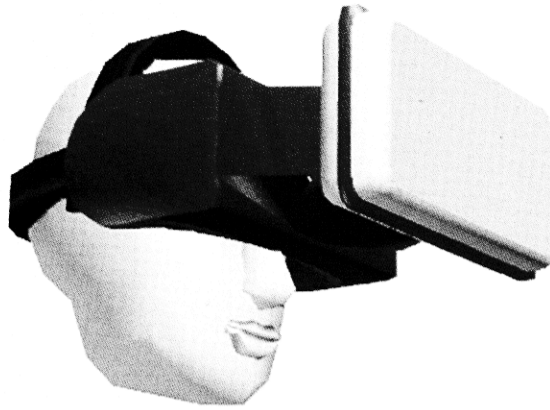



Figure 2.
Figure 4.



broadly used system is a small monoscopic display whose perspective is based on the Renaissance mythical camera model. This minimal visual interface may be enhanced for virtual reality use with the addition of stereo and viewer-centered perspective [10].

Head-Mounted Display (HMD): This is one of the most popular virtual reality visual interfaces. It consists of a stereo pair of small displays that cover the eyes. A head-tracking device provides the location and orientation of the viewer to simulate the correct view [4, 5, 11, 17].

Binocular Omni-Oriented Monitor (BOOM): Like the head-mounted display, the BOOM mounts small displays in front of the eyes, though more like binoculars than like goggles. Unlike the head-mounted display, the BOOM is suspended from an articulated arm, which measures its position and orientation in space and counterbalances its mass. Moreover, the user is expected to hold and position the BOOM manually throughout the virtual reality experience.

Audio-Visual Experience Automatic Virtual Environment (CAVE): This is the fourth visual paradigm for virtual reality and is a recursive acronym, also reminiscent of Plato's allegory of the cave [12]. The CAVE is a cube with display-screen faces surrounding a viewer. It is similar to surround systems such as OMNIMAX theaters and early flight simulators [14, 15]. Its more recent instance is coupled with a head-tracking device. As the

viewer moves within the bounds of the CAVE, the correct perspective and stereo projections of the environment appear on the display screens.

Immersion Issues:

Immersion is the degree of visual simulation a virtual reality interface provides for the viewer—the degree of the suspension of disbelief. The five main issues in creating a powerful suspension of disbelief are shown in Table 1.

Field of View: The field of view represents the visual angle a viewer can see without head rotation. The simplest formulation, using W as the width of the display and D as the distance from the viewer to the display, is derived as a single angle

$$\theta = 2 \tan^{-1} \frac{W}{2D}$$

which describes the horizontal visual angle.

The field of view is variable in the CRT paradigm and is based on the size of the CRT and the viewer's distance from it. Viewing a 19-inch diagonal CRT from 18 inches produces a 45° field of view. The field of view for each eye is fixed in the BOOM and HMD paradigms. In the HMD paradigm, field of view angles from 100° to 140° are common. The Fake Space BOOM allows the viewer to adjust the field of view anywhere from 90° to 120°.

The field of view of the CAVE display varies by viewer location for each individual screen but achieves a full 360° for the entire display. Hence, the viewer experiences a maximal field of view, which may be limited by display hardware such as stereo glasses.

Cinerama and IMAX theaters fall in the CRT interface paradigm. These noninteractive virtual reality visual interfaces create a larger field of view by increasing the size of the projection screens. Cinerama used three vertical-edge linked projection screens and three synchronized projectors, whereas IMAX uses one very large screen which is placed near the viewing audience.

The success of both of these systems in enhancing the suspension of disbelief characterizes the dependence of virtual reality on field of view.

Panorama: This is the ability of a display to surround the viewer and is crucial in creating a sense of immersion. It differs from field of view in that head rotation is used to view panorama. The CRT paradigm is not well-suited for panorama, and is generally treated as a window into some virtual environment. The OMNIMAX theater creates a sense of panorama by placing a large hemispherical screen about the viewer. The sense of panorama is strong in the BOOM and HMD interfaces, simulating everything the viewer sees. Viewer rotation is fast and smooth in the BOOM interface, due to its mechanical rotation-sensing equipment, though the inertia of the BOOM limits the rate of rotation.

The HMD is light, compact and easy enough to move quickly. Hence, the viewer can alter position and orientation much faster than present day tracking equipment. The result is a distracting lag: when the user turns, the environment turns with the user and then moves back to the correct orientation. Users of such systems are forced to move quite slowly and smoothly to avoid this problem.

