

Adrian Ilie
Kok-Lim Low
Greg Welch
Anselmo Lastra
Henry Fuchs

Department of Computer Science
University of North Carolina at
Chapel Hill

Bruce Cairns

Department of Surgery
University of North Carolina
Chapel Hill

Combining Head-Mounted and Projector-Based Displays for Surgical Training

Abstract

We introduce and present preliminary results for a hybrid display system combining head-mounted and projector-based displays. Our work is motivated by a surgical training application where it is necessary to simultaneously provide both a high-fidelity view of a central close-up task (the surgery) and visual awareness of objects and events in the surrounding environment. In this article, we motivate the use of a hybrid display system, discuss previous work, describe a prototype along with methods for geometric calibration, and present results from a controlled human subject experiment.

This article is an invited resubmission of work presented at IEEE Virtual Reality 2003. The article has been updated and expanded to include (among other things) additional related work and more details about the calibration process.

I Introduction

Today, the pace of surgical innovations has increased dramatically, as have societal demands for safe and effective practices. The mechanisms for training and retraining suffer from inflexible timing, extended time commitments and limited content. Video instruction has long been available to help surgeons learn new procedures, but it is generally viewed as marginally effective at best for a number of reasons, including the fixed point of view that is integral to the narration, the lack of depth perception and interactivity, and the insufficient information that is presented (Tachakra, Jaye, Bak, Hayes, & Sivakumar, 2000).

Apprenticeship (actually being there) can be very effective, but there are some difficult surgical procedures that are so rare and have such high mortality rates that there is no time to assemble the students before the procedure and no time to explain what is going on during the procedure. One example of such a procedure is the treatment of blunt liver traumas resulting from automobile accidents. Such traumas occur rarely (at the UNC-Chapel Hill Hospital there are about five such cases per year), have to be treated immediately (they are life-threatening), and are delicate and complex. As a result, it is typically impractical for trainees to participate in a meaningful way.




Figure 1. Henry Fuchs using a prototype of our hybrid display.

1.1 Immersive Electronic Books

A training paradigm that uses immersive Virtual Reality could be a more effective approach since it would allow trainees to witness and explore a past surgical procedure as if they were there. We are currently pursuing such a paradigm together with our medical collaborators at the UNC-Chapel Hill School of Medicine (Bruce Cairns, M.D., and Anthony Meyer, M.D.), and our computer graphics collaborators at Brown University (Andries van Dam, Sascha Becker, Loring Holden, and Jesse Funaro). Our approach is to record the procedure and to reconstruct the original time-varying events in an immersive 3D Virtual Environment that depicts the real scene, including the patient and expert physicians, augmented with relevant medical instructions and information such as vocal narration, 3D annotations, and illustrations, and to display this in a hybrid display as shown in Figure 1. As part of his Ph.D. thesis, Yang (2003) constructed a research prototype (see Figure 2) and developed new methods that allow 3D data capture and reconstruction. Our collaborators at Brown University are currently working on various methods for annotating the captured data and creating a complete Immersive Electronic Book (van Dam, Fuchs, Becker, Holden, Ilie, et al., 2002).




Figure 2. Dr. Bruce Cairns using the 3D data capture prototype described in (Yang, 2003).

1.2 Display Requirements

Besides the recording and the reconstruction, providing an effective way to display a 3D Virtual Environment to the user is also a major challenge. We envision a display system that can be used both by surgeons (to review a procedure during clinical decision-making, or to annotate it for training purposes), and by trainees (to observe and learn a procedure).

The visual needs of the users are twofold. The first need is for a high-quality stereo view of the objects and events that the trainees are directly viewing, such as the main surgical procedure as shown in Figure 3(a). High-quality and high-resolution views are needed to discern the great intricacy of the surgery, and stereo vision is needed for better spatial understanding. The second need is for a peripheral view of the surroundings. This is needed to maintain visual awareness of the surrounding events as shown in Figure 3(b). Our medical collaborators feel that visual awareness of the entire body of the patient and the surroundings is a critical component of surgical training. In particular, with trauma surgery

