THREE-DIMENSIONAL DISPLAYS

The human visual system perceives and comprehends the world in three dimensions by using both physiological and psychological depth cues. Physiological depth cues include accommodation (change in focal length of the eye lens), convergence (inward or outward rotation of the eyes), binocular disparity (differences between left and right eye images), and motion parallax (image changes due to motion of the observer).

Psychological depth cues include linear perspective (distant objects appear smaller), shading and shadowing (indicate positions relative to light sources), aerial perspective (distant objects appear less distinct), occlusion (nearer objects hide more distant objects), texture gradient (distant objects have less detail), and color (distant objects appear darker). A computer graphics display device must provide some or all of these depth cues in order to present a semblance of three dimensions (see Table 1).

Workers in the fields of computer graphics and volume visualization have created illusory three-dimensional (3D) images and scenes on two-dimensional (2D) display screens, typically cathode ray tubes (CRTs), by computing and displaying psychological depth cues. These images lack the physiological depth cues supplied by an actual 3D object, are limited to only a single angle-of-view, and require significant computation to "render" the depth cues (calculate perspective; remove hidden lines and surfaces; add shading, lighting, and shadows; etc.).

Techniques that include one or more physiological depth cues improve 3D perception, but suffer from their own limitations. For example, stereoscopic CRT approaches (in which the left-eye image is presented to the left eye only, while the right-eye image is presented to the right eye only) add limited stereopsis (simulating binocular disparity), but still lack motion parallax and large angles of view and require rendering twice, once for each eye.

Head-tracking technologies added to stereoscopic approaches (head-mounted display devices) provide motion parallax and angles of view, with the added benefit of unlimited display volume. They still suffer, however, from the need to render in silicon (i.e., compute the image) for each eye. An additional drawback is the current physical intrusiveness of the technology (e.g., bulky head gear).

The varifocal mirror display devices pioneered by Sher (1) provide increased angle of view, but are currently limited to a single color and a single view point. Computer-generated holography provides horizontal parallax and a moderate horizontal angle of view, but no vertical parallax (a computational limitation versus a fundamental one), in only very small display volumes and with a huge computational burden. The holographic approach requires rendering and calculating the holographic interference patterns for each individual angle-of-view (2,3).

Instead of trying to simulate various depth cues in an attempt to trick the eye into seeing 3D, direct volume display devices (DVDDs) (also called volumetric displays) actually generate light in a physical 3D space. The generated photons then propagate from their points of origin to the observer's eyes, just as photons from real objects do. The complete set of physiological depth cues is achieved, and several of the psychological depth cues also result.

Table 1 lists various display technologies and shows the depth cues they are able to generate. Table 2 shows the technologies organized by the techniques used to achieve the depth cues.

In Table 2, *monocular* means a single view is generated and presented to both eyes of the observer. *Binocular* means that the eyes of the observer are presented different views. *Stereoscopic* means that some gadgetry is used to present different views to the eyes of the observer, with the aim of generating depth cues. *Autostereoscopic* means that the observer seeing a different view with each eye is an intrinsic property of the device.

Table 1. Depth Cues Available with Various Display Technologies

Display Technology	Cathode Ray Tube	Stereoscopic CRT	Head-Tracking Technology	Varifocal Mirror Device	Computer Generated Holography	Direct Volume Display Devices
Psychological Depth Cues						
Linear perspective	Yes	Yes	Yes	Yes	Yes	Yes
Shading and shadowing	Yes	Yes	Yes	Limited	Yes	Limited
Aerial perspective	Yes	Yes	Yes	No	No	No
Occlusion	Yes	Yes	Yes	No	Yes	Limited
Texture gradient	Yes	Yes	Yes	Limited	No	Limited
Color	Yes	Yes	Yes	No	No	No
Physiological Depth Cues						
Accommodation	No	No	No	Yes	Yes	Yes
Convergence	No	Yes	Yes	Yes	Yes	Yes
Binocular disparity	No	Yes	Yes	Yes	Yes	Yes
Horizontal motion parallax	No	No	Yes	Yes	Yes	Yes
Vertical motion parallax	No	No	Yes	Yes	No	Yes

Table 2. Techniques Used to Generate 3D Displays

Monocular
Computer graphics (CRTs, flat panels, projection systems)
Ray tracing
Volume rendering
Hidden surface/Hidden line removal
Binocular
Stereoscopic
Stereoscope
Screen shutter and glasses
Colored
Polarized
Active shutter glasses
Autostereoscopic
Holography
Vibrating mirror
Head tracking/Helmet-mounted devices
Direct volume display devices
Swept volume
Passive screen
Symmetric double helix
Laser reflection
Nonsymmetric single helix
Laser reflection
Nonsymmetric double helix
Laser reflection
Vertical plane
Laser reflection
Active screen
Vertical plane
Light-emitting diodes
Light-emitting phosphors
Static volume
Discrete medium
Continuous medium

In Table 2, DVDDs are classified as *binocular* because they present different views to each eye of the observer, and they are classified as *autostereoscopic* because no special gadgets are required for the observer's eyes to perceive separate views. They are termed *direct volume display devices* because they generate light directly in 3D (volume) space, as opposed to faking it via mirrors and/or complex rendering on CRTs.

COMPUTER GRAPHICS RENDERING

The process of creating 2D images of 3D objects for depiction on a computer graphics screen using psychological and physiological depth cues is called *rendering.* As has been discussed by Krueger (4), rendering techniques are often based on an evaluation of the linear Boltzmann equation from transport theory. This equation describes the gains and losses in intensity as energetic particles traverse a volume. A very simplified form of the equation suitable for this discussion is

$$\partial I/\partial s = A - B \times I + C + D \tag{1}$$

where

I = intensity
s = direction of (position along) the ray
A = intensity source
B = opacity (absorption plus total scattering)
C = particle energy changing interactions
D = particle direction changing interactions

The terms A, B, C, and D describe the physics of the particles being transported. These terms are typically functions of position in the volume, position along the ray, and particle energy.

Rendering via the Boltzmann equation involves first mapping the variables of the abstract voxel space into the various parameters hidden by the simplicity of Eq. (1), and then integrating the equation along the rays through each pixel on the screen to the eye of the observer. The psychological depth cues (shading, shadowing, aerial perspective, occlusion, etc.) can be generated through proper manipulation of the terms on the right-hand side of Eq. (1).

According to Kaufman (5), ". . . the ultimate highly inspirational goal in equipment development [for volume visualization] is a novel 3D display technology or media for fast presentation of 3D volumes, as well as surfaces, in any arbitrary direction." Direct volume display devices approach this goal by displaying 3D volumes and surfaces in a volume, providing "depth rather than depth cues" (1).

DIRECT VOLUME DISPLAY DEVICE RENDERING

One of the defining characteristic of direct volume display devices is that they do not perform rendering by computation. Once the values of the abstract voxels are known, a DVDD merely translates each abstract voxel into intensity and/or color and then appropriately illuminates the corresponding voxel in the 3D display volume. Real particles (photons) traverse the volume to all observers where any required processing is performed by their visual systems.

