Occlusion-capable multiview volumetric three-dimensional display

Oliver S. Cossairt, Joshua Napoli, Samuel L. Hill, Rick K. Dorval, and Gregg E. Favalora

Volumetric 3D displays are frequently purported to lack the ability to reconstruct scenes with viewer-position-dependent effects such as occlusion. To counter these claims, a swept-screen 198-view horizontal-parallax-only 3D display is reported here that is capable of viewer-position-dependent effects. A digital projector illuminates a rotating vertical diffuser with a series of multiperspective 768 × 768 pixel renderings of a 3D scene. Evidence of near–far object occlusion is reported. The aggregate virtual screen surface for a stationary observer is described, as are guidelines to construct a full-parallax system and the theoretical ability of the present system to project imagery outside of the volume swept by the screen.

OCIS codes: 080.2740, 090.1760, 090.2870, 090.4220, 100.6890, 350.4600.

1. Introduction

Three-dimensional displays are used in fields spanning medical visualization, petroleum exploration and production, and military visualization. Two types of 3D displays pertain to this paper: volumetric and multiview. Although the field's vocabulary is not standardized, volumetric displays generate imagery from light-emitting, light-scattering, or light-relaying regions occupying a volume rather than a surface in space,[1] as averaged over the display's refresh period. For example, multiplanar volumetric 3D displays produce 3D imagery by projecting a series of 2D planar cross sections, or slices, of a 3D scene onto a diffuse surface undergoing periodic motion with a period equal to or less than the eye's integration time.[2,3] One commercially available multiplanar volumetric display is the Perspecta Spatial 3D Display (Actuality Systems, Inc., Bedford, Massachusetts). It reconstructs approximately spherical image volumes with a diameter of 25 cm, where each of 198 radially oriented slices has a resolution of 768 × 768 pixels.

Whereas volumetric displays reconstruct 3D scenes with a set of slices, multiview displays reconstruct scenes with a set of one or more pixelized fields, each transmitting in one or more ray directions, or so-called views. The image surface is usually stationary and planar. Multiview displays take many forms: Spatially multiplexed multiview displays include lenticular[4] and parallax-barrier[5] displays, and angle-multiplexed multiview displays include scanned-illumination systems.[6] Surveys of 3D displays suggest various taxonomies of the field[7,8] while others emphasize volumetric 3D displays.[1,3]

In this paper, we argue against the prevalent assertion that volumetric 3D displays are unable to depict instances of occlusion among scene elements, or more generally that volumetric 3D displays are incapable of any viewer-position-dependent reconstruction effects. To strengthen our argument, we assume that the display system is incapable of sensing the observer position(s) and is suitable for multiple simultaneous observers.

2. Volumetric Displays are Capable of Occlusion

The previous paragraph's alleged drawback appears in several prominent works in the field[8,9] and is generally attributed to the time-sequential architectural aspect of volumetric displays. For example, it is often argued that since many volumetric displays utilize a moving screen, the reconstructed volumetric pixels (or voxels) are somehow necessarily translucent and isotropically emissive. Thus, the argument goes, voxels of nearby objects cannot appear to block the light of occluded voxels of distant objects because, in an

When this research was performed all the authors were with Actuality Systems, Inc., 213 Burlington Road, Suite 105, Bedford Massachusetts 01730, USA. O. S. Cossairt is now at Columbia University, New York, New York, USA. R. K. Dorval is with Insight Technology, Inc., 9 Akira Way, Londonderry, NH 03053, USA. G. Favalora's e-mail address is favalora@actuality-systems.com.

Received 31 July 2006; revised 25 October 2006; accepted 29 October 2006; posted 31 October 2006 (Doc. ID 73641); published 20 February 2007.

0003-6935/07/081244-07$15.00/0

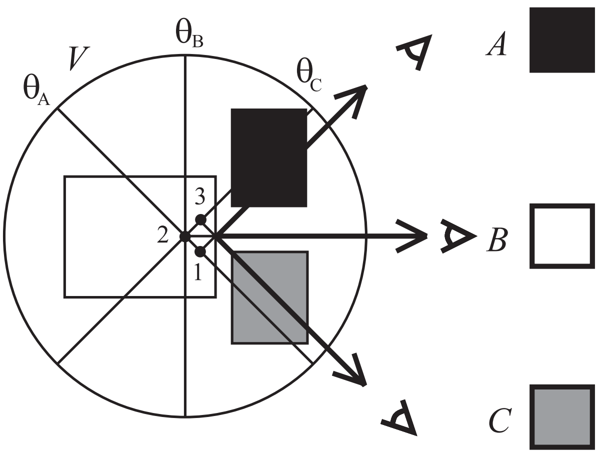

Fig. 1. Scene V composed of three opaque objects. Our multiview volumetric display reconstructs the A ray when the projection screen traverses point 1, the B ray at point 2, and the C ray at point 3 when the screen is at angles θ_A, θ_B, and θ_C, respectively. In a traditional volumetric display, the three rays have contributions from all objects intersecting those rays, usually resulting in incorrect reconstruction. Top view.

approximate sense, their intensities are integrated along the line of sight (see Fig. 1).