The CAVE solves this problem by showing all views from a fixed location simultaneously. Users of the CAVE experience the same viewer location and head rotation measurement delays as do users of the HMD, but since rotations only require a small alteration to the stereo projections the effect is less noticeable.

Viewer-Centered Perspective: This depends heavily on the speed and accuracy of viewer location sensing. The viewer-centered perspective of each of the CRT, HMD and CAVE environments suffers from the same delay problems due to slow viewer location sensing. Hence, viewers compensate by moving slowly. The mechanical



Figure 1. Cathode ray tube (CRT)

Figure 2. Head-mounted display (HMD)

Figure 3. Binocular-omni-oriented monitor (BOOM)

Figure 4. CAVE Audio visual experience automatic virtual environment (CAVE)




Figure 5. Representation of what viewer would see if he or she had visual acuity of a) 20/20, average eyesight b) 20/40, CRT interface c) 20/85, the Boom interface d) 20/425, the HMD interface e) 20/110, the CAVE interface. The lines on the eyechart are from top to bottom, 20/200, 20/100, 20/80, 20/60, 20/40, 20/20.

position-sensing of the BOOM overcomes this delay with faster location reporting at the expense of high inertia, which prevents any abrupt location changes where a delay would be noticeable.

Body and Physical Representations: Interaction in a virtual environment often requires a visual representation of the body of the viewer, particularly the hands. Additionally, other physical equipment in the interaction area may appear in the virtual environment. In the BOOM and HMD interfaces, the senses are completely restricted to the computer simulation. Here, body representation is explicit—the body must be simulated and rendered like any other kind of geometry in the environment. This requires detailed body-part position measurement and the allocation of extra rendering time.

In the CRT and CAVE environments, body representation is implicit—the body appears physically and does not require rendering. This also means the visual interface cannot alter the presence of the

viewer's body. Furthermore, the body or any other physical object will occlude a virtual object even if the virtual object is closer to the viewer.

Virtual reality applications focusing specifically on body representation are the video art installations that process the image of the viewer and reproduce it in some synthesized environment. Examples are "Videoplace" [9] and E. Tannenbaum's "Recollections" video installation, on permanent display at the Exploratorium in San Francisco.

Intrusion: The intrusiveness of a virtual reality interface indicates the severity of its restriction of the senses. The HMD interface has the highest intrusion, completely isolating the viewer from the real environment. Using half-silvered mirrors allows a viewer with an HMD to see the real environment superimposed by objects in the virtual environment. Such modifications decrease the intrusion caused by the HMD, but also severely reduce the field of view. The BOOM similarly isolates the viewer, but the viewer may easily remove the interface.

The CRT and CAVE interfaces are nonintrusive. In such an environment, the viewer is free to move at will, secure in the awareness of the real, as well as the virtual, aspects of the environment.

Visualization Issues

For virtual reality to become a more

useful tool, we must evaluate its visualization effectiveness. The five main visualization issues are identified in Table 2.

Visual Acuity: The quality of a computer graphics display is commonly measured by its *resolution*—the number of pixels, or individual points it uses to produce an image. The quality of a virtual reality interface is more properly measured by a combination of both resolution and field of view. This measurement is called the *visual acuity* of the display—the portion of a pixel taken from the center of the display that spans one minute (1/60th of a degree) of the field of view. Symbolically, a horizontal resolution of H pixels across a display W inches wide has a pixel pitch of $P = W/H$ inches per pixel. Given a viewer distance of D inches from the center of the display, the angle a single pixel creates on the retina is approximately $\tan^{-1} P/D$ and is measured in minutes. The visual acuity, the portion of a pixel that subtends an angle of one minute on the retina, is simply the inverse $1/\tan^{-1}(P/D)$.

The popular Snellen fraction, used to measure vision, is another unit for reporting visual acuity. A viewer whose visual acuity is indicated by the Snellen fraction 20/ X means the viewer can see at 20 feet what a viewer with average eyesight can see at X feet. For example, a Snellen fraction of 20/20 is average, meaning a visual angle of one min-

ute is perceptible, 20/10 is above average with a minimum perceptible visual angle of one-half minute and 20/40 is below average with a minimum perceptible visual angle of two minutes. Legally blind, for example, is the accepted term for vision that cannot be corrected to better than 20/200. Residents of Illinois need a visual acuity of 20/70 or better to drive in the daytime, and at least 20/40 to drive at night. The Snellen fraction, when divided, produces the correct visual acuity as previously defined.