For DVDDs the equivalent linear transport equation would be

$$\partial I/\partial s = A \tag{2}$$

DVDDs do not perform the integration of Eq. (2) in silicon. Conceptually, DVDDs perform the integration by *nature,* by generating light in the voxels and allowing it to propagate through the display volume to the observer. This means that the integration is performed at the speed of light, simultaneously for all rays going to each eye of the observer, including perspective projection, and simultaneously for multiple observers viewing the same display device at the same time. As observers change their positions, their views are all automatically and simultaneously updated in real time. DVDDs currently provide no capabilities at all for terms B, C, and D of Eq. (1). The point is that DVDDs can generate light in a voxel, but cannot currently absorb, scatter, reflect, or refract light in a voxel.

Even with the limitations of Eq. (2), it is still possible to simulate some of the effects of light sources—for example, geometric shadows and diffuse lighting via Lambert's law. This requires first ray tracing from light sources to surfaces in the abstract voxel space, then mapping the results of the surface brightness computations to the source term A. Surfaces can be displayed by detecting them in the abstract voxel space (see, for example, Refs. 6–8) and then mapping surface voxels to the source term A.

Any visual effect which involves the terms from Eq. (1) that are missing in Eq. (2) (i.e., *B*, *C*, and *D*) or requires knowledge of the position of the observer relative to the display volume cannot be produced. This includes almost all hidden surface and hidden volume representations.

A SURVEY OF DIRECT VOLUME DISPLAY DEVICES

As of 1998, operational DVDDs involve generating rapidly changing 2D images on a rotating screen which sweeps through the 3D volume. The motion of the screen is fast enough that the observer sees only a dim blur from it, with his visual system instead focusing on the light generated during the entire sweep. These devices effectively *slice* a physical cylindrical 3D volume with the rotating 2D screen and illuminate that screen with corresponding slices from a 3D data set. Some important performance characteristics of eight prototype DVDDs are listed in Table 3. These *swept volume* DVDDs can be further classified as to whether the screen is *passive* (consisting of a simple reflective surface) or *active* (containing embedded light emitting elements or material).

Passive Screen Swept Volume Devices

Four basic approaches have been pursued for implementing *passive screen swept volume* DVDDs. These approaches differ mainly in the shape of the screen utilized. All four use screens that merely reflect light generated elsewhere in the device.

Nonsymmetric Single-Helix Screen. Brinkman (9) briefly described an experimental device built at the IBM Heidelberg Scientific Center. This prototype used a nonsymmetric single helix illuminated from above by a single laser beam. *Helix* is defined as a cylindrical surface whose elevation is constant along any radius but increases uniformly with angle about the axis. A *nonsymmetric single helix* is a helix whose surface sweeps upwards in one rotation, as shown in Fig. 1(a). No performance data were presented by Brinkman, and apparently the effort was abandoned.

The Naval Research and Development Test and Evaluation Division (NRaD) of the Naval Command, Control and Ocean Surveillance Center (10) has produced a system that uses a nonsymmetric single-helix surface and three lasers to illuminate it from below (see Table 3 and Fig. 2).

Table 3. Performance Parameters for Eight Direct Volume Display Devices Employing the Swept Volume Approach

Development Organization	Device Designation	Display Surface	Rotation Rate (rot/s)	Effective Display Volume (cm^3)	Display Screen Diameter (cm)	Display Screen Height (cm)	Voxel Size (mm)	Volume Frame Rate (Hz)	Volume Image Resolution	Number of Discrete Colors	Maximum Number of Displayed Voxels $\times 10^3$
Raytheon/Texas Instruments, Inc.	OmniView™ (third generation)	Non-symmetric double helix	10	293,000[a]	91	46	2.5	10–20	370 × 370 × 185	3 (RGB) with mixing	11.6
Raytheon/Texas Instruments, Inc.	OmniView™ (fourth generation)	Passive vertical plane	20	51,100	51	25	1.0	40	512 × 512 × 256	1 (G)	1.4
NRaD	3D Volumetric Display (first generation)	Symmetric double helix	10	2,050[a]	33	15	1.0	20	256 × 256 × 256	2 (RG) with no mixing	8.0
NRaD	3D Volumetric display (second generation)	Symmetric double helix	10	24,600[a]	91	46	1.7	20	256 × 256 × 256	4 (RGBY) with mixing	120
NRaD	Transportable 3D volumetric display	Non symmetric single helix	20	26,400	41	20	0.7	20	256 × 256 × 256	3 (RGY) with mixing	80
LAMDA Systems Corp.	LAMDA	Passive vertical plane	15	25,400	30	36	1.0	30	256 × 256 × 256	1 (R)	0.6
Volumetric Imaging, Inc.	Matrix imager	Active vertical plane	20	39,600	41	30	~2	20	64 × 64 × 48	1 (R)	790
University of Canterbury	Cathode ray sphere	Active vertical plane	25–30	3,220	16	16	~1	25	256 × 256 × 256	3 with mixing	150

Discrete colors indicated are as follows: R, red; G, green; B, blue; Y, yellow.

[a] Display volumes are reduced to account for the central dead zone, or drawing is purposely limited to a volume on one side of the cylinder to avoid the central dead zone.

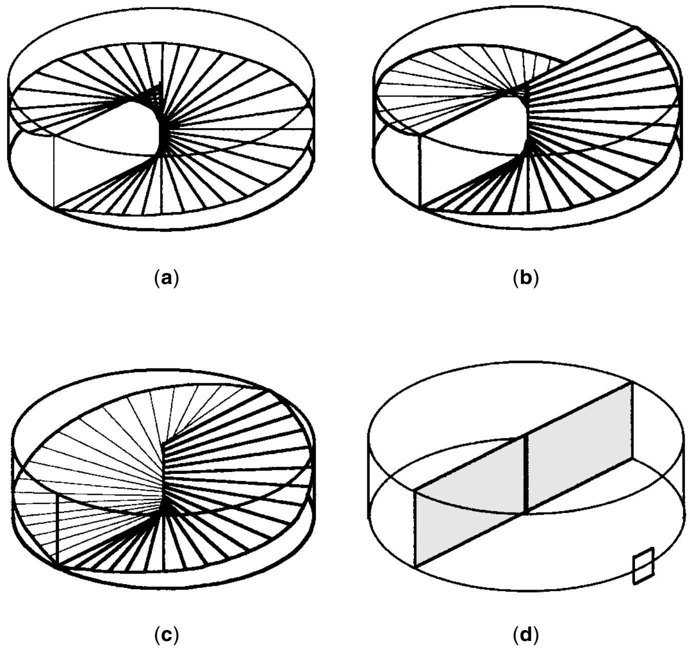

Figure 1. Four display surfaces used in direct volume display devices employing the swept volume approach. (a) Nonsymmetric single helix. (b) Symmetric double helix. (c) Nonsymmetric double helix. (d) Vertical plane. The three helix surfaces are indicated by a series of horizontal radial lines through the surfaces. As each surface is rotated about its vertical axis, the surface sweeps through all points in the cylindrical volume. Illuminating a spot on the surface as that spot passes through a desired 3D point causes the observer to see a voxel (3D pixel) floating in space at that point. Rotating at a speed of up to 20 rotations per second, the surface itself disappears into a blur. The helix surfaces are illuminated by scanned laser beams directed from either above or below by mirrors. The vertical plane surface may be illuminated by scanned laser beams directed up the rotation axis, outward to the edge of the volume, then back to the surface by small mirrors. Alternatively, the vertical surface may be covered by light-emitting diodes or coated by a phosphor material and scanned by electron beams.