While we agree that most volumetric displays built to date are indeed incapable of viewer-position-dependent reconstruction effects, it is a consequence of the properties of the projection surface, not its motion alone. Historically, volumetric displays have employed projection screens with highly diffuse surfaces that act to modulate incident light into a set of point sources. We argue that screens of different design, such as translucent screens with unidirectional diffusion[10] or fields of microlenses,[11] result in volumetric displays whose reconstructed scenes can have angle-dependent per-voxel radiation, as well as the interesting property of projecting scenes that occupy more than the volume swept by the projection surface itself.

Here we present the highest resolution occlusion-capable swept-screen display[12,13] that the authors are aware of, as described in a U.S. patent application that makes it one of the first multiview volumetric displays in general. Earlier work in the field includes a multiview volumetric display that was patented by one of the authors but never constructed.[11] Other work includes a rotating microlouver system,[14] simultaneous 24-view projection onto a rotating controlled-diffusion surface,[15] and a 12-projector system that illuminates a rotating screen composed of a vertically oriented louver.[16] There are a variety of cylindrical-surface multiview displays, such as cylindrical holographic stereograms[17] and the SeeLinder display.[18] However, these are arguably not volumetric displays because the illumination originates from a curved 2D surface, not a volume, even as integrated over the fusion period of human vision.

Notably, several curved static- and dynamic-surface multiview displays were proposed and constructed by Collender[19–21] several decades prior to the recent resurgence of work in this area, exploring a variety of camera-based scene recording methods, a customized microbead-based high-speed illumination assembly, and reconstruction techniques such as viewing imagery through narrow rotating slits or as reflected off a retroreflective cinema screen.

3. Experimental System

In April 2004 we built a 198-view multiview volumetric 3D display by modifying the projection surface and rendering software of the commercially available Perspecta Spatial 3D Display, Version 1.7.[22] We begin by summarizing the operation of an unmodified Perspecta in conjunction with Fig. 2. First, the 3D scene, such as a computer tomography (CT) scan or molecular model, is deconstructed into a series of 768 pixel $\times$ 768 pixel slices using software executed on an x86 computer and an NVIDIA GeForce 6800 Ultra graphics-processing unit.[23] The projection surface is an omnidirectional diffuser with nearly equal transmission and reflection coefficients. Its 25 cm diameter disk-shaped active area is oriented with its normal parallel to the floor and rotates at 900 rpm, centered on the axis of rotation. Since it sweeps two volumes for every 360° rotation, 396 slices are projected onto it in two sets of 198 images that are ideally perceived as superimposed. The slices are projected at approximately 6000 images/s by a group of three Digital Micromirror Devices, microelectromechanical-systems- (MEMS-) based spatial light modulators (Texas Instruments, Inc., Plano, Texas). To ensure proper focus regardless of screen angle, the screen is

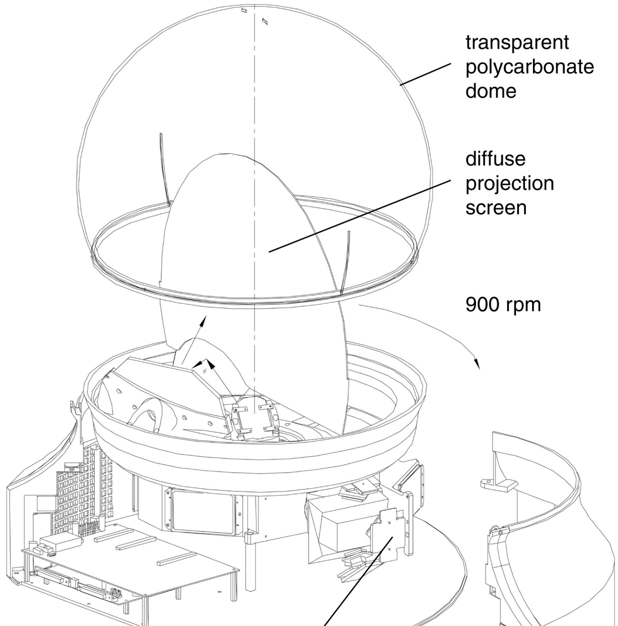

Fig. 2. Perspecta display projects a series of 2D images onto a rotating diffuse screen.