The following acuity measurements use the maximum (horizontal) dimension resolution whereas actual vision research tends to prefer the minimum (vertical) dimension resolution. The resolution of CRT's is considered relatively high, commonly $1,280 \times 1,024$ pixels. A viewer at a distance of 18 inches from a 19-inch-CRT ($\approx 15''$ horizontal) creates a pixel pitch of 0.0117 inches per pixel. The visual acuity of the CRT is $1/\tan^{-1}(0.0117/18) = 0.448$ pixels per minute giving the viewer 20/45 vision, which is almost good enough to drive at night.

The visual acuity of the HMD interfaces is currently limited by LCD technology. The fish-eye optics for our example, the LEEP CYBERFACE 2, complicate visual acuity computations. However, the angle subtended on the retina by a pixel from the center of the display is specified as 0.0062 radians or 21.3 minutes. This infers a visual acuity of $1/21.3 = 0.0469$ pixels per minute giving the viewer about 20/425 vision, which is undoubtedly legally blind. The LEEP optics improve the poor resolution of LCD. Other HMD interfaces with an equal field of view but lacking the resolution enhancements from the LEEP optics score an even worse acuity.

The resolution of the Fake Space BOOM interface is currently about $1,000 \times 1,000$ (in black and white) with flexible screen widths to trade acuity for field of view. For a narrow 90° field of view, the BOOM

screen width, coupled with the LEEP optics, generates center pixels subtending an angle on the retina of 0.00127 radians or 4.37 minutes. This infers a visual acuity of $1/4.37 = 0.229$ pixels per minute, giving the viewer about 20/85 vision, which, except for the limited field of view, is almost good enough to drive.

Finally, the resolution of an individual CAVE screen is the same as the CRT, 1,280 pixels over 7 feet for a pixel pitch of 0.00547 feet per pixel. The viewer in the default-centered position is 3.5 feet from the center of the display. Hence, the CAVE has a visual acuity of $1/\tan^{-1}(0.00547/3.5) = 0.186$ pixels per minute, giving the viewer about 20/110 vision, which is better than legally blind but is not sufficient for even daytime driving.

Linearity: Often, the field of view and resolution of a display are increased through optical devices. These devices increase the field of view by bending the light from the displays, in effect curving a flat display around the viewer. Resolution is increased by concentrating more pixels into a small central area of

the display, while leaving the edges of the display less well defined. Without such optics, the visual acuity of flat screens is worse in the center of the display and sharper at the edges—exactly the opposite of what our visual system needs.

Some manufacturers of HMD and BOOM devices use optics such as the LEEP systems [8]. These displays improve resolution at the center of the screen by as much as a factor of 2.9 and expand the field of view for each eye up to 140° .

These transformations also cause distortions that bend straight lines. The distortions are easily modeled and the inverse distortion can be computed and applied to the image to reduce the effect, though at a significant degradation in simulation speed.

Look Around: This is the ability to move about an object, viewing it from different angles. It is a useful property when using the virtual reality interface for visualization.

The CRT paradigm lacks a strong *look around* capability due to the size of its screen and its low field of view. When viewed from the side, the visible area of the CRT

Table 1.
Immersion Issues

	Field of View	Panorama	Perspective	Body Rep.	Intrusion
CRT	45°	None	Slow	Physical	None
BOOM	$90^\circ \leftrightarrow 120^\circ$	Fast	Fast	Virtual	Partial
HMD	$100^\circ \leftrightarrow 140^\circ$	Slow	Slow	Virtual	Full
CAVE	Full	Fast	Slow	Physical	None

Table 2.
Visualization Issues*

	Vis. Acuity	Linearity	Look Around	Prog. Refine.	Collab.
CRT	20/45	Linear	Limited	Fix Loc. only	Single Persp.
BOOM	20/85*	LEEP	Full	Fix Loc. and Rot.	Dup. Hardware
HMD	20/425	Either	Full	Fix Loc. and Rot.	Dup. Hardware
CAVE	20/110	Linear	Full	Fix Loc. only	Single Persp.

*At 90° field of view, black and white.

becomes much smaller, severely limiting the viewing angles from which an object is visible. One solution to this problem is to have the CRT rotate in order to always face the viewer.

The BOOM and HMD interfaces handle *look around* since they simulate everything the viewer sees. The CAVE also provides a sense of *look around*. These interfaces require some kind of virtual travel to *look around* distant objects.