Advantages of the nonsymmetric single-helix surface are as follows:

- Voxels are illuminated at the same rate (once per rotation).
- Voxels may be illuminated from either above or below.

Disadvantages of the nonsymmetric single-helix surface are as follows:

- The surface is mechanically nonsymmetric, hence unstable in rotation.
- Voxels can be illuminated only once per rotation, requiring a higher rotation rate to avoid flicker.
- Voxels to be illuminated must be sorted into the order in which the single-helix surface intercepts their positions.
- The surface has a dead zone at the center of rotation because of the steepness of the surface and the axle which may be necessary. A dead zone is an area in which voxels cannot be drawn.
- The surface may have other dead zones depending on the geometry of the illumination source.
- The surface is self-occluding; that is, there are areas in the display which the surface itself blocks the observer from seeing.
- The actual voxel shape on the surface changes with distance from the axis of rotation.
- The perceived voxel shape changes with position in the volume and the observer's viewing direction.
- Some vectors may not be accurately drawn because all voxels comprising them need to be illuminated simultaneously (e.g., a horizontal radial line).

NRaD illuminated their display surfaces with acousto-optically scanned laser beams (see ACOUSTO-OPTICAL DEVICES). There are also advantages and disadvantages associated with this illumination technique.

Advantages of lasers for voxel illumination include the following:

- Lasers provide high brightness and well-formed circular beams.
- Lasers are easily scanned in direction and modulated in brightness.
- Multiple colors can be provided by multiple lasers with different wavelengths.
- Multiple colors can also be provided by color mixing (by illuminating voxels simultaneously with two or more laser beams).

Disadvantages of using lasers include the following:

- Acousto-optic scanners have limited bandwidth which limits the number of voxels that can be displayed. NRaD uses beam splitters and multiple acousto-optic scanners in parallel to increase the number of displayable voxels.
- Voxels suffer from speckle due to the coherent nature of laser light.
- Alignment problems make color mixing prone to position errors.

Symmetric Double-Helix Screen. NRaD (10–13) has produced two other prototype devices: a laboratory setup that uses a symmetric double helix and two lasers to illuminate the surface from the side and a system that uses a symmetric double helix and three lasers to illuminate the surface from above. The same advantages and disadvantages of using lasers listed above apply. A *symmetric double helix* is a helix whose surface sweeps upwards in one half-rotation, drops vertically back to the bottom of the cylinder, then sweeps upwards again in the second half-rotation, as shown in Fig. 1(b).

Advantages of the symmetric double-helix surface are as follows:

- The surface is axially symmetric, hence stable in rotation.
- Voxels are illuminated at the same rate (twice per rotation).
- Voxels may be illuminated from either above or below.

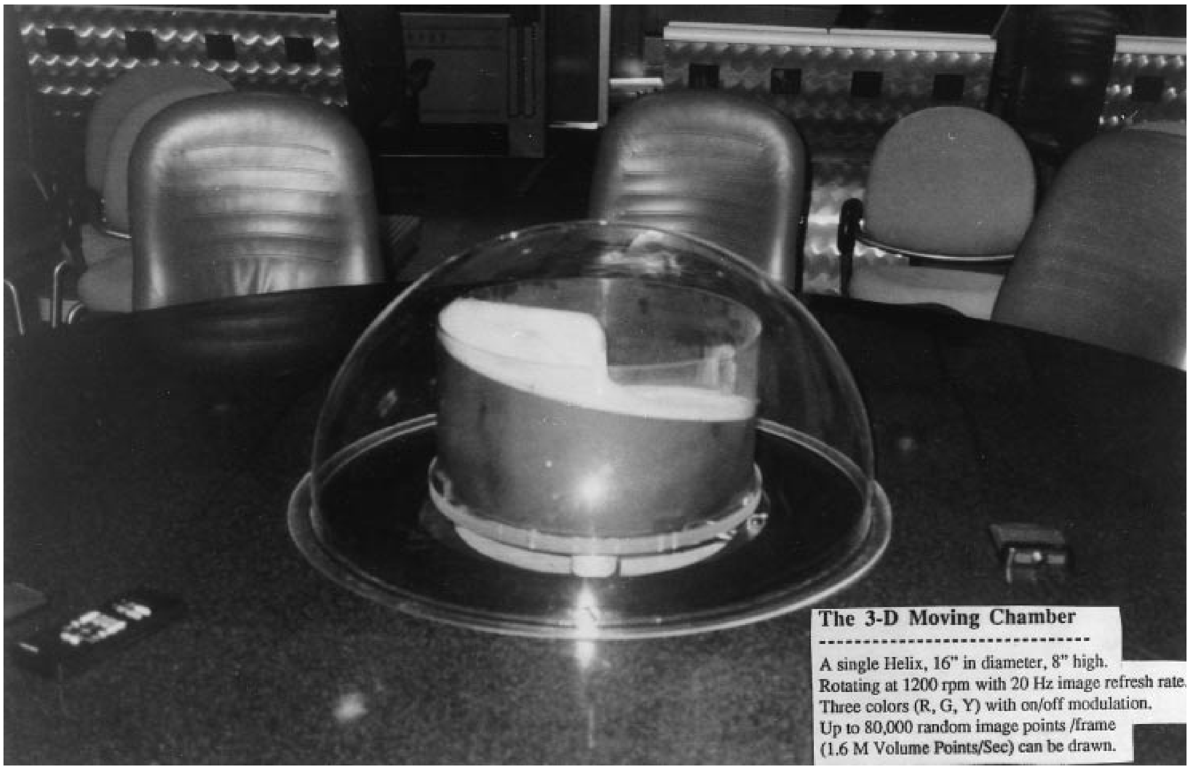

Figure 2. The transportable 3D volumetric display developed by researchers at NRaD is an example of a direct volume display device using a passive nonsymmetric single-helix surface. It is illuminated from below by two (red and green) lasers. The 41 cm diameter helix is 20 cm high and rotates at 20 rotations/s. The resulting volume frame rate (frequency of updating the entire volume image) is 20 Hz. Up to 80,000 voxels can be simultaneously displayed. (Courtesy of Parviz Soltan of "SPAWAR Systems Center—San Diego," "Simulation and Human Systems Technology Division, Code 44.")

