



The ability to look into an electronic display from any direction, walk around it, and see real 3D information being displayed is a dream-come-true technology for those who work with volumetric data sets.

REALIZING 3D VISUALIZATION USING CROSSED-BEAM VOLUMETRIC DISPLAYS

Since we live in a three-dimensional world, we continuously interact with 3D objects both near and far away. The human body can extract enormous amounts of information about its environment from a single sensory input: sight. Because we obtain over 70% percent of our sensory input visually, we judge shape, size, distance, relative position, movement, speed, and a host of other physical attributes using the binocular vision provided by our two eyes. In addition, the human visual system works in real time, enabling us to interpret a 3D world that is changing rapidly. Our incredible ability to perceive depth through stereopsis, motion parallax, focus, and eye convergence is

more reliable and useful than the visual cues present in current 2D displays using relative size, superposition, and lighting.

Electronic displays such as the cathode-ray tube (CRT) and the liquid crystal display (LCD) present visual information on flat surfaces. Despite powerful rendering techniques, such as perspective, shading, shadowing, and texturing that have been employed to increase the realism of these images,

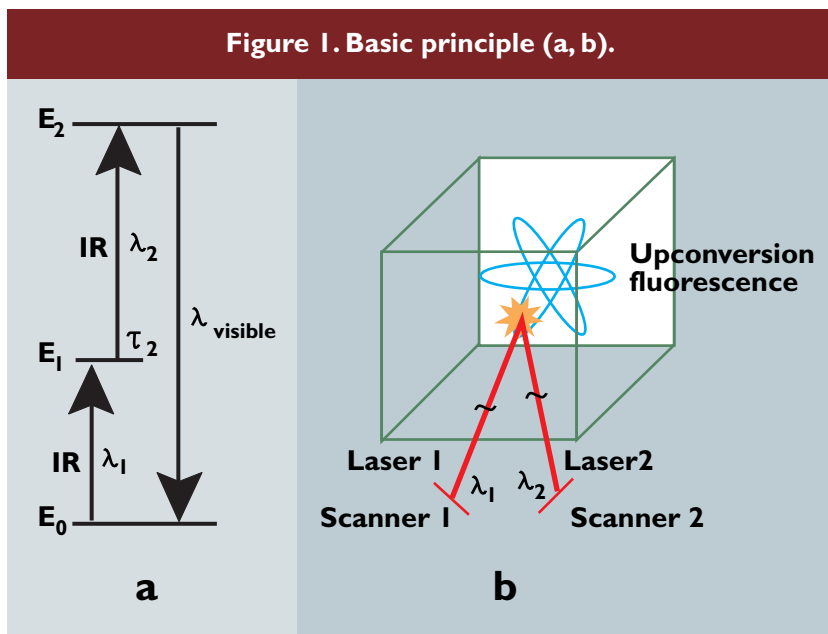
they are still flat. In addition, the data is presented from a single-user vantage point that prevents viewers from using motion parallax to gain slightly different views in order to extract depth information. Even stereo display techniques that provide separate left and right eye images do not provide users with sufficient depth cues, resulting in a conflict between focus and convergence that often causes unwanted physiological side effects.

Many graphic constructs have been developed to make flat display screens seem three-dimensional, but they cannot generate true volumetric images because the systems lack basic volume-addressing capabilities. This article describes a fundamentally different display hardware, called a crossed-beam display (CBD), which enables data actually to be addressed as a volume. These displays can be viewed from almost any direction by multiple viewers simultaneously, without the need for glasses or headgear, thereby providing real depth perception to users. The emerging electronic display modality of CBDs provides not only the promise of enhanced visualization, but challenges and opportunities for graphics software developers who wish to display new and existing data sets using a fundamentally different architecture.