Progressive Refinement: This is the ability to dynamically increase the computational expense of a model during a pause in viewer response [3]. The standard scheme simulates a fast coarse model for viewer interaction, then computes a much finer model when the viewer remains still. Coarse versus fine attributes are model resolution, such as the number of points or polygons, and rendering techniques, like adding shadows or progressive radiosity.

The HMD interface requires the viewer to remain absolutely still to refine the display. The BOOM also requires zero movement of the interface, but the high inertia and nonintrusiveness of the BOOM make this much easier than the HMD. The CRT and CAVE interfaces require only the viewer's location to remain fixed. Hence, in the CAVE in particular, the viewer is allowed to pan around during refinement.

Collaboration: One of the most important aspects of visualization is communication. For virtual reality to become an effective and complete visualization tool, it must per-

mit more than one user in the same environment.

The BOOM and HMD interfaces allow multiple users in their environment at the high cost of duplicating the interface hardware [1]. The CRT and CAVE environments allow multiple users to benefit from the experience without modification. In such a situation, the perspective accommodates only one of the viewers. If shuttered glasses are used for stereo, then the CRT and CAVE interfaces can simulate the correct perspective for all users, though n users would see the screen only $1/n$ th of the time, requiring a fast scan rate and a very bright image.

CAVE Implementation

At the time of this writing, our implementation of the CAVE uses two projection screens (two walls)—five screens (three walls, a ceiling and a floor) are expected for the Showcase '92 exhibition. The implementation of the CAVE interface requires computation of viewer-centered perspective projections, deployment of viewer tracking equipment, synchronization of displays, and overcoming any resulting projector and tracking limitations.

Viewer-Centered Perspective

The CAVE requires special perspective projections to simulate viewer-centered perspective. These projections are offset to simulate stereo, and thus require knowledge of the viewer's orientation.

Off-Axis Perspective Projections: The viewer-centered perspective,

as well as the projections used for stereo, are derived from the off-axis perspective projection [7]. The simplest derivation alters a standard on-axis perspective projection by two affine transformations. First, points are *sheared* in a direction parallel to the projection plane, by an amount proportional to the point's distance from the projection plane (points in the projection plane remain unchanged). Then, points are scaled along the axis perpendicular to the projection plane by an amount again proportional to the point's distance from the projection plane (and again points in the projection plane remain unchanged). Adding stereo consists of bifurcating this projection into two similar projections differing by opposite shears along the axis of disparity—the line through the two eyes of the viewer.

The Need for Orientation: The viewer's head must be oriented. There are two reasons for this and they both involve correct stereo projection. In theatrically released three-dimensional films the viewer's head is assumed to be vertical. In the CAVE, one may want to tilt one's head. Unless the viewer's rotation about the line of sight is accounted for, one's head could tilt 180° to find an inverse stereo effect, or 90° to find no stereo effect at all. These concerns become paramount when considering stereo projection onto the floor and ceiling.

The position of the viewer's eyes is needed to prevent inconsistencies at the edges of the CAVE walls. If stereo is computed assuming the user is looking perpendicular to the

One of the most important aspects of visualization is communication. For virtual reality to become an effective and complete visualization tool, it must permit more than one user in the same environment.

projection planes, the stereo disparities will not line up at the edges where two projection planes meet.

Display Hardware

Real-time rendering of the virtual world is achieved through six Silicon Graphics Inc. VGX workstations, each attached to a rear projection display. A Silicon Graphics "Personal Iris" serves as a master controller for the system and all workstations communicate via Ethernet.

Multiple Stereo Displays: These workstations display stereo, using the StereoGraphics "CrystalEyes." StereoGraphics divides the VGX frame buffer into two half-vertical-resolution fields, one for each eye. The user wears liquid crystal glasses that shutter at the field rate of the displays, synchronized by an infrared signal. At a rate of 60Hz (30Hz per eye) the display flickered noticeably at highly disparate areas. Hence, the update rate was doubled to 120Hz.

Multiscreen stereo requires synchronizing the video signals. The 120Hz screen update rate produced by the StereoGraphics hardware was not compatible with the VGX genlock input. The only prototype hardware unit used in the system, a stereo sync processor, fixed this problem. This processor filters out every other sync signal emitted by the source VGX, creating a 60Hz signal that the slave VGX genlock inputs could handle.

