An Interactive 360° Light Field Display

Andrew Jones Ian McDowall[†] Hideshi Yamada^{*} Mark Bolas[‡] Paul Debevec

USC Centers for Creative Technologies Fakespace Labs† Sony Corporation* USC School of Cinematic Arts‡



Figure 1: A 3D object shown on the display is simultaneously photographed by two stereo cameras (seen in the middle image). The two stereo viewpoints are from the 360° field of view around the display. The right pair is from a vertically-tracked camera position and the left pair is from an untracked position roughly horizontal to the center of the display. A filter was used to dim the projector. The images are left-right reversed for cross-fused stereo viewing.

1 Introduction

While a great deal of computer generated imagery is modeled and rendered in 3D, the vast majority of this 3D imagery is shown on 2D displays. Various forms of 3D displays have been contemplated and constructed for at least one hundred years [Lippman 1908], but only recent evolutions in digital capture, computation, and display have made functional and practical 3D displays possible.

We present an easily reproducible, low-cost 3D display system with a form factor that offers a number of advantages for displaying 3D objects in 3D. We develop and demonstrate the projection mathematics and rendering methods necessary to drive the display with real-time raster imagery or pre-recorded light fields so that they exhibit the correct cues of both horizontal and vertical parallax. The display is autostereoscopic, requiring no special viewing glasses, omnidirectional, allowing viewers to be situated anywhere around it, and *multiview*, producing a correct rendition of the light field with the correct horizontal parallax and vertical perspective for any viewpoint situated at a certain distance and height around the display. Furthermore, if head tracking is employed to detect the height and distance of one or more of the viewers around the the display, our display allows the rendered perspective to be adjusted on the fly to allow the tracked users to properly see objects from arbitrary heights and distances in addition to obtaining correct views from any angle around the display. Our display uses primarily commodity graphics and display components and achieves real-time rendering with non-trivial scene complexity across its entire field of view. Our contributions include:

- An easily reproducible, low-cost 360° horizontal-parallax light field display system leveraging commodity graphics and projection display hardware.
- A novel software/hardware architecture that enables real-time update of high-speed video projection using standard graphics hardware at kilohertz frame rates.

- A novel projection algorithm for rendering multiple center of projection raster graphics for a 360° horizontal-parallax light field display with correct vertical perspective for any given viewer height and distance.
- A light field display technique that is horizontally multiview autostereoscopic and employs vertical head tracking to produce correct vertical parallax for tracked users.

2 Background and Related Work

Recent surveys of the rich and varied field of three-dimensional display techniques can be found in [Travis 1997; Favalora 2005; Dodgson 2005]. Our display belongs to an emerging class of horizontalparallax 3D displays that combine one or more video projectors to generate view-dependent images on a non-stationary anisotropic screen. Viewers receive varying views of the scene depending on the position of their eyes with respect to the display.

The basic idea has existed within the field of holography for over a decade [Batchko 1994]. Recent systems that employ this idea include [Maeda et al. 2003], which uses an anisotropic privacy-guard film on a spinning LCD monitor to show six different viewpoints of a human subject. The Transpost system [Otsuka et al. 2006] uses a circle of mirror facets to reflect 24 images from the circumference of a video projected image onto a rapidly rotating anisotropic screen. The system aims for a similar form factor and effect as does ours, but achieves only 24 low-resolution (100x100) images around the circle. Their design does not scale well to additional views as the images must be arranged linearly around the circumference of a circle. However, it achieves 24-bit color images whereas we are limited to dithered binary images. The LiveDimension system [Tanaka and Aoki 2006] uses an inward-pointing circular array of 12 projectors and a vertically-oriented light-control film, similar to that used in [Maeda et al. 2003], to reflect each projector's image outwards to the viewer. While they achieve twelve full-color views, they do not produce a sufficient number of views for binocular parallax, and a greater number of views would require a greater number of projectors and use progressively less light from each of them. [Cossairt et al. 2004] describes a display that couples a threechip high-speed DLP projector with a moving slit and a large lens to direct images in 26 horizontal directions at 50Hz, but it uses highly specialized hardware and has a limited field of view. [Agocs et al. 2006; Balogh et al. 2006] place a horizontal array of projectors behind a large holographic diffuser, similar to ours, to create a multi-user horizontal-parallax display for a sizable zone in front of the diffuser. Their images are large, bright, interactive, and fullcolor, but the large number of projectors makes the system difficult to calibrate and expensive to reproduce. The Seelinder display [Endo et al. 2000; Yendo et al. 2005] uses the significantly different approach of spinning multiple 1D vertical arrays of LEDs past a cylindrical parallax barrier to produce 3D images. They achieve better than 1 degree view spacing at a pixel resolution of 128 pixels vertically, but require a very specialized hardware setup. None of these systems compensates for vertical perspective and parallax and require either many projectors or very specialized hardware.