CBDs open the door for doctors to eventually diagnose medical conditions and plan treatment therapies using MRI, CT scan, and ultrasound data that can be seen and interacted with in real 3D. Engineers using CAD systems to design complex parts will be able to do so using real 3D workstations. And mathematicians performing finite element analysis will be able to display the results of their multidimensional computations, on real-time, color, volumetric grids. Even entertainment will feel the impact when interactive 3D video games can be played from opposite sides of the table. The CBD, an elegant yet fundamentally simple technology that has finally emerged on the display scene, has the potential to enhance visualization in a wide range of diverse applications including medicine, science, CAD/CAM, remote exploration, vehicle guidance and control, microscopy, flow cytometry, education, entertainment, and more.

Crossed-Beam 3D Displays

The physical mechanism upon which CBDs are based is known as gated, two-frequency, upconversion fluorescence [4, 9]. This phenomenon occurs when an active ion that has been dispersed in an appropriate host material is optically excited to emit visible light by the sequential absorption of two low-energy infrared photons. In the more familiar form of fluorescence, called Stokes fluorescence, a material absorbs high-energy photons of one wavelength, then emits visible light at a slightly lower energy.



This is the underlying phenomenon behind the CRTs in television sets and computer monitors. In these devices, a high-energy electron beam is magnetically scanned across a screen of discretely placed phosphor dots called pixels, causing them to emit visible photons with brightness proportional to the intensity of the electron beam. In contrast, the excitation energy in a CBD is absorbed from two infrared lasers with different wavelengths. The sums of these wavelengths are different, but their sum, in energy, roughly corresponds to the visible emission wavelength. The first infrared laser excites the electrons from the ground state to an intermediate metastable level (Figure 1a), and the second (different) infrared wavelength excites the electrons to an even higher energy level. Though the excitation is sequential, involving a real intermediate level, visible light is emitted by the ions because spin coupling causes the excited electrons to fall directly down to the ground state, not back down the energy-level ladder. The ability to use two intersecting invisible laser beams to create a single point of visible light (Figure 1b) enables “voxels” (3D pixels) to be addressed anywhere inside a volume of active material simply by controlling when and where the lasers intersect. Rapidly scanning the point of intersection in a predetermined manner enables real 3D images to be drawn. The size of the voxel is defined by the diameter of the lasers at the point of intersection, and can be specified with focusing optics to provide resolution on the order of 50 μ m to 500 μ m. The emission wavelength, or color, is controlled through appropriate selection of both the active ion and the excitation wavelengths.

Visualization parameters. The value of 3D displays is vastly improved depth perception of electronic images and data sets. Greater storage densities and processing power, combined with sophisticated data acquisition techniques, have provided a means to extract and create 3D data in fields as diverse as weather mapping, CAD/CAM, scientific visualization, and medicine. Even video games contain 3D information.

Visual depth perception is a complex cognitive process involving the eyes, brain, and various muscles of the body, all working together as a visual system. Among the cues that the human visual system uses to perceive depth are *accommodation* and *convergence*, which are used together to give us binocular or stereoscopic vision [6]. Accommodation is the ability to focus the eye's lens on objects in the frontal field of view using the ciliary muscles behind the iris. Convergence is the ability of the eyes to pivot inward toward the object in focus, using the medial rectii muscles on the nasal side of the eyeball. Binocular vision is the ability for the visual cortex of the brain to process two disparate 2D images on the retinas to extract spatial information in conjunction with the signals from the focusing and converging muscles.