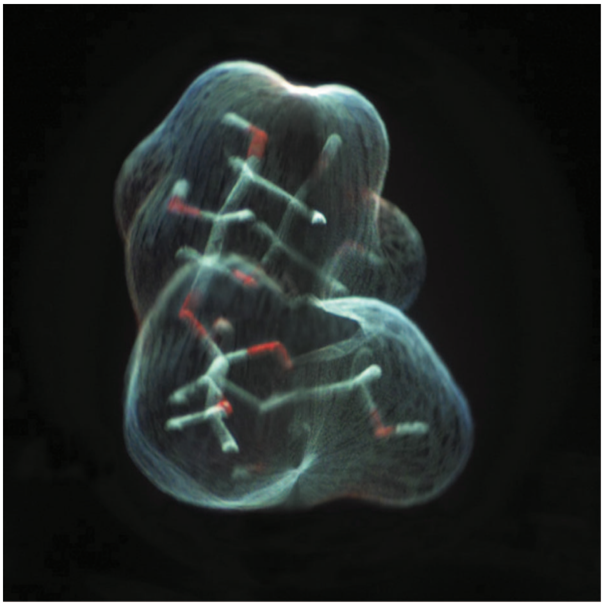

Fig. 3. (Color online) Photograph of a sugar molecule as reconstructed by the Perspecta multiplanar volumetric display before modification. The molecule's inner structure is visible through the necessarily translucent outer shell.

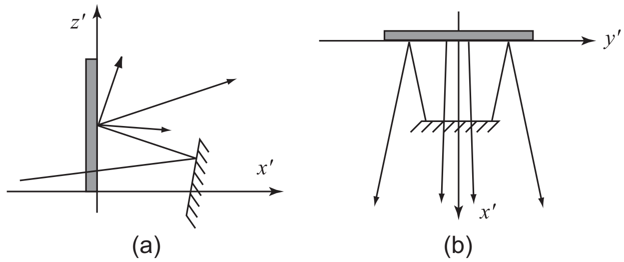

Fig. 4. (a) Side view of one ray showing the action of the final fold mirror and vertically diffusing screen. (b) Top view of the same, showing four rays incident on the screen.

illuminated by a series of fold mirrors that rotate with the screen.[24] Figure 3 shows a photograph of a molecule with a translucent surface, projected by Perspecta operating in its traditional, non-occlusion-capable mode.

We modified Perspecta to test the hypothesis that volumetric displays are indeed capable of viewer-position-dependent reconstruction effects. First, we altered the system's rendering software to generate views of the scene as would be observed from 198 center-looking observer viewpoints, approximately 0.5 m from the display's rotational axis, situated in a semicircle in a horizontal plane. There are several ways to render the scene; it could be captured using a computer-graphics camera with orthographic (parallel-beam) rendering in the horizontal and a perspective rendering in the vertical, to match the display's reconstruction geometry and intended viewer height. Such methods are well known in the field of computer graphics and holographic stereoscopy and are described elsewhere.[25] For simplicity, we used perspective rendering in both directions, effectively treating the display as a swept-pupil system rather than a piecewise projector of parallel rays.

Second, as pictured in Fig. 4, we replaced the projection surface with a diffuser (Physical Optics Corporation, Torrance, California) that has preferentially vertical diffusion and limited horizontal diffusion advertised to be $60° \times 0.1°$ (vertical $\times$ horizontal). The vertical diffusion acts to broaden the exit pupil's vertical extent, providing a wide vertical viewing zone. This is common practice in the field of horizontal parallax-only 3D displays. The restricted horizontal diffusion permits light to exit the display surface with only minimally broadening the narrow horizontal angular width of the light striking the screen, estimated to be 2° due to the slow *f* number of the system's projection optics.

Initially, we used a screen composed of a Mylar mirror covered by a vertical diffuser. In practice, the mirror layer was unnecessary because the system's steep vertical angle of incidence at the screen resulted in significant reflection off the diffuser. Thus the observer sees the imagery on the same side as the final fold mirror.