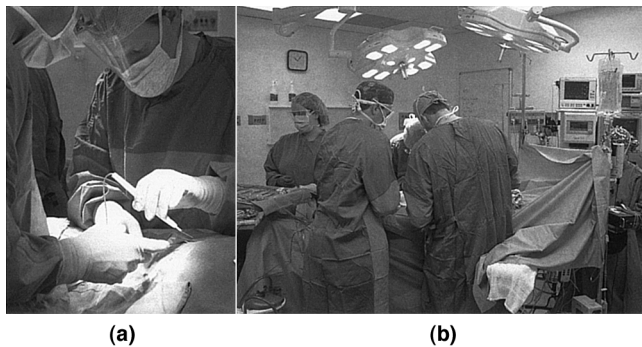


Figure 3. *Two different views of a surgical operation.*

there is typically a great deal of independent but relevant activity in the operating room as each member of the surgical team attends to the patient, monitoring machines change state, and so on. In virtual environments, this type of visual awareness needs to be provided for by a display that covers the user's peripheral view. However, because the effective resolution in the periphery of the human vision system is less dense than that of the foveal region (Davson, 1990), this peripheral display need not be as high-resolution and high-quality.

Traditionally, head-mounted displays (HMDs, also called head-worn displays) have been used to provide high-quality stereo visualization of 3D Virtual Environments. However, most HMDs offer limited fields of view (FOV), typically within the range of 40° to 60° horizontally and 30° to 45° vertically. Very wide FOV HMDs have been prototyped, but they are rare, expensive and cumbersome to wear (Kaiser Electro-Optics Inc., 2000). We are aware of no readily available HMD that can fully cover the human FOV of approximately 200° horizontally and 135° vertically (Werner, 1991). Although HMDs are good at providing high-quality stereo views, their typically narrow FOVs make them less than ideal for providing peripheral views. Common alternatives to HMDs are projector-based displays—either immersive displays such as the CAVE (Cruz-Neira, Sandin, & DeFanti, 1993), or planar displays such as the Responsive Workbench (Kruger, Bohn, Frohlich, Schuth, Strauss, et al., 1995). Most projector-based displays are capable of providing very wide FOV

visualization, and some of them, like the CAVE, are even capable of fully covering the human FOV. However, because the display surfaces—which are relatively large—are typically covered by only one projector, and because the user may move close to the surfaces, the image quality and resolution of such projector-based systems may be insufficient for applications that require fine detail.

1.3 Our Approach (Hybrid Display)

We decided that neither HMDs nor projector-based displays alone could meet the visual needs of a surgical trainee. However, we noticed potential advantages that could be exploited in each display system, especially the high-resolution stereo view in HMDs and the wide FOV in projector-based displays. These complementary characteristics led us to a hybrid approach that enjoys both a high-resolution stereo view through the HMD and a lower-resolution peripheral view from projectors. We believe this hybrid approach can effectively satisfy the visual needs of a surgical trainee.

As a proof of concept and a means of investigation, we implemented a simple prototype of the hybrid display system (see Figure 1). To quantify its effectiveness in combining the useful attributes of both the HMD and the projector-based display, we also conducted a human subject experiment. The preliminary results obtained from the experiment support our belief that the hybrid display is suitable for applications like 3D immersive surgical training, which involve visual tasks that require both up-close detail and peripheral awareness.

1.4 Contributions

The main contribution of this work is conceiving, articulating, and prototyping a hybrid display system that leverages the complementary merits of head-mounted and projector-based displays. A second contribution is our controlled human subject experiment, which offered some surprising (to us) results and valuable support to the idea. Finally, we also introduce a new approach that uses peripheral vision to calibrate opaque HMDs.

1.5 Paper Organization

In the next section, we look at some previous work related to our hybrid display. In section 3, we offer a detailed description of the hybrid display and our simple prototype, and include a discussion of the advantages, drawbacks, and implementation issues of the hybrid display. In section 4 we outline the calibration process, including our new opaque HMD calibration method. In section 5, we describe our human subject experiment and present and analyze its results. Finally, in section 6, we share some closing thoughts.

2 Related Work

Previous research related to our hybrid display prototype system lies in three main areas, corresponding to the two components of the prototype (projector-based displays and HMDs) and to other hybrid displays. The next subsections summarize relevant contributions in the three areas.