Because the human visual system uses these depth cues naturally, visual stimuli, which do not contain them or combine them in unusual ways cause errors in the interpretation process. These errors can sometimes be interesting, as in the case of simple optical illusions, or they can be problematic, as in the case of stereo and virtual reality displays that present two different flat images—one for each eye. Stereoscopic displays introduce a fundamental conflict between accommodation and convergence by providing two images that satisfy convergence, but only a single image plane for accommodation. This conflict is well known for its tendency to induce dizziness and nausea in stereo and VR display users, particularly after extended periods of use. Accurate interpretation of complex volumetric data clearly requires the presence of more depth cues than can possibly be provided by conventional two-dimensional, monoscopic and stereoscopic displays [3]. Depth cues such as motion parallax, which humans employ by moving their heads to acquire different views of real scenes, can not be supported using flat screen displays. Head-tracking, which has been the solution to providing motion parallax, is not only slow, because of the need to recompute the image and motion lags in the head-tracking system, but it constrains the display to a single user, thereby limiting collaborative interaction. The advantage of CBD technology is that it allows the viewer to use natural physiological and psychological depth cues such as accommodation, convergence, and

stereo disparity to determine the size, shape, and relative position of objects.

A variety of techniques such as perspective, shading, and shadowing have provided significant improvements for flat screens, but only from a single vantage point. In order to present volumetric data in a real 3D context a display in which data can be both addressed and viewed using three spatial dimensions is needed. While numerous clever approaches to providing volume-addressability and volume-viewability have been developed over the past few decades, the implemented concepts are intrinsically flawed. These techniques include images projected onto moving mirrors and images that are scanned onto rapidly rotating surfaces using visible lasers. One such device, the varifocal mirror display [12], uses an oscillating mirror to project a 2D display image at sequential depth planes within a 3D volume. Others include a multitude of spinning element displays [1, 8]. The drawbacks of these systems range from the need to rotate large surfaces at high speeds to the need to reflect visible laser light into the eye of the viewer. Because CBDs use infrared excitation to address visible voxels, which in turn emit harmless, incoherent light, they pose no “eye-fry” hazards to users. This is in sharp contrast to displays that use lasers, as even small amounts of scattered coherent radiation have been shown to cause retinal damage.

Crossed-beam displays provide the ability to selectively address voxels that emit light isotropically in a 3D volume of material. This enables visual information to be viewed from any direction by multiple users simultaneously, in much the same manner as a real 3D scene would be viewed. To gain a different perspective of a data set, a user need only move his head, thus enabling motion parallax to be employed as well. The image that is presented is independent of viewing perspective, free from conflicting visual depth cues, and does not require recomputation. In addition, the nonimmersive nature of CBDs provides collaborative and interactive volumetric viewing, a capability that is not possible with head-mounted VR systems.

Implementation. Unlike CRTs, which converged on a standard architecture decades ago, the fundamental system design of CBDs is still emerging. Notable differences between these devices are the type of excitation beam and the mechanism used for deflection. CRTs excite pixels by using electron beams that are rapidly scanned in both the horizontal and vertical directions via oscillating electromagnetic fields. CBDs rely on laser beams, which cannot be scanned using either electric or magnetic fields but must be reflected using mirrors, or diffracted using optical elements.

Figure 2. Photographs of images in CBD display chamber.




Figure 2 shows photographs of an existing prototype CBD engine. The image chamber is fabricated from glass, and the lasers and scanners are housed below the image chamber in the optical platform. The current size of the image chamber is 4in. x 4in. x 4in. (64in.³), with the amount of addressable data scaling linearly with pump laser power. Drive electronics for the lasers and scanners are packaged separately, and the data input to the device is currently provided via a PCI interface to a computer. The display is not intended to replace a CRT for displaying text and normal data processing operations, but rather to provide the added degree of visualization capability essential for many volumetric data sets.

In addition to improved visualization, CBDs feature a number of attributes that make them attractive from a manufacturing standpoint. Unlike two-dimensional LCD and CRT displays, the image chamber is non-pixelated, providing manufacturing ease beyond that of even the lowest resolution flat panel displays. The entire display volume of a CBD can be made from a solid, homogeneously doped material that requires no pixelization. Using remotely positioned lasers, voxels are optically (not electrically)