Figure 5 depicts the method of reconstruction for a scene containing a single line segment. As stated above, the scene is first recorded by a semicircular

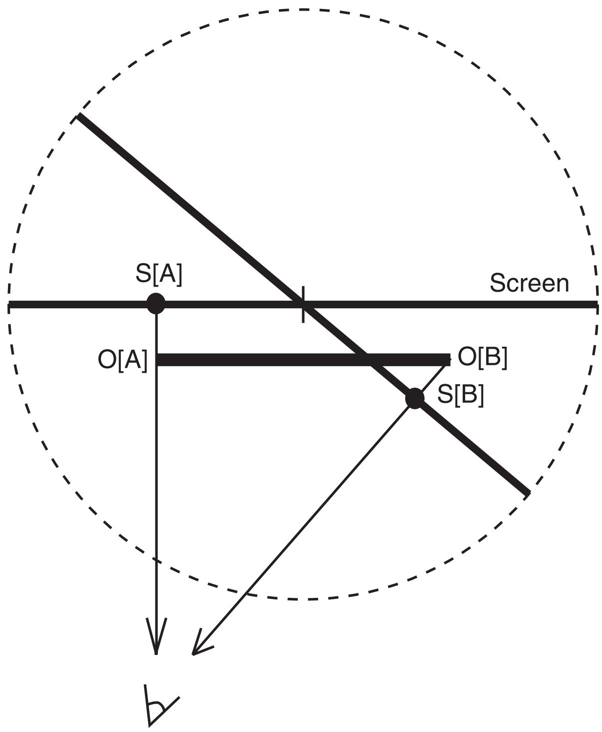

Fig. 5. (Top view) Reconstruction of a single line segment O. From the viewpoint of a stationary observer, end point O[A] of the line segment is visible when projected from screen point S[A]; likewise, end point O[B] is projected to that particular viewer position from screen point S[B] when the screen has rotated to a different location. For simplicity, this assumes light travels principally normal to the screen plane. Unlike many volumetric displays, the screen is often not colocated with the points it reconstructs.

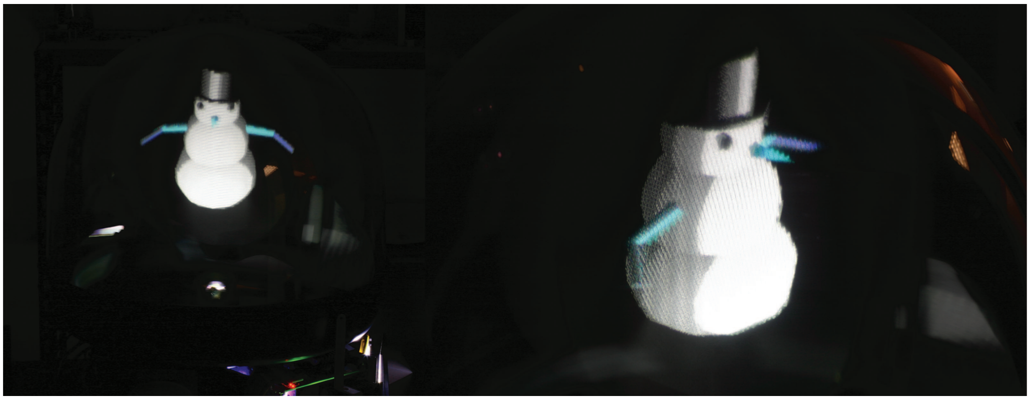

Fig. 6. (Color online) Photograph of multiview volumetric display. The snowman's right arm is clearly visible as a set of polygons that appear to occlude the snowman's white body.

arrangement of computer-graphics cameras, situated at the intended viewing regions, using the approach described above. The simplest method of scene reconstruction, at the expense of some accuracy, is the projection of the unmodified series of rendered images while the screen rotates.

Because of the screen's limited horizontal diffusion and the narrow NA at the screen, a stationary observer sees point O[A] of the object for the limited sector of screen angles. Assuming the screen is illuminated normal to the screen plane, the observer will see O[A] when the vertical ray fan centered on the screen normal through screen point S[A] intersects the observer's eye. A similar condition holds for viewing object point O[B] projected through screen point S[B]. The intermediate points of the line are projected from intermediate screen positions. Note that the screen points frequently do not have the same spatial location as the perceived object points.