Disadvantages of the symmetric double-helix surface are as follows:

- Voxels to be illuminated must be sorted into the order in which the double-helix surface intercepts their positions.
- The surface has a dead zone at the center of rotation because of the steepness of the surface and the axle which may be necessary. NRaD uses only a portion of the volume in their largest display device so that the central dead zone is avoided.
- The surface has other dead zones depending on the geometry of the illumination source (14).
- The surface is self-occluding.
- The actual voxel shape on the surface changes with distance from the axis of rotation.
- The perceived voxel shape changes with position in the volume and the observer's viewing direction.
- Some vectors may not be accurately drawn because all voxels comprising them need to be illuminated simultaneously.

Nonsymmetric Double-Helix Screen. During the early 1990s, Texas Instruments, Inc. (now Raytheon/Texas Instruments) produced two versions of DVDDs that used nonsymmetric double-helix surfaces [see Fig. 1(c)] illuminated from above or below by one or more acousto-optically scanned lasers (15–19). While Raytheon/Texas Instruments (RTI) has since withdrawn from this area of technology development, because of their historical importance, their two most recent display devices are mentioned here. A *nonsymmetric double helix* is a helix whose surface sweeps upwards in the first half-rotation, then sweeps downwards in the second half-rotation, as shown in Fig. 1(c).

Advantages of the nonsymmetric double-helix surface are as follows:

- Voxels may be illuminated from either above or below.
- Self-occluding areas are smaller than with the symmetric double-helix surface.
- Each voxel is illuminated twice per rotation.

Disadvantages of the nonsymmetric double-helix surface are as follows:

- The surface is mechanically nonsymmetric, hence unstable in rotation.
- Required extra bearings may generate considerable acoustic noise.
- Voxels to be illuminated must be sorted into the order in which the double-helix surface intercepts their positions.
- The voxel illumination rate varies with the vertical coordinate of the voxel.

- The surface has a dead zone at the center of rotation because of the steepness of the surface and the necessary axle (14).
- The surface may have other dead zones depending on the geometry of the illumination source.
- The surface is self-occluding.
- The actual voxel shape on the surface changes with distance from the axis of rotation.
- The perceived voxel shape changes with position in the volume and the observer's viewing direction.
- Some vectors cannot be accurately drawn because all voxels comprising them need to be illuminated simultaneously.

RTI, like NRaD, illuminated their voxels with acousto-optically scanned laser beams. The same advantages and disadvantages of using lasers listed above apply. RTI used a nonsymmetric helix in their third-generation device primarily because of the smaller self-occluding areas. In the envisioned applications (military situation displays) it was very undesirable for information to be displayed in the volume that the observer might not be able to see if the observer happened to be in the wrong position. The penalty for this approach was a surface that was mechanically nonsymmetric, requiring extra mechanical bearings to counteract its tendency to turn over, and these bearings resulted in higher acoustic noise.

Vertical Plane Screen. RTI built a fourth-generation device using a flat vertical plane illuminated by a single green laser reflected from a small mirror at the edge of the cylinder, as shown in Fig. 1(d). The small mirror rotated with the vertical plane (18,19). The image was transmitted up the rotational axis of the system via a dove prism in the hollow axle. This meant that the image rotated at one-half the rate of the display surface.

Advantages of the vertical plane surface are as follows:

- The surface is very easy to manufacture.
- Self-occluding areas do not exist.
- There are essentially no dead zones.
- The surface is nearly mechanically symmetric, hence stable in rotation.
- Voxels to be illuminated must only be sorted into the order in which the vertical plane intercepts their positions.
- Voxels can be illuminated twice per rotation.
- The voxel illumination rate is constant.
- The actual voxel shape on the surface changes only very slightly with distance from the axis of rotation.

Disadvantages of the vertical plane surface are as follows:

- The perceived voxel shape changes with position in the volume and the observer's viewing direction.
- Laser scanning and transmission up the axle of the device must compensate for the image rotation effects of the dove prism.

RTI again illuminated their voxels with an acousto-optically scanned laser beam. The same advantages and disadvantages of using lasers listed above apply. A fixed spherical mirror in a band around the base of the display volume was proposed by Shimada (20) to eliminate the small mirror which RTI rotated at the edge of the cylinder. LAMDA Systems Corp. (21) used an approach similar to RTI's fourth-generation device, with a phosphor grid added to the display surface to provide a reference grid.

Active Screen Swept Volume Devices

Two approaches have been pursued for implementing *active screen swept volume DVDDs*. These approaches differ mainly in the method of generating light on the surface. Both use screens that are rotating vertical planes, as shown in Fig. 1(d), but without the small mirror.

Volumetric Imaging, Inc. (22,23) used a vertical screen consisting of an active matrix of light emitting diodes (LEDs). This approach lends itself to a much higher degree of parallelism and has the advantage of being able to illuminate every voxel in the volume.

Advantages of the vertical plane LED surface are as follows:

- The surface is relatively easy to manufacture.
- Self-occluding areas do not exist.
- There are essentially no dead zones.
- The surface is nearly mechanically symmetric, hence stable in rotation.
- Communications with the rotating components are electrical, as opposed to optical; hence graphics memory and processors may be located behind the vertical plane and rotating with it.
- Voxels to be illuminated must only be sorted into the order in which the vertical plane intercepts their positions.
- Voxels once sorted into 2D slices can be stored in rotating memory.
- The voxel illumination rate is constant.
- Electrostatic scanning is fast.
- The actual voxel shape on the surface does not change with distance from axis of rotation.

Disadvantages of the vertical plane LED surface are as follows:

- The perceived voxel shape changes slightly with position in the volume and the observer's viewing direction.
- Voxels can be illuminated only once per rotation, requiring a higher rotation rate to avoid flicker.

Researchers at Canterbury University (New Zealand) have built experimental systems that use rotating vertical plane phosphor-coated screens enclosed inside CRTs (24–26). The screen is illuminated by two electrostatically scanned electron beams (see Fig. 3).

Advantages of the vertical plane phosphor surface are as follows:

- The surface is easy to manufacture.
- Self-occluding areas do not exist.
- The surface is mechanically symmetric, hence stable in rotation.