2.1 Projector-Based Displays

The CAVE (Cruz-Neira, Sandin, & DeFanti, 1993) is probably the most commonly cited example of a general-purpose projector-based display system for immersive virtual reality and scientific visualization. A CAVE usually consists of a cube (several meters on a side) made of rear-projection screens for the walls and the ceiling, and a down-projection screen for the floor. It is capable of providing stereo imagery to the user, typically using a time-multiplexing approach that alternates images for the left and right eyes.

Another popular and successful paradigm is that of the Responsive Workbench (Kruger, Bohn, Frohlich, Schuth, Strauss, et al., 1995). It consists of a back-projected planar display on which one or two time-multiplexed users can visualize stereo images. It provides an effective means of interacting with virtual environments for some applications and is easier to deploy than a CAVE. Like the CAVE, it has been widely used in research laboratories around the world.

Other approaches for projector-based displays have also emerged, such as the work described by Raskar, Brown, Yang, Chen, Welch, et al. (1999), that aims at reconfigurability, color uniformity, and seamlessness. In this work, non-planar display surfaces are casually arranged; then cameras are used to simultaneously calibrate the projectors and display surfaces to allow correct display of view-dependent (head-tracked) imagery at run time.

The projector-based component of our hybrid display builds on our previous work on life-sized projector-based dioramas (Low, Welch, Lastra & Fuchs, 2001). That work is a more general effort aimed at creating life-sized immersive virtual environments of real or imagined places. The basic idea is to size and arrange physical projection surfaces to roughly approximate the real geometry in the virtual environment, providing both visual and spatial realism.

There have been attempts to incorporate high-resolution insets into projector-based displays. The most notable are the *Escritoire* (Ashdown & Robinson, 2003), a front-projected desk surface that uses two overlapping projectors to create a foveal display that contains a high-resolution area, and *Focus Plus Context* screens (Baudisch, Good, & Stewart, 2001), which use an embedded LCD for the same purpose. They are extensions of the desktop metaphor to a large display, and users need to move objects to the high-resolution area if they want to visualize their details. Other approaches, with applications closer to Virtual Reality, are represented by the *Link ESPRIT* flight simulator (Dugdale, Fortin, & Turner, n.d.) and the *Sogitec APOGÉE* simulator family (Hurault, 1993), which use the pilot's head and/or eye movements to guide a high-resolution projected area-of-interest on a larger, low-resolution projected display.

2.2 Head-Mounted Displays

HMDs have been used as a means to display immersive Virtual Environments for the last three decades (Sutherland, 1968). Research related to HMDs has generally concentrated on improving the resolution and the optics, and minimizing negative effects such

as cybersickness (nausea). The Kaiser Full Immersion Head-Mounted Display (FIHMD) (Kaiser Electro-Optics, 2000) achieves wide FOV by tiling multiple screens to form the display. The FIHMD contains 12 LCD placed in a 3×2 arrangement for each eye. It provides a total FOV of approximately 176° horizontally by 47° vertically. Unfortunately, as far as we know, this system was a one-time prototype and is unavailable for purchase.

Research more closely related to our work includes attempts to incorporate peripheral vision into HMDs. Whitestone and Robinette (1997) concluded that this essentially means increasing the FOV of HMDs, which translates into an increase in the size and resolution of the display screen, an increase in the diameter of the lens, or a decrease in the focal length. In 1979, Eric Howlett designed the Large Expanse Enhanced Perspective (*LEEP*) Optics (LEEP Systems, 1996), which used spherical lenses and square screens, 69 mm on a side, to build an HMD with a 90° FOV. Images were optically predistorted using special lenses attached to cameras that imaged regular CRT displays, and then the distortion was reversed by the optics of the HMD. Slater and Usoh (1993) simulated a wider FOV using distortion of the displayed images in a way that compressed more of the scene into the outer regions of the display area. The foveal region of the display was rendered normally, while, in the periphery, a larger-than-normal portion of the FOV was imaged by adding a shear to the projection. A similar approach is to render high-resolution images in a small area-of-interest and low-resolution images elsewhere. Toto and Kundel (1987) and Yoshida, Rolland, and Reif (1995) constructed area-of-interest HMDs with eye-tracking systems that render the high-resolution inset to where the eye is looking at a given time. Watson, Walker, Hodges, and Reddy (1997) argued that there is no significant decrease in performance if the inset is fixed, since typical eye movements do not go farther than 15° from the fovea.

In his Ph.D. dissertation, Arthur (2000) provides a much more detailed overview of numerous attempts to incorporate peripheral vision into HMDs.