Continuing this simple example, a viewer at a different horizontal location sees the line's end points when they are projected by two different screen positions than for the first observer position. Since the partially diffuse screen minimally perturbs the horizontal trajectory of the illumination, an observer will see imagery from each on-screen pixel for a brief angular window. The width of this window is a function of the NA at the screen and the extent of the horizontal diffusion. For simplicity, this discussion assumes that light exits the screen normal to the screen plane. In practice it does not because the display's projection optics relay illumination to the screen spanning a horizontal included angle of 26°.

Other scene recording and playback algorithms can improve the fidelity of the reconstructed scene, but they are outside the scope of this paper. One set of approaches, such as extracting 2D surfaces matched to the display geometry out of a the scene's 3D spatioperspective volume, is described in a published patent application.[26]

4. Results and Discussion

Figure 6 illustrates a scene from two viewpoints by using the above method. In this example, the arm of a snowman clearly appears to block its body. The system exhibits an approximately 180° horizontal field of view and no vertical parallax. That is, an observer moving his head vertically sees the same aspects of the image. Observers whose vertical or radial position departs from the expectations encoded in the computer-graphics rendering stage will see distorted imagery, as is experienced with other horizontal-

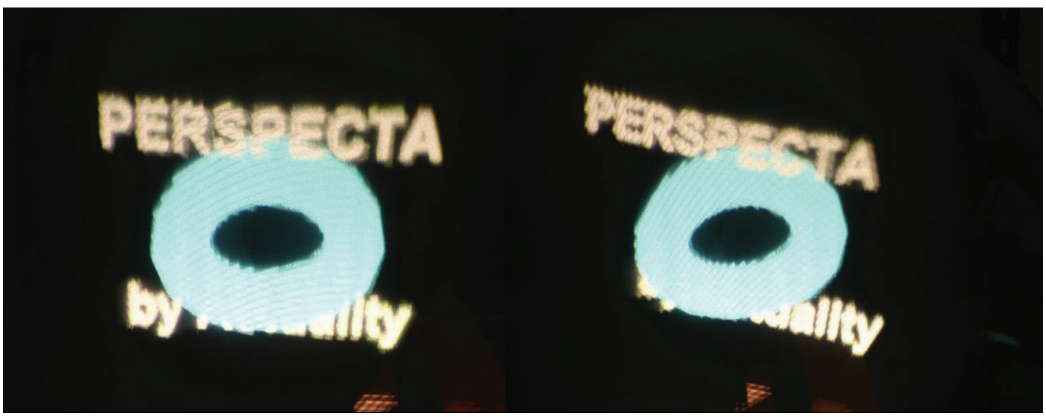

Fig. 7. (Color online) Scene in the multiview volumetric display, showing text in front of and behind a tilted, solid-shaded torus. The display's brightness contributed to the difficulty of minimizing the blur in this photograph.

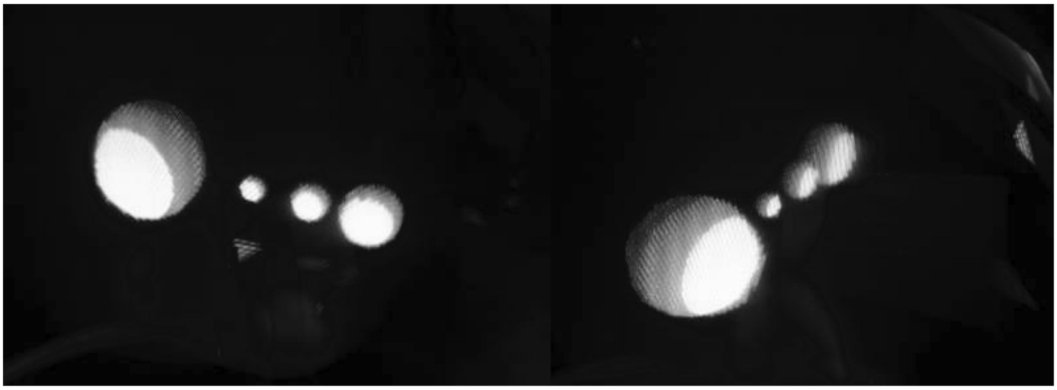

Fig. 8. Example of a sparse scene exhibiting occlusion; note rightmost moons in right half of figure.

parallax-only (HPO) systems.[27,28] Additional photographs are provided in Figs. 7 and 8. Although we did not measure its output, the system appeared considerably brighter than its unaltered counterpart. This made it difficult to photograph the system. Refer to Table 1 for a tabulation of the system specifications.