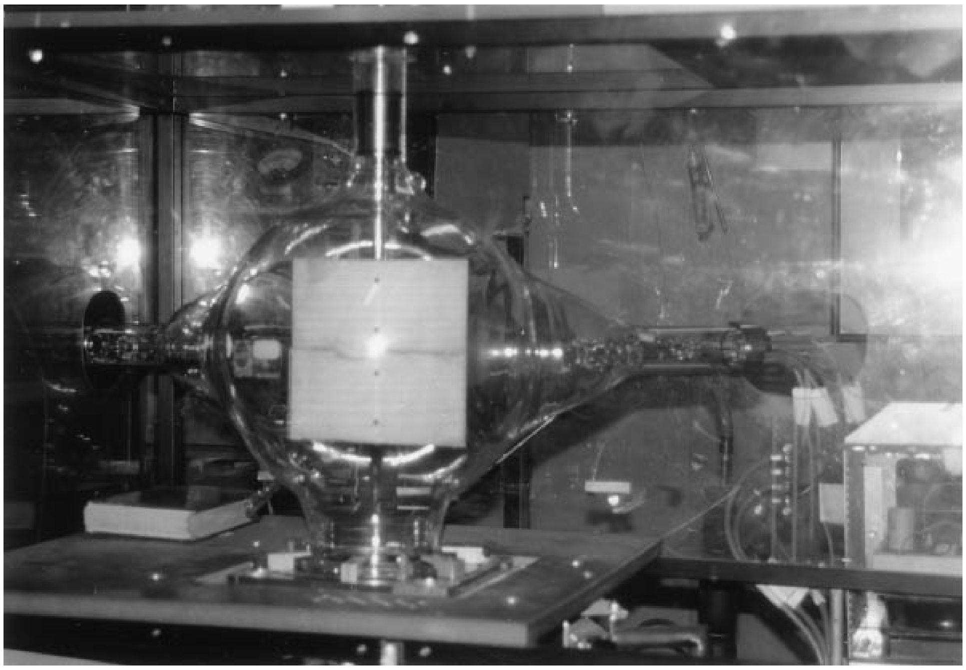

Figure 3. The cathode ray sphere developed by researchers at the University of Canterbury is an example of a direct volume display device using an active vertical screen surface. It is illuminated by scanning the phosphor-coated screen with two electron beams. The 16 cm^2 vertical plane rotates at 15 rotation/s inside an evacuated glass tube. The resulting volume frame rate is 25 Hz to 30 Hz. Up to 150,000 voxels can be simultaneously displayed. (Courtesy of Dr. Barry G. Blundell.)

- The voxel illumination rate is constant.
- Electrostatic scanning is fast.

Disadvantages of the vertical plane phosphor surface are as follows:

- The actual voxel shape changes with position on the surface and the angle of the surface relative to the electron beam source.
- Dead zones occur due to this changing voxel shape (voxel elongation). This problem has been greatly mitigated by using two electron guns separated by 90° and 120° in the two prototype devices, allowing the vertical plane to be addressed by whichever electron gun is more nearly perpendicular to the plane (14,27).
- Voxels to be illuminated must be sorted into the order in which the vertical plane intercepts their positions, and they must also be sorted with regard to which electron beam is to be used to illuminate them.
- The perceived voxel shape changes with position in the volume and the observer's viewing direction.
- Voxels can be addressed only once per rotation, requiring a higher rotation rate to avoid flicker.
- The rotating surface must be enclosed within a vacuum tube for proper functioning of the electron beams.

Static Volume Devices

An alternative to the *swept volume* approach is the *static volume* DVDD in which no physical moving screen is involved. The idea is to avoid all the mechanical, vacuum, and scanning problems associated with moving screens by using a nonmoving, all solid-state imaging chamber. That this is a difficult challenge is indicated by the fact that, while static volume DVDDs have been the subject of discussion and research since the late 1950s, no practical operational display device has yet been produced.

Static volume displays may be further classified as to whether the display medium is a *continuous medium* (solid, liquid, or gas) or a *discrete medium* (a 3D array of discrete emitters). Research in static volume DVDDs is briefly discussed below.

Discrete Medium. Alburger (28) briefly discussed an early prototype (called the "Electronic Crystal Ball") that used *electroflor materials*—that is, materials that become fluorescent or show visible colors at low voltages. The electroflor material was deposited on glass plates which were then stacked to provide the dimension of depth. The prototype had an addressability of only 10 × 10 × 4. Little information was provided on the performance parameters of this prototype device.

Researchers at General Electric (29) experimented with liquid crystal cells stacked in parallel. The resulting display volume was a cube 19 cm on each side. The addressability appears to have been 20 × 20 × 10 voxels with a volume frame rate of 2.5 Hz to 12 Hz. This embodiment of a static volume DVDD is interesting because a liquid crystal voxel in its clear state is transparent, while in its other state it is both visible and opaque. Little information was provided on the performance parameters of this prototype device.

Researchers at the University of Texas (30) experimented with voxels consisting of nuggets of an ultraviolet-cured optical resin doped with an organic dye and mated to thin glass sheets stacked vertically to provide a 3D array of voxels. Each voxel was attached to an optical fiber that carried ultraviolet light which pumped the dye, causing it to fluoresce at a visible wavelength. They constructed a prototype array of 11 × 11 × 5 voxels in a 300 cm^3 volume. Again little information was provided on the performance parameters of this prototype device.

Continuous Medium. Verber (31) discussed sequential excitation of fluorescence (SEF) as a method of directly generating light in voxels. The technique involves scanning infrared (hence invisible to the observer) laser beams in a medium that fluoresces at their intersection. Two circular laser beams may be used (31,32) to excite fluorescence at their intersection. An alternative approach uses a plane of laser light which sweeps through the volume while an orthogonal circular beam draws on that plane.

NRaD (32) has explored the use of ZBLAN, a fluorescent glass, for use in an SEF-based display device. Their prototype device used a small block of ZBLAN glass $7.5 \times 7.5 \times 3.7$ cm and two circular cross section infrared laser beams directed by galvanometer mirror scanners. The voxel diameter was 0.1 mm. Approximately 1000 voxels could be displayed.

Finally, researchers at Stanford University, in cooperation with the U.S. Navy and several corporations (32,33), have demonstrated red, green, and blue fluorescence in a single cube of ZBLAN approximately 1 cm on a side.

Discussion of Static Volume Devices. As stated above, no practical operational static volume display device has yet been produced; however, research continues in this area on many fronts, including the following:

- Finding a better display medium to provide brighter voxels, larger display volumes, and color mixing
- Developing faster scanning devices for laser pumping of *continuous media* to provide more voxels
- Index of refraction matching among the components within *discrete media* to reduce or eliminate internal reflections and refractions
- Index of refraction matching between the atmosphere and the display volume to reduce or eliminate image distortions

Perhaps the most serious drawback to all current prototypes is that the observer is looking into a block of glass. Even if research on the first three fronts listed above is successful, refraction effects depending on the position of the observer and the location within the display volume of the object being displayed will be noticeable and objectionable, as anybody who has ever looked into a tropical fish tank can attest. This problem would appear to be solvable only by a display medium with an index of refraction very close to unity.

The exciting possibility of the static volume approach is that a medium might be found that can be excited to emit light and also absorb light. Such a medium, assuming that the other problems can be solved, holds the possibility of displaying solid objects as solid objects, as opposesd to translucent objects as do swept volume DVDDs.