The Green Problem: An unexpected side effect of the 120Hz video rate was a lag in the green channel of the video. This is due to the long decay time of the green phosphors in the video projectors. The display would alternate from eye to eye faster than the green phosphors could handle. The resulting double image in the green channel caused a complete loss of stereo depth perception.

One immediate solution was to work in the hue space spanned by the red and blue axes which reduced brightness to 41%. Another solution was to fix green at some

level, say 50% which increased the brightness and allowed the full hue space at the expense of one brightness per hue and reduced contrast. We are now using faster green phosphors that completely solve the green problem.

Viewpoint Tracking: The position and orientation of the user's head is obtained with a 3SPACE Polhemus "Isotrack" sensor, whose transmitter is mounted on the StereoGraphics glasses. As expected, there is a noticeable sensing lag, which manifests itself during fast viewer motions. As stated earlier, this is not a problem for viewer rotation, but remains a problem when moving. Extrapolation techniques are currently being investigated to better predict user motion to reduce the interactive delays due to sensor lag.

Conclusion

The CAVE is a nonintrusive easy-to-learn high-resolution virtual reality interface. It is superior to other virtual reality paradigms across many issues, particularly in field-of-view, visual acuity and lack of intrusion. Moreover, it is open to limited use for collaborative visualization.

Applications: SIGGRAPH '92 Showcase

Several applications of the CAVE will be featured at Showcase. These include applications from the Electronic Visualization Laboratory, and others.

Regional-Scale Weather in Three Dimensions: This application is from the work of A. Campbell at the Argonne National Laboratory. It uses the PSU/NCAR Mesoscale Model in a parallelized form, running interactively on the Intel Touchstone Delta, to create a three-dimensional display of weather systems over a region of North America.

Graphical Planning for Brain Surgery: Brain-surgery-planning software, featured by R. Grzeszczuk, is currently undergoing clinical testing at the University of Chicago. It employs a three-dimensional local-

izer as means of interactively transferring spatial relationships from MR-derived three-dimensional anatomical models directly onto the patients.

The Visible Embryo: The viewer is taken on a trip through a human fetus via a simulation developed by L. Sadler and the Biomedical Visualization Laboratory at the University of Illinois at Chicago, providing a unique view of the human body that could aid in medical developments.

The Snowstorm: This project visualizes three-dimensional vector fields using interactive particle systems where small points traverse the vector field, each at speeds proportional to the magnitude at that point in the field. Predefined vector fields are provided as well as the ability to "comb" new vector fields interactively using a wand.

Fractal Exploratorium: A virtual laboratory of fractals and chaotic attractors is presented in this application by R. Hudson of the Electronic Visualization Laboratory. Participants can investigate chaotic forms from a variety of different perspectives, interactively altering their parameters, and hence, their shapes.

Bio Modeling: The interactive modeling of biological macromolecules is demonstrated in this application from K. Shulten of the Beckman Institute at the University of Illinois at Urbana-Champaign.

The Evolving Universe of Galaxies and Stars: A combination of stored database images and real-time computations from a remote CRAY will allow the viewer to fly through an evolving universe in this application developed by M. Norman at the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign.

Further Research

Most of the problems with the CAVE are a consequence of hardware shortcomings, such as lag time due to tracking delays, the green

problem due to insufficient phosphor decay times, and multiprojector stereo synchronization. Moreover, brighter screens with faster update rates would allow multiple-viewer-centered perspective.

The effectiveness of virtual reality interfaces, particularly the CAVE, need to be evaluated more quantitatively. The degree of immersion an interface creates as well as its ability as a visualization tool are difficult quantities to obtain and deserve a much more thorough treatment than given here.

Acknowledgments

We would like to thank Gary Lindahl and Robert Grzeszczuk for technical support and assistance. Graduate art students Michelle Miller, Kathy Koller and Lewis Siegel created the system diagrams.

Gary Oberbrunner of Thinking Machines Inc. simultaneously suggested the extrapolation technique. Thanks to Sumit Das for an implementation of this method.

The visual acuity and linearity issues used information graciously provided by Eric Howlett at LEEP Systems, Inc., for which we are grateful. We would also like to thank the Brenart Eye Care Center of Yorkville, Ill. for its assistance with the eye chart diagrams and vision assessment.

This publication, and the research and conclusions made thereby, were supported in part by a grant awarded by the Illinois Department of Commerce and Community Affairs. Representations made by this publication do not necessarily reflect the opinions and conclusions of the department.