We note three observations about this display architecture. First, the imagery is highly astigmatic. The voxels generated by traditional volumetric displays are emitted from a set of locations very close to their perceived positions from an ordinary diffuse surface. Therefore the wavefronts generated by each voxel are generally spherical or hemispherical and are believed to elicit natural accommodation and vergence responses in the viewer. In the present display, however, each voxel is reconstructed by potentially numerous vertical ray fans with different radial headings and usually more than one apex.

Second, the display is theoretically capable of reconstructing scenes that, for some observation locations, exist outside of the volume swept by the screen. This is subject to the restriction that all elements of the 3D scene must lie along the line extending from the observer through the swept volume.[8] This mode of operation is depicted in Fig. 9; screen pixels are illuminated such that their rays intersect the desired regions outside of the screen. However, this remains a conjecture, since the authors have not demonstrated this effect in the experimental system.

Third, assuming that a stationary observer's line of sight to a scene element is normal to the screen as it rotates, the on-screen emission points representing a single off-center voxel trace out a circular locus of points. This property has been noted with regard to another multiview volumetric display.[16] Furthermore, under the same assumption, the aggregate surface of screen sections responsible for reconstructing a single viewpoint is a curve as illustrated in Fig. 10. Traditional volumetric displays act differently: Each reconstructed voxel is generally projected once, at the angle and at the on-screen position best matching the voxel's location, and all slices are approximately equally visible.

This occlusion-capable volumetric display has advantages and disadvantages. It enhances work in fields such as industrial design, medical imaging, and advertising, since the imagery is perceived as more photorealistic than as fully translucent imagery. The viewer-dependent effects of occlusion and complex reflections, for example, enable the observer to see

Table 1. Specifications of the Multiview Volumetric Three-Dimensional Display

Visual volume refresh rate[a]	30 Hz
Per-view resolution	768 pixels × 768 pixels
Angular resolution	198 views/180°
Horizontal field of view	180°
Addressable image diameter	25 cm
Screen rotational frequency	900 rpm
Color depth: perceived	Hundreds of colors[b]
Color depth: physical	3 bit (binary R/G/B)
Electronic interface	SCSI-3 Ultra

[a]Not to be confused with the volume data refresh rate, a measurement of the time required to generate and upload data for a new 3D scene.
[b]Using dithering.

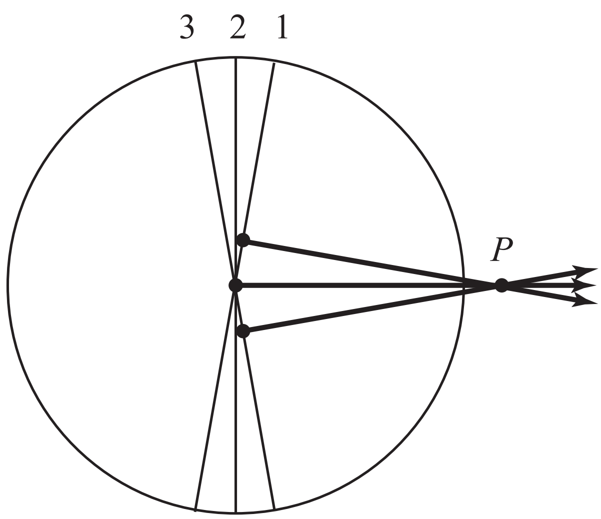

Fig. 9. Reconstruction of a point P perceived to be external to the volume swept by the projection screen. Top view.

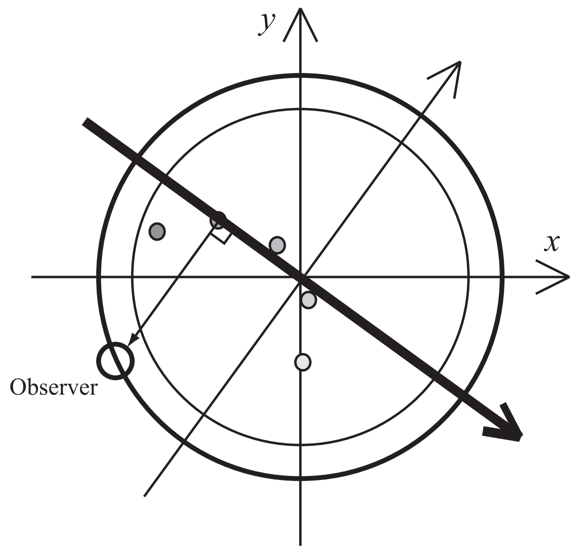

Fig. 10. Assuming an observer principally sees illumination exiting normal to the screen, the virtual aggregate projection surface for a single stationary observer is curved. As the screen rotates, it passes through the marked regions, which depict intersections with the observer's shortest line of sight to the screen. Top view. Screen is dark axis.

the front-facing surfaces within complex human anatomy or the shiny outer surfaces of a proposed sports car. The system's primary benefit over integral photography[29,30] (IP) and lenticular sheet displays is its broad horizontal field of view, which is demonstrated here to be 180° and theoretically is able to reach 360°. In addition, the system's 768 × 768 per-view resolution at the screen plane and 198 views across 180° matches or exceeds the performance of most IP and lenticular displays.