A Commercially Viable DVDD

Direct volume display devices have already proven useful in certain specialized applications. The goal of DVDD research is, however, the development of a commercially viable DVDD that would be useful in many application areas. Such a device would need to have approximately the characteristics listed below for the reasons given in parentheses:

- Voxel refresh rate of 40 Hz (to avoid flicker)
- Image update rate of 10 Hz (to provide relatively smooth dynamics)
- Display volume of at least 100,000 cm^3 with essentially no dead zone (to provide a useful volume)
- Voxel size of 1.0 mm or less (to provide crisp images)
- Number of voxels and voxel addressability of at least $512 \times 512 \times 256$ (67 million voxels) (to provide adequate image detail)
- At least 64 colors (6 bits) (to provide object differentiation)
- Number of displayable voxels in each color of at least 100,000 voxels (to handle complex images)
- Voxel throughput rate of at least one million voxels/s (to support the above parameters)
- Office environment operation (brightness, size, power, cooling, noise, etc.)
- Application program interface based on industry standards (to make it easy to program)

As in many areas of computer technology, more tends to be better. Hence more colors would be better, more displayable voxels would be better, and so on. The parameter values listed above are intended for use in assessing whether or not current technologies can result in a commercially viable DVDD. As such, variations of 50% or more in the listed numeric values are probably of no consequence, since the real requirements are the subjective ones listed in parentheses.

COMPUTER GRAPHICS TECHNIQUES

Having reviewed the hardware aspects of DVDDs, we turn now to topics related more to firmware and software. Two techniques have been used to represent 3D data in computer graphics: the *boundary representation* and the *volume representation.* In the boundary representation (abbreviated B-rep) the 3D nature of objects is represented by specifying the boundaries between volumes via a geometric model. In the volume representation (abbreviated V-rep), the 3D nature of objects is represented by dividing the 3D space into small cells and specifying the properties of the material that fills each cell.

Kaufman (5) has described the image display process of a traditional 2D computer graphics display device processing V-rep data as involving the following eight steps:

- 2D enhancement (image processing on 2D slices of the image)
- Reconstruction (interpolation in the third dimension)
- 3D enhancement (3D image processing)
- Manipulation (geometric and domain transformations)
- Classification (finding surfaces in the volume)
- Mapping (mapping the 3D voxel data into display primitives)
- Viewing (projecting the display primitives onto the 2D screen)
- Shading (shading the 2D projection)

For processing B-rep data the first three steps are replaced by a single process called *voxelization,* in which B-rep data are converted to V-rep data. The strength of the traditional approach is that a huge amount of volumetric data is rendered down to one 2D scene; however, even a simple change in perspective requires an entire recomputation to render the scene from the new point of view.

Like traditional computer graphics display devices, DVDDs must be able to handle application data in either representation. The DVDD graphics pipeline is relatively straightforward. After the 3D enhancement step, all voxels are mapped to the color (or intensity) and addressability of the display volume. The mapped voxel dataset is then divided into the appropriate 2D slices, pipelined through a data-to-light conversion, and recombined in the display volume. While eliminating the computational burden of rendering from 3D down to 2D, DVDDs have increased the amount of information that must be transferred to the displayed image.

DVDD Image Display

DVDD image display can be divided conceptually into a number of distinct processes, characterized primarily by the data conversion performed by each (34). The approach here is to start with the application data, which may be in either representation (B-rep or V-rep), convert to a single representation (V-rep), and then delineate the remaining processes required to generate the final displayed volume image. In the following presentation, what each process does is first listed, then discussed in a little more detail.

1. Voxelization and mapping process
   a. Receives image data from the application
   b. Converts B-rep data to V-rep data (if required)
   c. Maps V-rep data to the addressability of the volume image
   d. Sends the volume image to the image deconstruction process via the image interface
2. Image deconstruction process
   a. Buffers the volume image received from the voxelization and mapping process via the image interface
   b. Slices the 3D volume image into 2D electrical frames
   c. Sends the frames to the illumination process via the frame interface
3. Illumination process
   a. Receives frames from the image deconstruction process via the frame interface
   b. Converts 2D electrical frames to 2D light frames
   c. Sends the light frames to the image reconstruction process via the display interface
4. Image reconstruction process
   a. Receives 2D light frames from the illumination process via the display interface
   b. Displays 2D light frames to reconstruct the 3D image

Voxelization and Mapping Process. The voxelization and mapping Process (VMP) receives image data from the application via the application program interface (API). The VMP maps received geometric data (e.g., B-rep data from OpenGL) and/or received volumetric data (e.g., V-rep data from a magnetic resonance imaging system) into the DVDD-specific voxel space. The result of this process is a set of 3D arrays of voxels sized appropriately for the given DVDD, with appropriate color and/or intensity values for each voxel. The VMP communicates with the image deconstruction process across the image interface.

Conceptually, the VMP could execute on clients which would use the image interface to request 3D display services from the DVDD. Given the addressability and color characteristics suggested above, a full displayed volume image (i.e., all 67 million addressable voxels) would exceed 50 Mbytes. Updating a 50 Mbyte image at 10 Hz is beyond current local area network (LAN) capabilities. Ethernet can support about 1 Mbyte/s, and the fiber distributed data interface (FDDI) can support about 10 Mbytes/s, but neither approaches the required rate of 500 Mbyte/s.

Current DVDD displayed volume images, however, tend to be sparse; that is, only a small percentage of the voxels are illuminated. This property can be exploited in the image interface. For example, with only 100,000 voxels to be illuminated, it is more efficient to transfer the Cartesian coordinates and colors $\{x, y, z, c\}$ of only the voxels that need to be illuminated. With 6-bit color (i.e., 2 bits each for red, green, and blue), this reduces the image to less than 400 kbytes, which can be transferred at 10 Hz via FDDI. Note that as applications require more complex images, the image interface speed will need to be increased accordingly.

Image Deconstruction Process. The image deconstruction process (IDP) buffers the volume images received from the VMP, slices the volume images into 2D frames, and provides these frames to the illumination process via the frame interface. This process would likely be performed by a processor resident in the DVDD itself. The exact form of the data of the 2D frame depends on the nature of the illumination mechanism in the illumination process and on the geometry of the display surface (for swept volume DVDDs) used in the image reconstruction process. For purposes of this discussion, it can be characterized by $\{h, v, c\}$, where h and v are the horizontal and vertical indices for the pixel and c represents pixel color. Memory requirements depend on the number of voxels to be displayed. Double-buffering a full volume image would require only 100 Mbytes of random access memory (RAM).

The algorithm for slicing the 3D volume image into 2D frames depends on details of the image reconstruction process. Assuming a vertical plane display surface, it is anticipated that approximately 800 frames (actually 1600 but they occur in mirrored pairs) would be required to achieve comparable voxel diameters in the three dimensions. Each of these frames would require an addressability of 512×256. With 6-bit color an additional 160 Mbytes of RAM would be required to store the frames (double-buffered). If the display surface is a helix, then the 2D frames are not strictly 2D in nature; however, each "2D frame" represents all the pixels that need to be illuminated during that short interval of time.

Again taking advantage of the sparse nature of current DVDD images, we can reduce the memory requirements to more manageable levels. The slicing algorithm could simply store *in the volume image* the h index (since $v = z$ for a vertical screen) plus a frame number f for each voxel, thus eliminating the need to again store the color. The entire voxel description is $\{x, y, z, c, h, f\}$ and fits in 7 bytes, thus an image of 100,000 voxels could be stored in 1.4 Mbytes of RAM (double-buffered). This amount of memory is easily provided; however, again as applications require more complex images, the amount of required memory will increase accordingly.