References

- Blanchard, C., Burgess, S., Harvill, Y., Lanier, J., Lasko, A., Oberman, M. and Teitel, M. Reality built for two: A virtual reality tool. *Comput. Graph.* 24, 2 (Mar. 1990), 35–36.
- Brooks, F.P. Grasping reality through illusion: Interactive graphics serving science. In *Proceedings of SIGCHI '88* (May 1988), pp. 1–11.
- Brooks, F.P., Ouh-Young, M., Batter, J.J. and Kilpatrick, P.J. Project GROPE: Haptic displays for scientific visualization. *Comput. Graph.* 24, 4 (Aug. 1990), 177–185.
- Chung, J.C., Harris, M.R., Brooks, F.P., Fuchs, H., Kelley, M.T., Huges, J., Ouh-Young, M., Cheung, C., Holloway, R.L. and Pique, M. Exploring virtual worlds with head-mounted displays. In *Proceedings of SPIE*, Vol. 1083, Feb. 1990, pp. 42–52.
- Fisher, S. The AMES Virtual Environment Workstation (VIEW). *SIGGRAPH '89 Course #29 Notes*, Aug. 1989.
- Fisher, S. Viewpoint dependent imaging: An interactive stereoscopic display. In *Proceedings of SPIE*, Vol. 367, Feb., 1982.
- Hodges, L.F. Time multiplexed stereoscopic computer graphics. *IEEE Comput. Graph. Appl.* 12, 2 (Mar. 1992), 20–30.
- Howlett, E.M. Wide angle orthostereo. In *Proceedings of SPIE*, Vol. 1256, Feb. 1990, pp. 210–223.
- Krueger, M.W. *Artificial Reality II*. Addison-Wesley, 1991.
- Lippman, A. Movie Maps: An application of the optical videodisc to computer graphics. *Comput. Graph.* 14, 3 (1980).
- Pausch, R. Virtual reality on five dollars a day. In *Proceedings of SIGCHI '91*, 1991, pp. 265–270.
- Plato. *The Republic*. The Academy, Athens, circa 375 BC.
- Rheingold, H. *Virtual Reality*. Summit, N.Y., 1991.
- Rolfe, J.M. and Staples, K.J. *Flight Simulation*. Cambridge Univ., 1986.
- Schachter, B.J. Computer image generation systems. In *Computer Image Generation*, B.J. Schachter, Ed., Wiley & Sons, N.Y., 1983.
- Sutherland, I.E. The ultimate display. In *Proceedings of IFIP 65*, 2, pp. 506–508, 582–583.
- Teitel, M.A. The Eyephone: A head-mounted stereo display. In *Proceedings of SPIE*, Vol. 1256, Feb. 1990, pp. 168–171.
- Venolia, D. and Williams, L. Virtual Integral Holography. Tech. Rep. 90-10, Apple Computer Inc., Feb. 1990.

CR Categories and Subject Descriptors: I.3.2 [Computer Graphics]: Graphics Systems—*Distributed/network graphics*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—*Virtual reality*

General Terms: Human Factors

About the Authors:

CAROLINA CRUZ-NEIRA is a Ph.D student at the Electronic Visualization Laboratory. Her current research interest is in interactive computer graphics and virtual reality.

DANIEL J. SANDIN is codirector of the Electronic Visualization Laboratory and a professor in the School of Art and Design at the University of Illinois at Chicago. His current research interests are in interactive computer graphics and rendering of diffraction. email: sandin@siggraph.org

THOMAS A. DEFANTI is codirector of the Electronic Visualization Laboratory, professor in the department of electrical engineering and computer science, and director of the UIC Software Technologies Research Center. His current research interests are in interactive systems and computer graphics languages. email: tom@siggraph.org

ROBERT V. KENYON is an associate professor in electrical engineering and computer science at the University of Illinois at Chicago. His research interests include sensory-motor adaptation, visuomotor control, effects of microgravity on vestibular development, flight simulation, computer graphics, and digital signal processing using parallel computers. He is currently working on problems of modeling complex visual/motor coordination in the human. email: kenyon@uicbert.eecs.uic.edu

JOHN C. HART holds a joint postdoctoral research associate position at the Electronic Visualization Laboratory, University of Illinois at Chicago and the National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign. His current research interests are in the modeling and rendering of fractals. email: hart@uicbert.eecs.uic.edu

Authors' Present Address: Electronic Visualization Lab., EECs Dept. M/C 154, Univ. of Ill., Chicago, IL. 60680-4348; fax: 413-7585

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© ACM 0002-0782/92/0600-064 \$1.50