

Projects in VR

Editors: Lawrence Rosenblum and
Michael Macedonia

New Research and Explorations into Multiuser Immersive Display Systems

Mark Bolas

Stanford
University

Ian McDowall

Fakespace Labs

Dan Corr

Known
Associates

In a familiar scene from an old movie, generals huddle around a large map, pushing models of tanks and infantry regiments about to indicate the current battle situation. Today, the scene might include electronic displays and networked sensing technology, but the basic form would remain the same: A small group of domain experts surround and gesture toward a common data set, hoping to achieve consensus. This mode of decision making is pervasive, ranging in use from US Marine Corps command and control applications to product design review meetings. Such applications demonstrate the need for VR systems that accommodate small groups of people working in close proximity.¹

Yet, while non-head-mounted, immersive displays perform well for single-person work, when used by small groups they are hampered by an unacceptably large degree of distortion between the head-tracked viewpoint and an untracked collaborator's perspective (see Figure 1). What looks like a sphere to one user will look like an egg to another.² Solving this problem is critical. Decision makers and designers cannot jointly view and respond to data when all but one see incorrect images.

Working with Stanford and Fraunhofer, our group first addressed this problem in 1996 with the Duo system, which let two users share a virtual environment. Our earliest prototype demonstrated the power of VR systems that provide different users with correct-perspective viewpoints. Using this prototype, two of us pointed to a feature in a virtual model; when our fingers touched at precisely the same point in both real and vir-

tual space, something clicked in our minds and the virtual became real.

This powerful experience compelled us to experiment with a wide range of approaches to non-head-mounted, co-located, multiuser systems.³ Here, we present a framework for understanding these and other groups' approaches, discussing their effectiveness in achieving powerful, useful immersive display systems for multiple users working together.

Solution framework

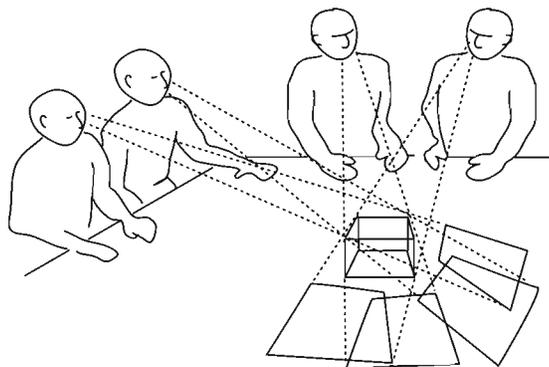
The basic design challenge is to create a system that can display multiple images in a common area, occluding all but the appropriate pair of these images for each user. Thus, the system would deliver a unique image for each eyeball viewing the scene. Solutions fall into four general categories:

- *Spatial barriers* use the display's physical configuration and user placement to block users from seeing each other's view.
- *Optical filtering* involves systems that filter viewpoints using light's electromagnetic properties, such as polarization or wavelength.
- *Optical routing* uses the angle-sensitive optical characteristics of certain materials to direct or occlude images based on the user's position.
- *Time multiplexing* solutions use time-sequenced light and shutters to determine which user sees an image at a given point in time.

Systems also can mix solutions from these categories. Time multiplexing, for example, could serve to create stereoscopic images, with spatial barriers employed to ensure that each user sees only the correct image.

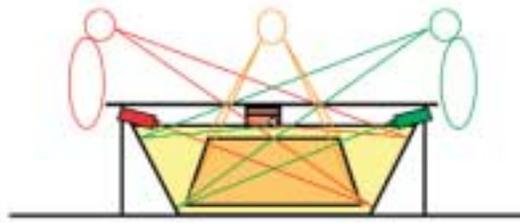
Each solution will work with either screen- or pixel-based approaches. Screen-based solutions work on an entire image at once. Active stereoscopic systems, for example, display alternating left and right images synchronized with occluding shutter glasses worn by the user. Pixel-based solutions use any of these methods on a pixel level. For example, a traditional barrier-slit autostereoscopic display blocks adjacent pixels with a spatial barrier to provide two distinct views, one for each eye. Here we will highlight our and others' work with screen-based solutions, in particular examining their

1 To avoid misleading imagery, each user's perspective requires a distinct image.





(a)



(b)



(c)

2 Pyramid approach: (a) front-projected pyramid surface; (b) inverted pyramid surfaces provide four distinct viewpoints; and (c) single user's view.

success in meeting two important user requirements: number of users and mode of collaboration.