The frame interface must transfer 2D frames to the illumination process at a rate such that each frame is redisplayed at about 40 Hz to eliminate flicker. While current prototype DVDDs operate at about half this rate (see Table 3), and these rates seem adequate during short demonstrations, the resulting flicker becomes very annoying with prolonged viewing. This means that frames must be transferred from the IDP to the illumination process at a very high rate: 32 kHz (that's 32,000 Hz!) if the display screen is passive, 64 kHz if it is active.

If a full volume image is to be sent from the IDP across the frame interface to the illumination process, then the required data rate is 4.2 Gbytes/s. This is 20 times the throughput rate of the high-performance parallel interface (rated at approximately 0.2 Gbytes/s). The solution is again to transfer data across the interface for only the pixels that are to be illuminated. Doing this in the form $\{h, z, c\}$ for a volume image comprised of 100,000 voxels requires an average data rate of only 3 Mbytes/s, which is in the range of small computer standard interface bus speeds.

Illumination Process. The illumination process (ILP) accepts image data (frames) from the frame interface and translates them into the corresponding 2D images, sending them to the image reconstruction process via the display interface. The display interface specifies the kind of optical signal expected by the image reconstruction process and any status data returned. The pixel addressability in the 2D frame needs to be at least 512×256 to be compatible with the voxel addressability of the volume image. The 40 Hz refresh rate required to eliminate flicker and the 800 frames per rotation required to achieve comparable voxel diameters in the three dimensions results in very high frame rates of 32 kHz or 64 kHz.

A frame rate of 32 kHz means that the entire screen must be drawn in only about 30 μs, a speed near the limit of the capabilities of CRTs. Voxel illumination may be achieved passively (light is reflected off of a moving translucent display screen) or achieved actively (light is emitted directly from a moving opaque display screen).

For a passive screen, several techniques have been considered, ranging from scanning an individual light beam, through various levels of parallel beam scanning, to the ultimate in parallelism, a frame addressed system (19). Beam scanning currently depends on acousto-optic scanners, rotating polygonal mirrors, or galvanometers, none of which is fast enough. Color mixing has been difficult to achieve because of the difficulty of accurately registering two or more beams.

Even if the scanning rates could be achieved and the registration problems solved, the beam intensity modulation rates required are very high. For example, a single-beam raster scan system requires a beam intensity modulation rate of 4.2 GHz. Taking advantage of the sparse nature of current DVDD images does not solve this problem. Increases in acousto-optic bandwidth by orders of magnitude are required for a calligraphic system to be able to display an image containing only 100,000 voxels with a single laser. Luckily, frame addressed systems with performance parameters approaching DVDD requirements are under development.

For an active screen the display surface is illuminated on one side only; hence any individual observer sees each voxel illuminated only once per rotation. The surface must be rotated twice as fast to achieve flicker-free images, and this necessitates a higher frame rate of about 64 kHz to achieve an image refresh rate of 40 Hz. Light-emitting diodes attached in a matrix on the display surface easily provide the required intensity modulation rates; however, LEDs currently suffer from at least two problems: low resolution and addressability (due mainly to low-density packaging) and low light output in general. Luckily these are problems being addressed by ongoing research and development by LED manufacturers.

Image Reconstruction Process. The image reconstruction process (IRP) builds the 2D images it receives from the ILP via the display interface back into a 3D volume image. Current DVDDs accomplish this by rapidly rotating a display surface within a cylindrical display volume. The surface sweeps through all possible Cartesian coordinates at least once per rotation, providing the medium on which graphics are drawn by the ILP. In order to eliminate flicker in the volume image, the surface must rotate at about 20 rotations/s (1200 rotations/min) if the display screen is passive and 40 rotations/s (2400 rotations/min) if the display screen is active. The rotation rate of the surface determines many of the data rates discussed above. For a DVDD with a vertical plane surface the slicing algorithm, in its simplest form, involves an inverse transcendental function and a square root (19). It is, however, susceptible to table lookup and parallel processing, since at the IDP level each voxel can be processed completely independently.

The IRP must rotate the display surface (a significant mass) at a rate of 20 rotations/s in a manner that is safe, is vibration-free, does not distort the surface, does not obstruct the observer's view of the display volume, and generates little acoustic noise. For a passive screen the display surface must provide a reflection of the incident image that is as close to isotropic as possible, to avoid observer-centric effects. For an active screen the display surface must be rotated even faster, and the pixels must provide an emission that is as close to isotropic as possible, again to avoid observer-centric effects. These potential problems seem to be solvable by careful engineering practices.

Discussion

The division of image display into the four processes presented here is somewhat arbitrary. For example, three subprocesses comprise the voxelization and mapping process, one providing an API for B-rep data, one providing an API for V-rep data, and a subprocess to map V-rep data to a common addressability. Alternatively, these could be considered three separate processes. Nevertheless, the division into processes and the process flow presented here has proven useful in structuring design considerations. Wefer (19) provides additional detail by beginning at the bottom of the process flow (i.e., at the image reconstruction process) and working up the graphics pipeline to the application program interface, looking at more details along the way.

From the investigation of process flow and interfaces, we can conclude that a DVDD with the capability of displaying a full volume image with a resolution/addressability of 512 $\times$

512 × 256, with a voxel refresh rate of 40 Hz, with an image update rate of 10 Hz, and with a voxel size of about 1 mm is beyond the capabilities of current technology.

Limiting the 3D images to 100,000 voxels (versus the 67 million voxels of a full volume image), however, brings the problems within the range of current technologies. DVDDs with the ability to display an order of magnitude fewer voxels have been used in real applications [see, for example, Hobbs (35)]; hence a DVDD with the capability to display 100,000 voxels should be viewed as a vast improvement, rather than as a serious limitation.

CONCLUSION

While development efforts have been in progress for many years, considerable further effort appears to be required before a commercially viable DVDD becomes available. Some of the approaches discussed in this presentation will lead eventually to this goal, if some other technology does not displace the whole concept. But at some point the attitude of the researchers in this field needs to migrate from one that might be characterized by "I will overcome the technical problems associated with my chosen approach, even if it kills me" to one of "I will implement the approach which has the most advantages and the fewest disadvantages and which will produce a commercially viable DVDD for which users can develop interesting applications."

Based on the above discussion, the DVDD approach with the most advantages and the fewest disadvantages appears to be that of the active screen swept volume device with a vertical plane surface illuminated by LEDs. Technical problems seem to be both nonfundamental and manageable. There are, of course other issues in addition to technical problems, one of these being related to patents. The situation is reminiscent of that surrounding the development of the sewing machine in the nineteenth century (36), a situation that resulted in the postponement of a commercially viable sewing machine for years by wrangling over patent claims and counterclaims.