2.3 Hybrid Displays

In 1991, Feiner and Shamash introduced the idea of hybrid displays as a way to move beyond the (then) conventional one-user/one-display metaphor. They demonstrated a prototype hybrid display that used a see-through HMD to provide low-resolution context via visually overlaid imagery and a fixed flat-panel display (viewed through the HMD) to provide high-resolution information. Others have since added low-fidelity displays to a conventional HMD to provide peripheral information with apparent success. For example, the Ambient HMD project at Monterey Technologies has shown that adding peripheral LCD displays improves performance in a helicopter simulator (Monterey Technologies, 1999). At CeBIT 2002, researchers from the Fraunhofer Heinrich-Hertz Institute demonstrated a prototype visualization workstation that combines a central autostereoscopic 3D display with surrounding 2D displays (Pastoor & Liu, 2002). The system is designed to allow a seated user to visualize 3D imagery in a (fixed) central region, and conventional 2D computer “desktop” imagery to the left and right. The system also includes gesture-recognition technology that allows a user to drag objects back and forth between the 2D and 3D display areas.

3 A Hybrid Head-Mounted and Projector-Based Display

Our hybrid display combines an HMD and a projector-based display. The idea is to address the visual needs of a user of the 3D immersive surgical training system by leveraging the complementary merits and shortcomings of HMDs and projector-based displays.

We have implemented a prototype as a proof of the concept. In this prototype, we use a Kaiser ProView 30 stereo HMD (Kaiser Electro-Optics, 2000), which does not have baffling material around the displays that would block the wearer’s peripheral view of the real physical world (see Figure 4). This allows users of the prototype to also see peripheral imagery projected on display surfaces around them.

Ideally, to completely cover the peripheral vision of a



Figure 4. A user wearing the HMD component of our prototype hybrid display with an attached HiBall tracker sensor (Welch et al., 2001).




Figure 5. The projector-based component of our prototype hybrid display.

user, the projector-based display should surround the user, even above and below as in a CAVE. Instead, for our prototype and user study, we have built only a partial surround-display as shown in Figure 5. The display surfaces were constructed using vertically stacked, white Styrofoam blocks and a horizontal piece of white foam-reinforced cardboard. When in use, four projectors

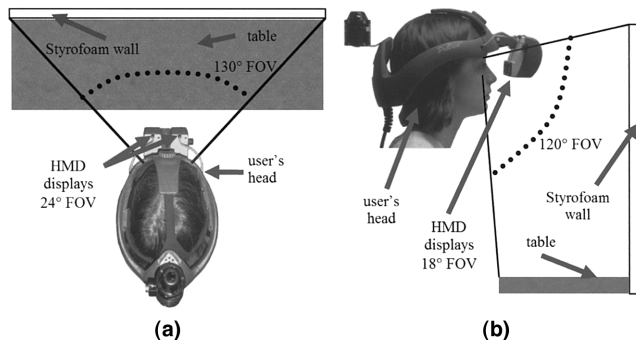


Figure 6. The layout of the hybrid display components to cover the user's FOV.

project view-dependent imagery on top of the cardboard and on the vertical Styrofoam wall. In the spirit of life-sized projector-based dioramas (Low, Welch, Lastra, & Fuchs, 2001), the horizontal cardboard surface corresponds to a surgical table in the virtual operating room. Besides providing the user with a stronger sense of spatial connection between the real and virtual environments, the main purpose of the horizontal cardboard surface is to increase the user's vertical FOV when the user's head is directly above the table.

When a user is standing at the center of the display (approximately 65 cm from the Styrofoam wall, which corresponds to a location where a user would usually stand in the virtual environment), the projector-based display has a FOV of approximately 130° horizontally and 120° vertically (see Figure 6, not drawn to scale). In comparison to the FOV of the Kaiser HMD, which is only 24° horizontally and 18° vertically, we believed the projector-based display was wide enough to offer improvement over the HMD alone.