Our design closely parallels the work of [Cossairt et al. 2007]. Both systems use a single high-speed DLP projector to project patterns onto a spinning anisotropic surface. Unfortunately, [Cossairt et al. 2007] is limited by the transfer speed of SCSI-3 Ultra and is not interactive. Our display can generate and update images in real-time by decoding a DVI signal. Our display is the first to show animated scenes for a 360° view region and dynamically adjust to changing vertical viewer positions.

Our rendering algorithm builds on previous work in holography and light field rendering. [Halle et al. 1991] proposed a method where static holographic stereograms account for the viewer distance but not height. Much of the existing light field literature [Levoy and Hanrahan 1996; Gortler et al. 1996; Isaksen et al. 2000] is useful for acquiring, storing, and sampling multi-view content. Results from [Chai et al. 2000; Zwicker et al. 2006] informed our choices for the amount of horizontal diffusion, the number of views we render around the circle, and the camera aperture used to record our light fields. Our technique for multiple-center-of-projection view rendering using GPU vertex shaders is informed by the recent work of [Hou et al. 2006].

3 System Overview

Our 3D display system consists of a spinning mirror covered by an anisotropic holographic diffuser, a motion-control motor, a high-speed video projector, and a standard PC and graphics card interfaced to the projector using a custom FPGA-based image decoder. As seen in the overview, (Figure 2), the spinning mirror tilted at 45° reflects rays of light from the projector to all possible viewing positions around the device. A group of people can also be seen viewing images on the device. The rest of this section provides details of the system components.

High-Speed Projector We achieve high speed video projection by modifying an off-the-shelf projector to use a new DLP drive card with custom programmed FPGA -based circuitry. The projector decodes a standard DVI signal from the graphics card. Instead of rendering a color image, the projector takes each 24-bit color frame of video and displays each bit sequentially as separate frames. Thus, if the incoming digital video signal is 60Hz, the projector displays $60 \times 24 = 1,440$ frames per second. To achieve faster rates, the video card's output is set to rates of 200Hz and above. At 200Hz,





Figure 2: (*Left*) *The display showing a virtual object in 3D to an audience standing around the device.* (*Right*) *Schematic showing the display's high-speed projector, spinning mirror, and synchronized motor.*



Figure 3: *Twenty-four consecutive binary frames of interactive OpenGL graphics (Left) are packed into a single 24-bit color image (Right).*

the projector displays 4,800 binary frames per second. We continuously render new horizontal views of the subject (288 images per rotation). These views are encoded into 24-bit images and sent to the projector. A complete kit consisting of the FPGA and DLP boards is now available from Tyrex Service, Inc, a company that distributes Texas Instruments DLP prototyping boards for research laboratories.

Spinning Mirror System Previous volumetric displays projected images onto a spinning diffuse plane which scattered light in all directions. Such displays could not recreate view-dependent effects such as occlusion. Instead, our projection surface is a anisotropic holographic diffuser bonded onto a first surface mirror. Horizontally, the diffused mirror is sharply specular to maintain a 1.25 degree separation between views. Vertically, the mirror scatters widely so the projected image can be viewed from multiple heights. This surface spins synchronously relative to the images being displayed by the projector. The PC video output rate is the master signal. The projector's FPGA decodes the current DVI frame rate and interfaces directly to an Animatics SM3420D Smart Motor. As the mirror rotates up to 20 times per second, persistence of vision creates the illusion of a floating object at the center of the mirror.