Luckily in the case of DVDDs, the development activities and the wrangling have now taken so long that the relevant patents are expiring, opening the way for an organization with the technical know-how, the financial backing, and the marketing and sales abilities to produce a commercially viable device. We can look forward to a time when applications like air traffic control, submarine situation display, molecular modeling, medical imaging, computer-aided design, and a host of others will benefit from the group display of 3D information in true 3D.

BIBLIOGRAPHY

1. L. D. Sher, The oscillating-mirror technique for realizing true 3D, in D. McAllister (ed.), *Stereo Computer Graphics and Other True 3D Technologies,* Princeton, NJ: Princeton Univ. Press, 1993, Chap. 11, pp. 339–364. (See also Bolt Beranek and Newman Systems and Technologies Corporation, SpaceGraph *Display System,* marketing brochure, 1991, 8 pages.)
2. S. A. Benton, Experiments in holographic video imaging, *SPIE Inst. Ser.,* **IS 8**: 247–267, 1990.
3. S. A. Benton and M. Lucente, Interactive computation of display holograms, *Proc. Comput. Graph. Int. '92,* Tokyo, Japan, 1992.
4. W. Krueger, Volume rendering and data feature enhancement, *Computer Graphics,* **24** (5): 21–26, 1990.
5. A. Kaufman, Introduction to volume visualization, in A. Kaufman (ed.), *Volume Visualization,* Los Alamitos, CA: IEEE Computer Society Press, 1991, Chap. 1, pp. 1–18.
6. W. E. Lorensen and H. E. Cline, Marching cubes: A high resolution 3D surface construction algorithm, *Comput. Graph.,* **21** (4): 163–169, 1987.
7. H. E. Cline et al., Two algorithms for the three-dimensional reconstruction of tomograms, *Med. Phys.,* **15** (3): 320–327, 1988.
8. R. A. Drebin, L. Carpenter, and P. Hanrahan, Volume rendering, *Comput. Graph.,* **22** (4): 51–58, 1988.
9. U. Brinkman, A laser based three dimensional display, *Lasers and Appl.,* **March**: 55–56, 1983.
10. P. Soltan et al., Laser projected 3D volumetric displays, *Naval Command, Control and Ocean Surveillance Center Report,* NRaD, San Diego, CA, 1997.
11. P. Soltan et al., Laser based 3D volumetric display system, *SPIE / IS&T Symposium on Electronic Imaging Science and Technology, High-Resolution Displays and Projection Systems,* San Jose, CA, 1992, pp. 177–192.
12. P. Soltan, W. Robinson, and W. Dahlke, Laser based 3D volumetric display system (2nd generation), *SID 1994 Dig.,* 1–4, 1994, pp. 191–194.
13. P. Soltan et al., Laser-based 3D volumetric display system, *Naval Command, Control and Ocean Surveillance Center Report,* NRaD, San Diego, CA, 1996.
14. A. J. Schwarz and B. G. Blundell, Regions of extreme image distortion in rotating-screen volumetric display systems, *Comput. Graph.,* **18** (5): 643–652, 1994.
15. R. D. Williams and F. Garcia, Jr., A real-time autostereoscopic multiplanar 3D display system, *SID '88 Dig. Tech. Papers,* **19**: 91–94, 1988.
16. R. D. Williams and F. Garcia, Jr., Volume visualization displays, *Inf. Display,* **5** (4): 8–10, 1989.
17. R. D. Williams, F. L. Wefer, and T. E. Clifton, Direct volumetric visualization, *Proc. Visualization '92,* Boston, MA, 1992, pp. 99–106.
18. T. E. Clifton and F. L. Wefer, Direct volume display devices, *Comput. Graph. Appl.,* **13** (4): 57–65, 1993.
19. F. L. Wefer, Computer graphics direct volume display devices, MITRE Corp., McLean, VA, Technical Report 92W0000141, 1994.
20. S. Shimada, A new approach to the real-time 3D globe display, *SID '93 Dig. Tech. Papers,* **24**: 1001–1004, 1993.
21. R. G. Batchko, Multiplanar displays, *Electron. Imaging,* **August**: 1992; Volumetric displays, *Inf. Display,* **8** (9): 8, 1992; 2-D Techniques for 3-D images, *Lasers Opt.,* **September**: 8, 1992; and Volumetric imaging meets information display needs, *Photonics Spectra,* **August**: 118, 1992.
22. D. J. Solomon, Volumetric Imaging, Inc., *Three-dimensional volumetric display system,* US Patent 4,983,031, 1991.
23. D. J. Solomon, Volumetric imaging launches graphics into a 3-D world, *Photonics Spectra,* **July**: 129–135, 1993.
24. B. G. Blundell, A. J. Schwarz, and D. K. Horrell, The cathode ray sphere: A prototype volumetric display system, *Proc. Eurodisplay 1993 (Late News Papers),* 1993, pp. 593–596.
25. B. G. Blundell, A. J. Schwarz, and D. K. Horrell, The cathode ray sphere: A prototype system to display volumetric three-dimensional images, *Opt. Eng.,* **33** (1): 180–186, 1994.
26. B. G. Blundell and A. J. Schwarz, A graphics hierarchy for the visualization of 3D images by means of a volumetric display system, *Proc. IEEE Tencon 1994,* 1994, pp. 1–5.
27. A. J. Schwarz and B. G. Blundell, Considerations for accurate voxel positioning on a rotating-screen volumetric display system, *IEE Proc. J: Optoelectronics,* **141** (5): 336–344, 1994.

28. J. Alburger, Crystal ball plots 3D curves in color, *Electron. Ind. Tele-Tech,* **February**: 50–53, 1957.
29. H. S. Cole, D. W. Skelly, and C. R. Stein, A three dimensional display for radar returns, *SID Symp.,* **May**: 52–53, 1974.
30. D. L. MacFarlane, Volumetric three dimensional display, *Appl. Opt.,* **33** (31): 7453–7457, 1994.
31. C. M. Verber, Present and potential capabilities of three-dimensional displays using sequential excitation of fluorescence, *Three-Dimensional Imaging, SPIE Proc.,* **120**: 62–67, 1977.
32. P. Soltan and M. Lasher, Non-moving 3D volumetric displays using upconversion materials, *Naval Command, Control and Ocean Surveillance Center Report,* NRaD, San Diego, 1997.
33. D. P. Mahoney, Full-color display in a cube, *Comput. Graph. World,* **March 20** (3): 19, 1997.
34. T. E. Clifton and F. L. Wefer, Functional architecture and component technologies for direct volume display devices, *SID '93 Dig. Tech. Papers,* **24**: 1031–1034, 1993.
35. B. A. Hobbs, *A user interface to a true 3-D display device,* Master's thesis in Computer Engineering, Air Force Institute of Technology, AFIT/GCE/ENG/92D-06, 1992.
36. D. J. Boorstin, *The Americans: The Democratic Experience,* New York: Vintage Books, 1973, pp. 92–97.

FRED L. WEFER
The MITRE Corporation