For the first prototype, we created a simplified synthetic model of an operating room with a virtual patient lying on the surgical table and three virtual surgeons/nurses moving around near the table (see Figure 1). Periodically, one of the virtual surgeons/nurses would extend its arm toward the center of the table, mimicking a real nurse handing a surgical tool to a surgeon during a real surgical procedure. When we tested the prototype ourselves, we would attempt to act in the way a surgical trainee would when using an actual training system. We

tried to concentrate on the virtual patient's abdominal area and maintain awareness of the movements of the virtual surgeons/patients around us. For comparison, each of us had opportunities to try two different setups, one with the use of both the HMD and projectors (the hybrid), and another with the use of the HMD alone. For the HMD-only setup, the projectors were switched off and the laboratory was kept dark.

After having experienced the two setups, every one of us felt that the hybrid display setup made it much easier to maintain awareness of the movements of the virtual surgeons/nurses while concentrating on the patient's abdominal area. In contrast, with the HMD-only setup, it was almost impossible to know the positions of the virtual surgeons/nurses without frequently turning our heads to scan the surroundings. Based on these informal preliminary findings, we decided to undertake a more objective, formal, human subject experiment. We describe this experiment in section 5.

Besides serving as a means to investigate and demonstrate the effectiveness of the hybrid display approach, another objective of the prototype was to identify or confirm potential limitations and implementation issues. These limitations and implementation issues, as well as the advantages of the approach, are discussed in the following subsections.

3.1 Advantages

Through the HMD, a hybrid display can achieve high resolution that is arguably hard to match by any immersive projector-based display. Thanks to the HMD, the user of the hybrid display can selectively increase the image details of any part of the virtual environment by moving closer to it. This directly corresponds to what a person would naturally do in the real world. Projector-based tiled displays (Li, Chen, Chen, Clark, Cook, et al., 2000; Raij, Gill, Majumder, Towles, & Fuchs 2003) have long been used to display high-resolution images; however, their use for our application would be relatively inefficient, as users do not need high resolution everywhere around them at the same time.

As mentioned before, the hybrid display includes a traditional immersive projector-based surround display, allowing it to cover the entire visual field of the user. However, because the users only view the projected imagery in their periphery, that projected imagery can be relatively low-resolution. This means fewer projectors are needed overall. Moreover, since the responsibility to provide stereo vision is already taken over by the HMD, there is no need for stereo projection. This offers further savings in terms of the number of projectors and required rendering resources.

With a hybrid display, there is no need to use a wide-FOV HMD, which would typically be very expensive and cumbersome. Instead a cheaper, smaller, and less cumbersome HMD can be used. A narrower FOV in the HMD also allows some rendering acceleration techniques, such as view frustum culling, to be more effective.

In summary, for our application, we believe that a hybrid display can be a more versatile, cost-effective, efficient, and functionally effective approach than either an HMD or a projector-based system alone.

3.2 Drawbacks, Limitations, and Issues

One drawback to our hybrid approach is that users still have to wear an HMD on their heads, which sometimes can be cumbersome. However, because the HMD does not need to be wide-FOV, and does not have baffling material around the display, it can be made smaller and lighter more easily. We also note that even users of immersive, projector-based display systems typically have to wear something on their heads—stereo glasses and tracking devices.

An artifact that we have noticed in practice is the dark border (opaque frame) around each of the HMD's displays. As with any Augmented Reality application, such borders break the continuity between the central (HMD) and peripheral (projected) views. Our experience is that for commercial systems available today, these borders tend to be relatively thick. For our HMD in particular, we estimate that each vertical segment of the border is approximately 7° wide, and each horizontal segment is approximately 12° tall. We did not per-

form any experiments to investigate how these borders affect a user's performance or sense of presence.

There is an overlap region between the left and right eyes that forms the human stereo visual field (Davson, 1990). This stereo visual field spans almost the entire vertical FOV of the eyes. In our hybrid display prototype, much of the user's stereo visual field is outside the HMD, and is served by non-stereo projected imagery from the projectors. We consider this an important issue for further investigation.

A further limitation of a hybrid system using an opaque HMD is that users cannot see their own hands when they are manipulating virtual objects in front of the HMD. To address this problem, some researchers have developed methods to incorporate real objects, such as the user's hands, into the virtual environments. One such method is described by Lok (2001). Lok also presents a comprehensive survey of other work in the area. Fortunately, the users of our hybrid display system can still see their bodies and legs in the peripheral view. We believe this helps significantly when the users are navigating in the virtual environment. As part of future work, we plan to investigate the use of a see-through HMD in the hybrid display. To address occlusion conflicts between the user's hands (or other real objects) and the virtual objects, we would need either real-time hand tracking or dynamic 3D hand models such as the ones obtained using Lok's approach.