4 Rendering Interactive Raster Graphics

A key element of the design was the ability to read data from a standard DVI graphics card which can render images at extremely fast rates using standard graphics calls. This approach has worked very well, allowing a low-cost desk side computer to render and communicate imagery at thousands of frames per second.

The images sent to the projector can either be pre-computed or rendered in real-time using raster OpenGL graphics. The simplest rendering algorithm proposed in [Cossairt et al. 2007] projects a sequence of perspective images from a camera rotating around the scene. This perspective geometry assumes that all rays reflected off the mirror reconverge at a single viewpoint. In reality, frames reflected off the mirror diverge towards multiple viewpoints in space. The result (shown in Figure 4 (a)) is a 3D image where the horizontal field of view is exaggerated and straight lines appear curved. A second artifact is due to the lack of vertical parallax. Just as an image on a piece of paper will appear smaller when viewed from an oblique angle, the projected image will appear to stretch vertically when the viewer changes height. A more subtle error is that the vertical perspective stays constant even when the viewer distance changes, causing a mismatch between vertical and horizontal perspective cues (as noted by [Halle et al. 1991]).

We have developed an improved rendering algorithm (shown in Figure 4 (b)) that leverages our system's unique real-time update to dynamically adjust the imagery based on each viewers height and distance. Instead of using a simple perspective camera, we trace each reflected projector ray to find the multiple correct viewer positions for a given frame. We implement this computation as a GPU vertex shader, bypassing the traditional OpenGL projection matrix. The entire mesh is rasterized in a single render pass allowing simple scenes to be effectively rendered at rates of over 5000 fps. If combined with a passive tracking system, the system allows multiple tracked viewers to each experience his own correct view of a three-dimensional scene, with horizontal and vertical parallax, without 3D glasses. In the accompanying video we demonstrate vertically tracked raster graphics using a Polyhemus magnetic tracking system.

5 Light Field Rendering

Rendering horizontal parallax We can capture photographic data of a real object for 3D display using an inexpensive motorized turntable (Figure 5(a)). We capture a movie sequence of at least 288 frames of the object rotating 360° on the turntable, which takes a few seconds. Alternatively, a multiple-camera system [Yang et al. 2002; Wilburn et al. 2005] or a high-speed single-camera system [Jones et al. 2006] using spinning mirrors to change the viewpoint could be used to capture such data.

We then *rebin* the images in the horizontal direction in order to compensate for the gain in perspective introduced by the diverging reflected rays of the projector. To rebin, we form images which, when projected by the projector, faithfully project rays of light into space that are consistent with the rays captured by the camera around the object. We project a ray through each pixel of each reflected projector position until it intersects the viewing circle formed by the camera positions. We take the radius of the viewing circle to be the distance of the camera from the center of the turntable so that correct vertical perspective will be seen by untracked viewers at this distance. We find the two nearest camera positions and project the ray/mirror intersection into each camera. We bilinearly interpolate between the two cameras to determine the final pixel value.



Figure 4: A scene is viewed from above (top) and straight-on (below) while projecting (a) simple perspective images [Cossairt et al. 2007], and (b) multiple-center of projection images warped using the vertex shader and correct perspective cues.

Once rebinning is performed, we halftone each rebinned image using the error-diffusion dithering algorithm [Ostromoukhov 2001], and we pack sets of 24 halftoned images into 24-bit color images, yielding twelve 24-bit color images for the object's rotation. For display, these packed images are loaded into the texture memory of the graphics card and successively drawn into the frame buffer at the DVI refresh rate. At 512×512 resolution, one rotation requires just 9MB of texture memory. An image of a light field displayed in this manner is shown in Figure 5.

Displaying a 4D Light Field As with raster graphics, we perform real-time vertical parallax correction to reproduce the full 4D light field for any viewer whose height and distance are known. We capture a 4D light field by shooting multiple rotations of the turntable from several heights. We preprocess each rotation as described above. Since each rotation is just 9MB, we can easily load



Figure 5: (*Left*) One frame of an object filmed rotating on a turntable. (*Right*) A frame of the object seen on the 3D display.

twenty or more vertical light field slices into the nVIDIA graphics card texture memory. While the display is running, we perform dynamic vertical rebinning that combines photographed images from different heights to match the viewer's height and distance. A sequence of dynamically-rebinned 4D light field imagery displayed for a tracked camera viewpoint, as well as a visualization of the dynamic light field slices, is shown in the accompanying video.