The intended number of users changes how the system operates. Our work focuses on groups of two to eight people, roughly the size found in real-life design and decision-making situations. In larger groups, the dynamic moves from collaboration to presentation, calling for a different type of system.

Collaboration mode is a subtler requirement. Some people use a system to share a model; think again of the movie scene with generals arranged around a table. Other applications require users to share a viewpoint; for example, when one person calls the others around to his or her side of the table so they can see the proposed angle of attack, or designers stand shoulder to shoulder, viewing the model from the same viewpoint.

Spatial barriers

The University of North Carolina's Protein Interactive Theater is a good example of a spatial barrier system. Simply by arranging time-multiplexed stereoscopic display surfaces at a large angle, this elegant solution naturally occludes images for users looking across a common volume on an L-shaped screen. As long as the users restrain their gaze to the intersecting volume between them, this configuration allows two independent views.

Early in 1999, we used a similar approach to accommodate four users with distinct views by producing a simple pyramid-shaped projection surface (See Figure 2a). This solution worked, but was awkward to use. The projection surfaces were too close to the users, blocking opposing users from seeing each other.

Next, we tried inverting the pyramid, sinking it into a table and adding a surrounding skirt to occlude off-axis views and yet let everyone easily see each other across the table (Figure 2b). Unfortunately, this step also limited each user's viewing cone to a relatively narrow region, forcing users away from the display and apart from one another (Figure 2c). These limitations made it unsuitable for most applications, including command and control situations requiring both a large virtual model and the ability to share viewpoints between users.

Such configurations do appear to work for small objects. At Siggraph in 2001, Osaka University demonstrated the Illusion Hole, which was optimized for close interaction with small objects. Fraunhofer's Virtual Showcase uses a clever pyramid-mirror arrangement above a 120-Hz stereoscopic display to emulate a show-

case display, making it feel natural to users who would stand, each to a side, to view smaller objects.

Optical filtering

Standard 120-Hz flicker stereo can be extended to support two users by coupling circular polarization filters with both the projectors and the flicker glasses. Circular polarization is required as it can work in conjunction with the glasses' linear polarization. This type of solution tends to cause unacceptable bleed between users caused by the poor extinction ratio of circular polarizers. (However, we have recently begun exploring polarizers that appear to work fairly well.)

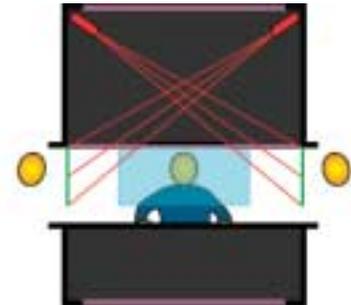
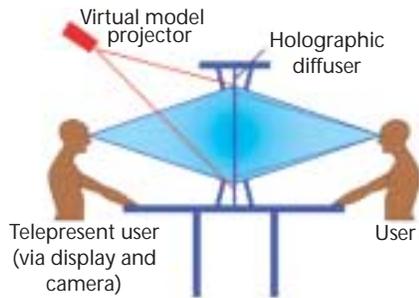
Tan (now Barco) and DaimlerChrysler have demonstrated the Infinitec system, which used three-color notch filters centered about different frequencies for each eye. This arrangement creates a unique set of three primary colors for each eye. Using a pair of projectors with corresponding filters, this approach overcomes the lack-of-color limitation of traditional anaglyph techniques and could be extended to support two users.

Optical filtering is a relatively straightforward method of extending one-person systems to accommodate two users, and it does not limit users to working in physically limited spaces. However, in general, the analog characteristics of optical materials make it difficult to accommodate more than two users.

Optical routing

We obtained a holographic screen material that diffuses light impinging its surface at 60 degrees, but is transparent for light perpendicular to its surface. We placed this material in front of the participant to create a display surface on which 120-Hz stereoscopic imagery is projected (at a 60-degree angle) to create a virtual workspace. The participant sees this imagery and can simultaneously look straight through the screen at another user standing on the system's other side (see Figure 3a).