There are also other issues that are inherent in the use of projector-based displays, such as intensity blending at overlapping regions, color uniformity across different projectors and back-projection versus front-projection. We refer the reader to Low, Welch, Lastra, and Fuchs (2001), and Raskar, Brown, Yang, Chen, Welch et al. (1999) for further details.

Finally, because the HMD and projected images originate at different distances from the eyes, users might experience accommodation or focus difficulties if they attempt to rotate their eyes back and forth between HMD and projected imagery. While we believe the users of our system would primarily use the projected imagery for peripheral awareness only, we consider this an important area warranting further investigation.




Figure 7. A user's view of the hybrid display (highlight added).

4 Construction, Calibration, and System Configuration

To create a hybrid HMD-Projector display system, we first establish a global coordinate frame using our tracking system. Afterward, we plan and arrange the projectors and physical display surfaces as in Low, Welch, Lastra, and Fuchs (2001), and calibrate the projector(s) and the HMD. These steps are described in the following subsections. The resulting view from a user's perspective is shown in Figure 7 (photographed with a fish-eye-lens camera placed inside the HMD). The hardware and software configuration of our prototype is described in section 4.5.

4.1 Global Coordinate Frame

To ensure that the combined imagery presented to the user is contiguous, we establish and use a single global coordinate system (virtual world coordinates) throughout the calibration process. To establish the global coordinate system, we compute a 3D rigid-body transformation between the physical tracker coordinates and the virtual world coordinates. We first choose at least four points in the real world and measure their tracker coordinates, and then put them in correspondence with appropriately-chosen points in the virtual world. We then solve for the transformation using a least-squares method described in Arun, Huang, and Blostein (1987). This transformation and its inverse are used later in the calibration process to transform between tracker and virtual world coordinates.




Figure 8. A screenshot of the design program used for arranging the projection surfaces.

4.2 Projector and Display Surface Arrangement

To aid with the construction of projector-based displays, we developed a program to allow us to manually design an arrangement of projectors and surfaces that we use afterward as a blueprint for the physical setup. Figure 8 shows a screenshot of the program, displaying the design we used for the human subject experiment detailed in section 5 (c.f., Figure 1). The projection surfaces are outlined in white. The steps of the process are as follows: (1) we use our program to design the projection surfaces; (2) using this design, we roughly arrange the physical surfaces (Styrofoam blocks); (3) we then position and orient the projectors, and adjust their focusing distances to get the best focus on the display surfaces.

4.3 Projector Calibration

The next step is to calibrate the projectors to find their geometric relationships with the virtual world coordinate frame and the display surfaces. To calibrate a projector, we need to find a set of pair correspondences between the coordinates of 3D points in space and the coordinates of the corresponding pixels on the projector's image plane.




Figure 9. Fine-tuning the physical setup by projecting wireframe outlines.

With six or more pair correspondences (no four coplanar 3D points), we can solve for the projection parameters using linear least-squares methods or nonlinear optimization approaches (Faugeras, 1993). To find a pair correspondence, we display a known 2D point using the projector. This 2D point emerges from the projector as a ray. Then, we move a small, rigid, temporary display surface to intersect the ray and use a tracker to measure the 3D position of the intersection. The 3D point is then transformed from the tracker coordinate system into the virtual world coordinate system using the previously computed transformation. This set of pair correspondences is used to solve for the projector's projection parameters with respect to the virtual world coordinate system. The above manual approach to finding pair correspondences can be tedious and error-prone at times. In the future, we might implement a procedure for automatic projector calibration using cameras such as the one by Raskar, Brown, Yang, Chen, Welch, et al. (1999), possibly as part of an automated projector placement algorithm.

When the projectors are calibrated, we project wireframe outlines of the designed display surfaces onto the physical surfaces to help us fine-tune and correct the physical setup by moving the surfaces slightly to match the projected outlines (see Figure 9; c.f., Figure 8).




Figure 10. *Eyepoint calibration.*

4.4 HMD Calibration

Geometric calibration of the HMD is necessary to properly align the images seen in the HMD with the projected imagery from the projectors. While researchers have calibrated optical see-through HMDs (Tuceryan & Navab, 2000), we are aware of no published method to calibrate an opaque HMD. Here, we present a relatively simple and novel approach that makes use of the user’s peripheral vision. We divide the calibration of each HMD’s display into two stages: *eyepoint calibration* and *view frustum calibration*.