6 Conclusion

Our display prototype is good for showing small (12cm) objects to a group of people who are looking at the virtual object. The display creates an image which appears to hang in space and provides viewers with an ability to walk anywhere around the object. Our 360° display creates a light field which exhibits correct horizontal parallax anywhere around the display. Occlusion behaves correctly, and the resolution is sufficient so that one is not even conscious of the transitions between vertical zones. Novel rendering techniques support real time interactive graphics. We also demonstrated a method for capturing and rendering a real object's light field on the display. Among the significant enabling contributions, we develope requisite multiple-center-of-projection mathematics for rendering to the display's convenient form factor, and a unique method of rendering single-bit imagery at over 5000Hz from a commodity graphics card. Further mathematical details are available in our SIGGRAPH 2007 paper [Jones et al. 2007].

Our goal is to disseminate this system architecture broadly to researchers in the field of computer graphics and immersive displays. The core components should continue to decrease in price while increasing in capability in line with display and rendering technologies. This system makes it possible to render and display real-time 360° light fields with a useful number of viewpoints. Such a capability will hopefully allow the pursuit of capturing and rendering more accurate and realistic models.

Acknowledgements

The authors wish to thank Bruce Lamond, Naho Inamoto, Tim Hawkins, Pieter Peers, Dell Lunceford, Tom Pereira, Jacki Morie, Sean Bouchard, Bill Swartout, Randy Hill, and Randolph Hall for their support and assistance with this work. This work was sponsored by the University of Southern California Office of the Provost and the U.S. Army Research, Development, and Engineering Command (RDECOM). The high-speed projector was originally developed by a grant from the Office of Naval Research under the guidance of Ralph Wachter. The content of the information does not necessarily reflect the position or the policy of the US Government, and no official endorsement should be inferred.

References

- AGOCS, T., BALOGH, T., FORGACS, T., BETTIO, F., GOBBETTI, E., ZANETTI, G., AND BOUVIER, E. 2006. A large scale interactive holographic display. In VR '06: Proceedings of the IEEE Virtual Reality Conference (VR 2006), IEEE Computer Society, Washington, DC, USA, 57.
- BALOGH, T., DOBRANYI, Z., FORGACS, T., MOLNAR, A., SZLOBODA, L., GOB-BETTI, E., MARTON, F., BETTIO, F., PINTORE, G., ZANETTI, G., BOUVIER, E., AND KLEIN, R. 2006. An interactive multi-user holographic environment. In SIGGRAPH '06: ACM SIGGRAPH 2006 Emerging technologies, ACM Press, New York, NY, USA, 18.
- BATCHKO, R. G. 1994. Three-hundred-sixty degree electro-holographic stereogram and volumetric display system. In *Proc. SPIE*, vol. 2176, 30–41.
- CHAI, J.-X., TONG, X., CHAN, S.-C., AND SHUM, H.-Y. 2000. Plenoptic sampling. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, 307–318.