In creating an initial two-person prototype system, we found that if a projected—rather than live—participant sits a few feet behind the holographic diffusing screen, an effective telecollaboration system results. Users can collaborate as if they were sitting across a table, with a shared virtual workspace between them (see Figure 3b). The prototype was surprisingly engaging, possibly because the visual accommodation and convergence cues approximate those of the represented environment. Users must refocus their eyes if they are



(a) (b) (c)
3 Creating a virtual workspace: (a) a holographic diffuser allows simultaneous viewing of projected and real images; (b) shared virtual workspace projected between collaborating users; and (c) four-person configuration.



4 The FLEA: (a) 120-Hz stereo DLP projector; (b) stacked for two-person collaborative use.

This system's success indicated that if a faster time-multiplexed system was technically feasible, it could offer shoulder-to-shoulder view-point sharing among eight or more users. To reach such rates, we needed a display technology that could overcome the decay limitations of phosphor-based displays, which are constrained to roughly 180 Hz.

Toward this end, we were fortunate to work with Christie Digital and Texas Instruments during the Mirage projector's development

looking at the screen displaying the virtual model (at arm's length) or at the remote participant (4 feet back). This effect has exciting implications for increasing verisimilitude in virtual display systems. We intend to study this approach further to determine the utility of merging a few distinct accommodation distances with typical convergence cues.

Limiting users to sitting across from each other worked well for eye-contact and gesture cues, but they had difficulty collaborating on the details of a small model. They could only see different sides of the model and could not share a similar viewpoint. Software techniques such as virtual pointers and viewpoint portals can help ameliorate this problem.

We can extend the system to support more than two users by orienting multiple sheets of the holographic material around a central virtual workspace (Figure 3c). Additionally, time-multiplexed diffusion screens can substitute for the 60-degree holographic material, and time-multiplexed lighting can increase contrast while reducing screen reflections.

Time multiplexing

Standard occluding shutter glasses used for active stereoscopic systems typically run at 120 Hz. In our earliest work, we created shutter glasses capable of flicker rates greater than 1 kHz, integrating them with CRT projectors running at 180 Hz. Working with Stanford and Fraunhofer, this let us display stereoscopic images to two simultaneous users at roughly 45 Hz and begin looking at the design issues presented by multiuser collaboration systems.²

We researched how best to create a stereoscopic projector based on TI's digital micromirror devices and our rapid-rate flicker glasses. This work helped lead to a feature which controls the *dark interval* timing, letting the system regain signal lost during the vertical retrace period, without incurring crosstalk. The results helped cement our desire to turbocharge digital light processing (DLP) projectors to achieve extremely fast refresh rates, approaching our glasses' kilohertz rates.

Time-multiplexed FLEA projector

Toward this end, we modified the drive electronics of off-the-shelf single-chip DLP projectors to produce stereoscopic imagery at 120 Hz (see Figure 4a). The first images our fast light engine apparatus (FLEA) projector produced were stunning. There was a complete lack of bleed between the left- and right-eye images—now a well-known advantage of three-chip DLP light engines. This was a critical result. If two or more users were to use the design, we had to eliminate bleed. A multiperson system accumulates more bleed as users increase in number, eventually rendering the system unusable (as opposed to a standard single-user stereoscopic system that only bleeds from one eye to the other).

With this core capability in hand, we developed two approaches for a multiperson display. First, we simply designed the FLEA housing to permit easy stacking and alignment of multiple FLEA projectors, one for each user (see Figure 4b). Optical filtering techniques that we described earlier could then separate two users.³

Alternatively, we integrated FLEA's modified single-chip displays with the optical engine from a three-chip

DLP projector. This approach lets three distinct stereoscopic images overlap on a single screen with fixed alignment and a shared single lens, avoiding alignment issues during system setup. This way, we can achieve a three-person collaborative system running at 120 Hz per eye in a single projector.

Given its small size, FLEA projector variants enable new forms for immersive displays. Large VR systems such as workbenches or small wraparound systems can fit into small conference rooms and individual offices. At the other end of the design spectrum, the small form lets us tile multiple stereoscopic projectors in a small space to create

large tiled images with little distance between the projectors and the screen. Shown in Figure 5, our proposed gazillion-pixel array of stereo projectors (GASP) uses 96 projectors to provide 75 to 125 million pixels in a space smaller than that possible with traditional projectors.