addressed, eliminating the need to embed wires, transistors, fibers, or transparent electrodes in the image chamber. These features are significant as it is the low yield of pixelated LCD displays that is still the major cost driver, and complex volume-manufacturing techniques of a similar nature would be cost-prohibitive. The image chambers of CBDs are solid-state, with no rotating mechanical parts to fail or pose safety risks. Resolution is a function of the focusing properties of the excitation lasers, and is governed by diffraction. The use of infrared excitation sources enables unused pump energy to be easily and discretely filtered from the displayed images with infrared filters. The visible light emitted from the voxels in CBDs is incoherent, like a CRT, eliminating any potential risk of retinal damage: at no time do viewers look at direct, scattered, or reflected coherent radiation. Though many issues must still be addressed in order for CBDs to solve a wide range of visualization problems, the intrinsic fundamental attributes of CBDs position them to be a powerful new addition to the display infrastructure of the electronic and information age.

Computer Science Challenges and Opportunities

In addition to display quality, ease of use is important to the CBD's success. Requiring CBD application developers explicitly to control scanner and laser hardware is unreasonable. Instead, what is desirable is a support architecture (software and hardware) that provides a simple application programming interface (API) that accepts compact representations of 3D objects and converts these representations into the appropriate hardware control signals for the CBD. This structure correlates roughly to the combination of graphical APIs (for example, OpenGL) and graphics hardware accelerators currently available for 2D displays.

The proposed support architecture contains three levels. The API layer receives 3D object and environment specifications from an application and converts them into an intermediate representation. The translation layer converts the intermediate representation into hardware commands. The hardware layer translates the hardware commands into analog signals and sends the signals to the CBD hardware. Finally, the CBD hardware executes the signals to render representations of the original 3D objects.

API layer. The API layer needs to be a general, extensible graphics specification interface that supports specification of both surface and volumetric data. A volumetrically extended version of OpenGL will be used initially, and Java3D and Fahrenheit in the future, to provide a standardized interface to application programs. The API extension will allow specification of volumetric data, (voxel-based, CSG), volumetric display specific information (display projection), and environment information (light location). While traditional perspective projection is not needed with a volumetric display, other viewing projection transformations may be useful to accommodate for non-Cartesian object data and possibly to accentuate depth cues using linear perspective [1]. Specifying the location of a head-tracked user may be necessary for developing techniques for simulating view-dependent perspective and illumination effects, such as specular illumination and refraction.

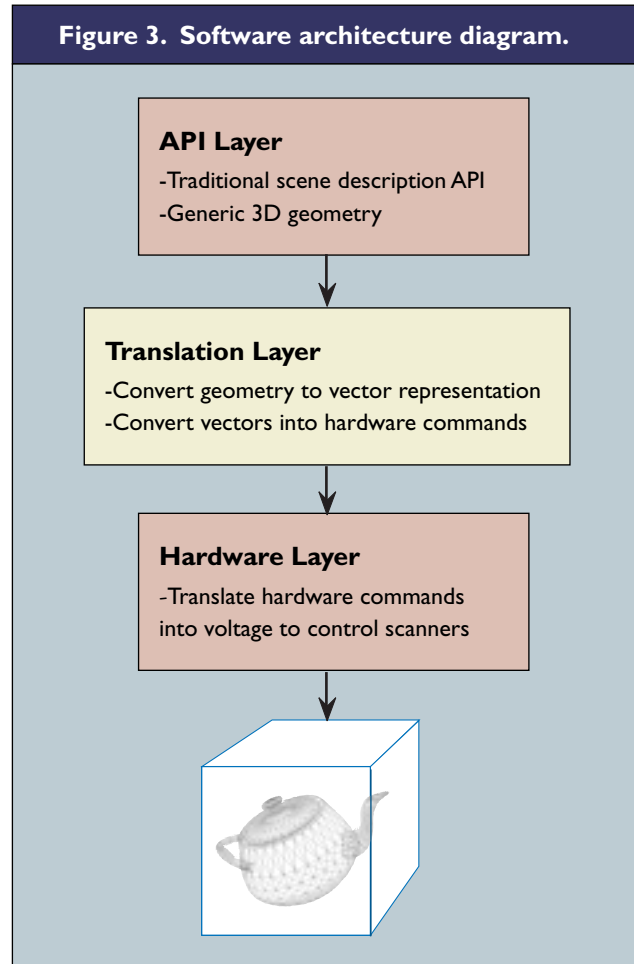
These input parameters to the API layer will be “rendered” to an intermediate representation through a combination of software and hardware-accelerated graphics techniques. The use of existing hardware graphics engines to process the input data into a suitable format for the CBD display is being explored. Rendering for this display technology raises many interesting questions, such as:

- How do you simulate viewer-dependent shading characteristics when observers may be simultaneously viewing the scene from different locations?
- Since the display medium emits light but is transparent, how can you simulate semiopaque objects?
- What are the most appropriate techniques for anti-aliasing given the characteristics of the display technology?

The answers to these questions will involve quality/performance trade-offs that are currently being explored. The goal is to provide a flexible API layer that contains the information needed for making hardware-specific decisions. However, the API layer interface will be device-independent to isolate changes in hardware technology from application software. The output of the API layer is dependent on whether the volumetric display hardware is vector-based or raster-based.

Vector display. In general, a vector-based system takes 3D object specifications (for example, “sphere of radius 1 meter”) as input and calculates the resulting image using only lines and points as drawing primitives. The display hardware then steps sequentially through the list of lines and points and draws them.

Figure 3. Software architecture diagram.



For the vector-based CBD, lines are additionally decomposed into points. This is done for two reasons. The first is to ensure straight lines are drawn by circumventing inaccuracies and nonlinear responses in the scanners when they are tasked to traverse relatively large distances. The second reason is to control render characteristics (intensity, opacity, color) of all voxels along a line by changing laser and scanner characteristics (for example, laser dwell time, intensity, beam focus, and beam frequency).

Raster display. Raster-based systems divide the display surface into regularly spaced pixels or voxels. Typically, an array of memory, called a frame buffer, is used to store information such as color and opacity for each pixel or voxel. The display system steps through each pixel or voxel in the display and each corresponding element in the frame buffer, and illuminates the pixel or voxel in accordance to the frame buffer information. To fill the frame buffer with information, the display system takes 3D object specifications as input (for example, “blue square of size 0.5 inches”) and calculates the resulting image using only points (typically at the resolution of the frame buffer), and stores the point information in the frame buffer. Note that since the

display system looks at all frame buffer entries, locations in the frame buffer that correspond to having nothing drawn need to be explicitly set to a value that means “draw nothing”.

The software architecture for the raster-based CBD display (see Figure 3) utilizes a 3D frame buffer. The API layer produces a 3D array of points representing the entire display volume. The points are stored in the frame buffer as $\{r, g, b, a\}$, where $\{r, g, b\}$ is the color and intensity of the point, and $\{a\}$ is the alpha value, or opacity of the point. A $\{r, g, b\}$ value of $\{0, 0, 0\}$ indicates nothing should be drawn.

Translation layer. The translation layer takes the output of the API Layer and converts it into hardware commands. Since this layer is closely associated with the hardware, it is customized for the various CBD configurations.

For vector CBDs, the translation layer is given a list of seven-tuples from the API layer, which represents the points to be drawn (x, y, z, r, g, b, a) . The translation layer may elect to, depending on performance trade-offs, sort the list for various optimizations depending on characteristics of the lasers, scanners, and display volume. These include minimizing scanner movements [10], minimizing color or intensity changes, and perhaps maximizing opacity quality by drawing from the center of the display volume outwards or vice-versa. Then the translation layer traverses the sorted list of seven-tuples and computes the angles for each scanner from the position $\{x,y,z\}$ and the color, dwell time, and laser intensity from the alpha value $\{a\}$ and color elements $\{r,g,b\}$. This information is translated into hardware commands and sent to the hardware layer.