The system has disadvantages as well. Vertical parallax is absent, and the scene appears distorted for all but the correct viewing height. The horizontal field of view in the experimental system is 180° rather than 360° because the embedded projector is blanked during the sector responsible for rear projection.

Practical considerations may place limits on the size and refresh rate of this display architecture. As the diameter of the rotating volume increases, it may become more difficult to balance, will require stronger illumination, and will demand an embedded projector with higher spatial and temporal resolution in order to maintain high spatial resolution. Regarding the last point, the projector speed and volume refresh rate together determine the angular resolution of the display. Experimentally, we have found that the screen's rotational rate and optomechanical precision can be limited by the mass of the rotating subassembly. For example, the relay mirrors at the spinning platter's circumference tilt as the unit comes up to speed. This is a repeatable effect that can be compensated for in software and mechanical alignment. In short, these issues and historic results[31] suggest that the display size may reach up to 1 m in diameter, with a volume refresh rate (screen rotational rate) of 900–1500 rpm.

The preceding discussion explores our HPO multiview volumetric display. A full-parallax system can be constructed with a projection screen consisting of a lenticular lens sheet or parallax barrier with the long axis oriented perpendicular to the axis of screen rotation. Although this requires a higher-resolution image source, it can reconstruct multiple vertical ray trajectories for each small emissive region in the image volume. Typically, the pixel array projected onto the screen would have a much greater vertical density than horizontal density, since the vertical spatial dimension is mapped to a vertical angular dimension by the lens array.

5. Conclusion

We describe and demonstrate a hybrid volumetric 198-view 3D display with nearly XGA (1024 × 768) per-view resolution of 768 × 768 that reconstructs HPO 3D scenes with viewer-position-dependent effects such as object occlusion. We note the system's theoretical ability to reconstruct image regions outside the volume swept by the screen, describe the virtual projection surface for a stationary observer, and suggest modifications for full-parallax operation.

We acknowledge the help of Michael J. Richmond and Steven Kerry in modifying a Perspecta Display for this research, and Won-Suk Chun, T. J. Purtell II, and Sandy Stutsman for the underlying software platform. This work was supported in part by U.S. Army contract W56HZV-04-C-0144.

References

1. B. Blundell and A. Schwarz, *Volumetric Three-Dimensional Display Systems* (Wiley, 2000).
2. G. E. Favalora, J. Napoli, D. M. Hall, R. K. Dorval, M. G. Giovinco, M. J. Richmond, and W. S. Chun, "100 million-voxel volumetric display," in *Cockpit Displays IX: Displays for Defense Applications,* D. G. Hopper, ed., Proc. SPIE **4712,** 300–312 (2002).
3. G. E. Favalora, "Volumetric 3D displays and application infrastructure," Computer **38,** 37–44 (2005).
4. W. Hess, "Stereoscopic picture," U.S. patent 1,128,979 (16 February 1915).
5. F. E. Ives, "Parallax stereogram and process of making same," U.S. patent 725,567 (14 April 1903).
6. A. R. L. Travis, "Autostereoscopic 3-D display," Appl. Opt. **29,** 4341–4343 (1990), reprinted in *Selected Papers on Three-Dimensional Displays,* S. A. Benton, ed., SPIE Milestone Series **MS 162,** 249–251 (2001).
7. *Selected Papers on Three-Dimensional Displays,* S. A. Benton, ed., SPIE Milestone Series **MS 162** (2001).
8. M. Halle, "Autostereoscopic displays and computer graphics," Comput. Graphics **31,** 58–62 (1997).
9. *Selected Papers on Three-Dimensional Displays,* S. A. Benton, ed., SPIE Milestone Series **MS 162,** xxii (2001).
10. O. S. Cossairt and J. Napoli, "Radial multiview three-dimensional displays," U.S. patent application 2005/0180007 A1 (provisional 16 January 2004), (nonprovisional 14 January 2005), (published 18 August 2005).
11. G. E. Favalora, "Volumetric three-dimensional display architecture," U.S. patent 6,487,020 (26 November 2002).