The FLEA's design mandates a reduction in the final image's relative color depth to provide two images. To produce intensity gradation, DLP-based projection systems modulate each pixel's duty cycle during the frame time to achieve the desired pixel brightness. By effectively doubling the projector's frame rate, half as much time is available per frame to modulate the color intensity. This important limitation not only cuts the color depth but also prevents further segmentation of the projected images in time. The projectors are hard coded to require at least 120th of a second to achieve each pixel's color intensity, so we cannot simply stack individual FLEA projectors and then switch them in time to add yet more users.

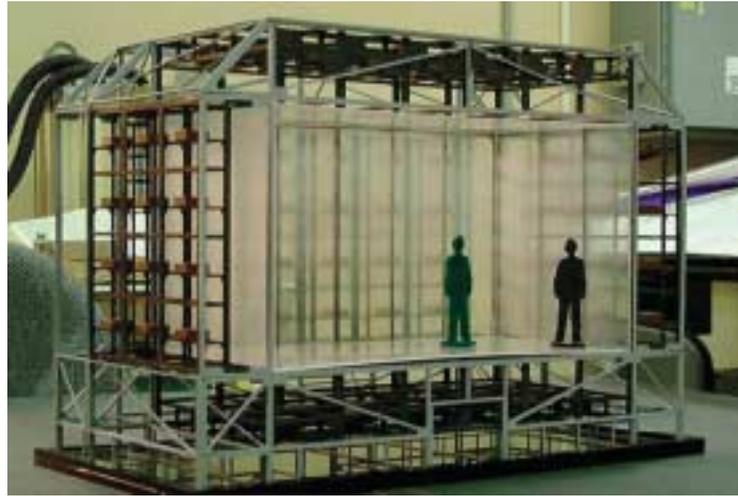
Mule projector

To overcome this limitation, we are working on a single-chip projector that will be fully programmable for a number of applications. Called the multiuser light engine (MULE), this system will take in a digital stream and display images at 2- to 8-kHz frame rates. Although there are associated color-depth costs, this approach will create a single projector that can display multiple pairs of stereoscopic images at flicker-free rates. At the extreme, if we were to display text and simple graphics with single-bit monochrome pixels, perhaps as many as 40 stereo-pairs could be produced, letting the system accommodate more people than can physically fit around the display.

When completed, the MULE will be available to other research groups for experimentation. For example, monochrome refresh rates greater than 1 kHz could be achieved for testing if faster frame rates will improve the perceived sense of presence.

Conclusions

In coming years, VR systems will be deployed to meet the requirements of a wide range of real-world applications. For that to happen, they must present accurate



5 Gazillion-pixel array of stereo projectors showing small footprint and overall form

representations of the models and worlds being visualized, especially as these systems support group collaboration. By employing the technology presented here, we can begin creating systems that let small groups of people fluently communicate spatial concepts as easily as if they could simply draw them in air—because now they can. ■

Acknowledgments

We especially thank Ralph Wachter at the Office of Naval Research and Rich Kaste at the Army Research Lab and the SBIR program for pushing us in right directions. We also thank Oliver Reidel, Pat Hanrahan, Terry Schmidt, Scott Fisher, Marc Levoy, Andreas Simon, and Nanci Anderson. Support has come from the US Office of Naval Research contract number N00014-99-C-0122 and Stanford's MediaX program.

References

1. J. Durbin et al., "Battlefield Visualization on the Responsive Workbench," *Proc. IEEE Visualization Conf.*, IEEE CS Press, 1998, pp. 463-466.
2. M. Agrawala et al., "The Two-User Responsive Workbench: Support for Collaboration through Individual Views of a Shared Space," *Proc. Siggraph Conf.*, ACM Press, 1997, pp. 327-332.
3. I. McDowall et al., "Single and Multiple Viewer Stereo with DLP Projectors," *Proc. SPIE Eng. Reality VR Conf.*, SPIE Press, 2001, pp. 418-425.

See www.fakespacelabs.com/papers.htm for additional references.

Readers may contact the department editors by email at rosenblu@ait.nrl.navy.mil or michael_macedonia@stricom.army.mil.

Readers may contact Mark Bolas at Stanford Univ., Dept. of Mechanical Eng., Stanford, CA; bolas@well.com and Ian McDowall, Fakespace Labs, Mountain View, CA; ian@well.com.