For raster CBDs, the translation layer is given a 3D frame buffer from the API layer, which represents all the drawable points in the display. The translation layer traverses the regions and creates hardware commands for the laser intensity, color, and dwell specifications. Display locations are implicitly known, given the location in the frame buffer, and are used to calculate the angle for each scanner. The information for each region is translated into hardware commands and sent to the hardware layer.

The raster architecture allows for some interesting variations, such as a second frame buffer. With a second frame buffer the translation layer can read from one frame buffer while the API layer writes to the other. The translation layer and API layer then switch buffers and repeat the process. This is called double buffering and is a prevalent method of eliminating display flicker created when a single frame buffer is written to at the same time it is being read by the rendering system. This architecture also allows optimiza-

tion using a dual-access single frame buffer.

A second variation is to interlace the display. Ideally, the entire contents of the frame buffer can be repeatedly displayed fast enough to avoid flicker to viewers (approximately 60Hz) [11]. This will not be possible when the number of voxels overwhelm the capabilities of the support architecture and hardware. To help this situation, the translation layer can interlace the frame buffer by dividing the frame buffer into halves and drawing the halves sequentially. The first half consists of every other voxel, while the second half consists of the remaining voxels. This method fills the entire display space with half the points, but twice as quickly, thus eliminating the flicker. Interlacing is an effective technique that is used in broadcast television, among other places.

Hardware layer. The hardware layer takes the hardware control commands and converts them into electrical laser and scanner control signals. The implementation of this layer is tightly coupled to the laser and scanner hardware.

Data bottlenecks. The data transfer needed for interactive update of volumetric displays can quickly become a bottleneck. Different implementations of the software architecture are being explored to address this problem, including methods for real-time volume encoding to decrease bandwidth requirements and the use of special purpose hardware (digital signal processing chips, programmable gate arrays) for implementation of portions of the API and translation layers.

User interaction. Since the CBD will provide true 3D display of volumetric data, 3D interaction techniques can be incorporated to complete the 3D experience and provide a natural interaction metaphor to this new display technology. As mentioned previously, head-tracking of users may be useful for view-dependent shading simulations. Six-degrees-of-freedom input devices may also be useful for interacting with the volumetric data.

Future Efforts

In addition to the new software architecture that is being developed at present, there remain a number of hardware challenges that must be overcome to make CBDs an affordable mainstream technology. The current size (a four-inch cube) is smaller than many applications require, and the upconversion fluorescence efficiencies of the image chamber materials are low. These issues are being addressed with higher efficiency optical materials that can be manufactured into larger image chamber volumes, combined with advanced system architectures capable of displaying higher density data sets.

Crossed-beam displays are based on a fundamental concept that has been envisioned by creative display and software developers for many years. The technology is now viable because of the commercial availability of components like laser diodes and high-speed processors, and because of the recent identification of materials in which the concept could be demonstrated to work. CBDs currently stand in much the same position as the CRT did over 70 years ago [5]. CBDs present a feasible option for viewing and visualizing 3D data for a variety of applications. The ability to allow several users to simultaneously view 3D data sets provides a collaborative research environment. The three-dimensionality of the display also can provide important spatial information for many applications and can decrease the time necessary to understand 3D spatial relationships among objects within the data set. Medical visualization, surgical training, scientific visualization, command and control applications (air traffic control, for example), and many other applications will benefit from the improved 3D information display of CBDs.

Though CBD technology is quite new, early adopters at the NASA Goddard Space Flight Center are exploring the use of the technology as an alternative to traditional stereoscopic displays. The researchers there have already begun to assimilate and embrace CBDs, and are leading the application development effort.

NASA has numerous applications for a volumetric display in the areas of modeling, satellite measurements, spacecraft design, and spacecraft operation. Modeling refers to the use of a mathematical model to simulate natural phenomena. NASA is interested in numerous phenomena, including Earth's atmosphere, Earth's magnetosphere, solar wind, planetary, stellar and galactic formation and evolution, aircraft and spacecraft body and engine performance, and so forth. The output of these mathematical models is typically large and three-dimensional.