- COSSAIRT, O., TRAVIS, A. R., MOLLER, C., AND BENTON, S. A. 2004. Novel view sequential display based on dmd technology. In *Proc. SPIE, Stereoscopic Displays* and Virtual Reality Systems XI, A. J. Woods, J. O. Merritt, S. A. Benton, and M. T. Bolas, Eds., vol. 5291, 273–278.
- COSSAIRT, O. S., NAPOLI, J., HILL, S. L., DORVAL, R. K., AND FAVALORA, G. E. 2007. Occlusion-capable multiview volumetric three-dimensional display. *Applied Optics* 46, 8 (Mar), 1244–1250.
- DODGSON, N. A. 2005. Autostereoscopic 3d displays. Computer 38, 8, 31-36.
- ENDO, T., KAJIKI, Y., HONDA, T., AND SATO, M. 2000. Cylindrical 3d video display observable from all directions. In 8th Pacific Conference on Computer Graphics and Applications, 300–306.
- FAVALORA, G. E. 2005. Volumetric 3d displays and application infrastructure. Computer 38, 8, 37–44.
- GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. 1996. The lumigraph. In *Proceedings of SIGGRAPH 96*, Computer Graphics Proceedings, Annual Conference Series, 43–54.
- HALLE, M. W., BENTON, S. A., KLUG, M. A., AND UNDERKOFFLER, J. S. 1991. The ultragram: A generalized holographic stereogram.
- HOU, X., WEI, L.-Y., SHUM, H.-Y., AND GUO, B. 2006. Real-time multiperspective rendering on graphics hardware. In *Rendering Techniques 2006: 17th Eurographics Workshop on Rendering*, 93–102.
- ISAKSEN, A., MCMILLAN, L., AND GORTLER, S. J. 2000. Dynamically reparameterized light fields. In *Proceedings of ACM SIGGRAPH 2000*, Computer Graphics Proceedings, Annual Conference Series, 297–306.
- JONES, A., DEBEVEC, P., BOLAS, M., AND MCDOWALL, I. 2006. Concave surround optics for rapid multiview imaging. In *Proceedings of the 25th Army Science Conference*.
- JONES, A., MCDOWALL, I., YAMADA, H., BOLAS, M., AND DEBEVEC, P. 2007. Rendering for an interactive 360 degree light field display. In ACM Transactions on Graphics.
- LEVOY, M., AND HANRAHAN, P. M. 1996. Light field rendering. In Proceedings of ACM SIGGRAPH 96, Computer Graphics Proceedings, Annual Conference Series, 31–42.
- LIPPMAN, G. 1908. Epreuves reversibles donnant la sensation du relief. Journal of Physics 7, 4 (Nov), 821–835.
- MAEDA, H., HIROSE, K., YAMASHITA, J., HIROTA, K., AND HIROSE, M. 2003. All-around display for video avatar in real world. In ISMAR '03: Proceedings of the The 2nd IEEE and ACM International Symposium on Mixed and Augmented Reality, IEEE Computer Society, Washington, DC, USA, 288.
- OSTROMOUKHOV, V. 2001. A simple and efficient error-diffusion algorithm. In Proceedings of ACM SIGGRAPH 2001, Computer Graphics Proceedings, Annual Conference Series, 567–572.
- OTSUKA, R., HOSHINO, T., AND HORRY, Y. 2006. Transpost: A novel approach to the display and transmission of 360 degrees-viewable 3d solid images. *IEEE Transactions on Visualization and Computer Graphics* 12, 2, 178–185.
- TANAKA, K., AND AOKI, S. 2006. A method for the real-time construction of a full parallax light field. In Stereoscopic Displays and Virtual Reality Systems XIII. Edited by Woods, Andrew J.; Dodgson, Neil A.; Merritt, John O.; Bolas, Mark T.; McDowall, Ian E. Proceedings of the SPIE, Volume 6055, pp. 397-407 (2006)., A. J. Woods, N. A. Dodgson, J. O. Merritt, M. T. Bolas, and I. E. McDowall, Eds., 397-407.
- TRAVIS, A. R. L. 1997. The display of three-dimensional video images. Proceedings of the IEEE 85, 11 (Nov), 1817–1832.
- WILBURN, B., JOSHI, N., VAISH, V., TALVALA, E.-V., ANTUNEZ, E., BARTH, A., ADAMS, A., HOROWITZ, M., AND LEVOY, M. 2005. High performance imaging using large camera arrays. ACM Transactions on Graphics 24, 3 (Aug), 765–776.
- YANG, J. C., EVERETT, M., BUEHLER, C., AND MCMILLAN, L. 2002. A real-time distributed light field camera. In *Rendering Techniques 2002: 13th Eurographics* Workshop on Rendering, 77–86.
- YENDO, T., KAWAKAMI, N., AND TACHI, S. 2005. Seelinder: the cylindrical lightfield display. In SIGGRAPH '05: ACM SIGGRAPH 2005 Emerging technologies, ACM Press, New York, NY, USA, 16.
- ZWICKER, M., MATUSIK, W., DURAND, F., AND PFISTER, H. 2006. Antialiasing for automultiscopic 3d displays. In *Rendering Techniques 2006: 17th Eurographics* Workshop on Rendering, 73–82.