12. Actuality Systems, Inc., *2nd Progress Report for A032-1449* (Army SBIR Progress Report, Actuality Systems Inc., 2004).
13. G. E. Favalora, "The ultimate display: What will it be?" presented at ACM SIGGRAPH (Los Angeles, 2005).
14. J. S. Kollin, "Time multiplexed auto-stereoscopic three-dimensional imaging system," U.S. patent 4,853,769 (1 August 1989).
15. R. Otsuka, T. Hoshino, and Y. Horry, "Transpost: all-around display system for 3D solid image," in *Proceedings of the ACM Symposium on Virtual Reality Software and Technology* (Hong Kong, 2004), pp. 187–194.
16. K. Tanaka and S. Aoki, "A method for the real-time construction of a full parallax light field," in *Stereoscopic Displays and Virtual Reality Systems XIII,* A. J. Woods, N. A. Dodgson, J. O. Merritt, M. T. Bolas, and I. E. McDowall, eds., Proc. SPIE **6055,** 605516 (2006).
17. S. P. McGrew, "System for synthesizing strip-multiplexed holograms," U.S. patent 4,206,965 (10 June 1980).
18. T. Honda, Y. Kajiki, K. Susami, T. Hamaguchi, T. Endo, T. Hatada, and T. Fujii, "Three-dimensional display-technologies satisfying 'super multiview condition,'" in *Three-Dimensional Video and Display: Devices and Systems,* B. Javidi and F. Okano, eds., Proc. SPIE **CR76,** 218–249 (2000).
19. R. B. Collender, "Three dimensional unaided viewing method and apparatus," U.S. patent 3,178,720 (13 April 1965).
20. R. B. Collender, "The stereoptiplexer," Inf. Disp. **4,** 27–31 (1967).
21. R. B. Collender, "Methods for electronic 3-D moving pictures without glasses," Displays **8,** 193–209 (1987).
22. R. K. Dorval, M. Thomas, and J. L. Bareau, "Volumetric three-dimensional display system," U.S. patent 6,554,430 (29 April 2003).
23. W.-S. Chun, J. Napoli, O. S. Cossairt, R. K. Dorval, D. M. Hall, T. J. Purtell II, J. F. Schooler, Y. Banker, and G. E. Favalora, "Spatial 3D infrastructure: display-independent software framework, high-speed rendering electronics, and several new displays," in *Stereoscopic Displays and Virtual Reality Systems XII,* A. J. Woods, M. T. Bolas, J. O. Merritt, and I. E. McDowall, eds., Proc. SPIE **5664,** 302–312 (2005).
24. M. Hirsch, "Three-dimensional display," U.S. patent 3,822,938 (9 July 1974).
25. M. W. Halle, "Multiple viewpoint rendering for three-dimensional displays," Ph.D. dissertation (Massachusetts Institute of Technology, 1997).
26. W.-S. Chun and O. S. Cossairt, "Data processing for three-dimensional displays," U.S. patent application US2005/0230641 A1 (filed 16 March 2005).
27. M. W. Halle, S. A. Benton, M. A. Klug, and J. S. Underkoffler, "The ultragram: a generalized holographic stereogram," in *Practical Holography V,* S. A. Benton, ed., Proc. SPIE **1461,** 142–155 (1991).
28. D. J. DeBitetto, "Bandwidth reduction of hologram transmission systems by elimination of vertical parallax," Appl. Phys. Lett. **12,** 176–178 (1968), reprinted in *Selected Papers on Three-Dimensional Displays,* S. A. Benton, ed., SPIE Milestone Series **MS 162,** pp. 402–404 (2001).
29. M. G. Lippmann, "Épreuves réversibles. Photographies intégrales," Compt. Rend. **146,** 446–451 (1908), reprinted in *Selected Papers on Three-Dimensional Displays,* S. A. Benton, ed., SPIE Milestone Series **MS 162,** pp. 305–310 (2001).
30. H. Choi, Y. Kim, J.-H. Park, J. Kim, S.-W. Cho, and B. Lee, "Layered-panel integral imaging without the translucent problem," Opt. Express **13,** 5769–5776 (2005).
31. M. Lasher, P. Soltan, W. Dahlke, N. Acantilado, and M. McDonald, "Laser projected 3-D volumetric displays," in *Projection Displays II,* M. H. Wu, ed., Proc. SPIE **2650,** 285–295 (1996).