In addition to modeling various behaviors of nature, NASA also directly measures them—most often with an instrument on a satellite. Traditionally, satellites have generated 2D measurements (surface temperature, for instance, since the temperature is recorded for each location in longitude-latitude pairs). However, recent satellites can generate 3D measurements. This is the case with the Tropical Rainfall Measuring Mission Satellite, which measures, among other things, water moisture content within a large, 3D region of the Earth's atmosphere.

Developing complex spacecraft such as the International Space Station requires detailed analysis of large quantities of 3D data. Making these spacecraft

simple to operate and maintain requires advanced monitoring and control mechanisms, such as a status display that incorporates a 3D representation of the spacecraft.

All of these activities could benefit from a volumetric display. NASA plans to use early implementations of the CBD for vector representations of atmospheric wind models and particle animations of modeled flows around wings. As more capabilities are incorporated into the display, NASA will use the display for more complex visualization tasks.

As the technology continues to develop, the use and potential of these volumetric displays will be further explored for new application areas. Surgical simulation, display-guided surgical operations, medical diagnosis and treatment planning, command and control applications, and even home-entertainment systems have the potential to exploit this new technology. **□**

REFERENCES

1. Blundell, B., Schwarz, A., and Horrell, D. Volumetric three-dimensional display systems—Their past, present and future. *IEEE Science and Engineering Education* 2, 5 (1993).
2. Blundell, B.G., and Schwartz, A.J. A graphics hierarchy for the visualization of 3D images by means of a volumetric display system. In *Proceedings of IEEE Tencon '94*.
3. Dodsworth, C., Ed. *Digital Illusion: Entertaining the Future with High Technology*. ACM Press, NY 1998.
4. Downing, E.A., et al. A three-color, solid-state, three-dimensional display. *Science* 273, (Aug 30, 1996), 1185–1189.
5. Farnsworth, P.T. Television System. U.S. Patent #1,773,980, Aug. 26, 1930.
6. Gregory, R.L. *Eye and Brain: The Physiology of Seeing*. McGraw-Hill, 1966.
7. Kameyama, K. and Ohtoni, K. A shape modeling system with a volume scanning display and multisensory input device. *Presence* 2, 2 (1993).
8. Lasher, M. et al. Laser-projected 3D volumetric displays. In *Proceedings of Projection Displays II*, 2650, SPIE, 1996.
9. Lewis, J., Verber, C., and McGhee, R. A true three-dimensional display. *IEEE Trans. Elec. Devices*, 18 (1971), 724–729.
10. Schwarz, A. and Blundell, B. Optimizing dot graphics for volumetric displays. *IEEE Computer Graphics and Applications* 17, 3 (May–June 1997).
11. Sekuler, R. and Blake, R. *Perception*. Knopf, NY, 1985.
12. Traub, A.C. Stereoscopic display using varifocal mirror oscillations. *Applied Optics* 6, 6 (June 1967).

DAVID EBERT (ebert@umbc.edu) is an associate professor of computer science and electrical engineering at the University of Maryland, Baltimore County

EDWARD BEDWELL (ted@umiacs.umd.edu) is a software engineer at the University of Maryland Center for Advanced Computer Studies at the University of Maryland, College Park.

STEPHEN MAHER (stephen.maher@gsfc.nasa.gov) is a computer scientist in the Scientific Visualization Studio at NASA's Goddard Space Flight Center.

LAURA SMOLIAR (lsmoliar@3dtl.com) is Vice President for Research and Development at 3D Technology Labs in Sunnyvale, CA.

ELIZABETH DOWNING (elizabeth@3dtl.com) is President and CEO of 3D Technology Labs in Sunnyvale, CA.