4.4.1 Eyepoint Calibration. In the first stage, we find the position of the eyes in the *tracker sensor* coordinate frame. (The tracker sensor is the part of the tracker that is rigidly attached to the user’s head, so that the *sensor coordinate frame* becomes fixed with respect to the user’s head.) The position of each eye can be found by finding at least two lines in the tracker sensor coordinate frame that intersect at the location of the eye.

We begin by physically marking two points in space that are reasonably far apart but within the tracking range of the tracker. Let these two points be \mathbf{P} and \mathbf{Q} (see Figure 10, not drawn to scale). One of the two points, say \mathbf{P} , should be just below eye level, and there should be a clear line of sight from \mathbf{P} to \mathbf{Q} . Next, we measure the positions of \mathbf{P} and \mathbf{Q} with respect to the *tracker coordinate frame* (the tracker coordinate frame is fixed with respect to the laboratory). Let \mathbf{P}_T and \mathbf{Q}_T be the positions of \mathbf{P} and \mathbf{Q} with respect to the tracker coordinate frame, expressed as three-element column vectors containing the coordinates of points \mathbf{P} and \mathbf{Q} .

We then attach a tracker sensor firmly to the user’s head using a rigid mount placed on the back of the HMD. The following steps are identical for the two eyes, so we describe them for the right eye only. The user positions his or her right eye at a convenient distance to point \mathbf{P} and tries to line up points \mathbf{P} and \mathbf{Q} , as shown in Figure 10. When the user declares that \mathbf{P} and \mathbf{Q} are aligned, the pose of the sensor with respect to the tracker coordinate frame is recorded. At this point, \mathbf{P} and \mathbf{Q} form a line that passes through the right eye. Since we already know \mathbf{P}_T and \mathbf{Q}_T (the positions of \mathbf{P} and \mathbf{Q} with respect to the tracker coordinate frame), and also the pose of the tracker sensor at that moment, we can now express the positions of \mathbf{P} and \mathbf{Q} in the sensor coordinate frame as \mathbf{P}_1 and \mathbf{Q}_1 , respectively. More specifically, we use the sensor’s pose \mathbf{S} (its position and rotation in the tracker coordinate frame, expressed as a 4×4 homogeneous matrix) to transform \mathbf{P}_T and \mathbf{Q}_T in order to get \mathbf{P}_1 and \mathbf{Q}_1 as follows. First, we express all the vectors in homogeneous coordinates, and then compute $(\mathbf{P}_1^T \ 1)^T = \mathbf{S}^{-1} * (\mathbf{P}_T^T \ 1)^T$ and $(\mathbf{Q}_1^T \ 1)^T = \mathbf{S}^{-1} * (\mathbf{Q}_T^T \ 1)^T$. Finally, we extract \mathbf{P}_1 and \mathbf{Q}_1 as the first three scalar components of the homogeneous coordinates column vectors $(\mathbf{P}_1^T \ 1)^T$ and $(\mathbf{Q}_1^T \ 1)^T$. When the user moves his or her right eye away from the line formed by \mathbf{P} and \mathbf{Q} , the points with coordinates \mathbf{P}_1 and \mathbf{Q}_1 remain the same in the sensor coordinate frame and form a line that still passes through the right eye. In other words, the line passing through the points with coordinates \mathbf{P}_1 and \mathbf{Q}_1 has become “rigidly attached” to the sensor coordinate frame and will always pass through the right eye no matter how the user moves. In order to find the position of the right eye in the sensor coordinate frame, we need to capture at least one additional line so that its intersection with the first can be used to determine the viewpoint’s position. The user needs to line up \mathbf{P} and \mathbf{Q} again, but with a different head orientation from the first time. This requires the user to roll his or her right eyeball in another direction. The new pose of the tracker sensor at that moment is recorded and used to transform \mathbf{P}_T and \mathbf{Q}_T into the sensor coordinate frame as \mathbf{P}_2 and \mathbf{Q}_2 , respectively. The points with coordinates \mathbf{P}_2 and \mathbf{Q}_2 form a second line that passes through the right eye. We can repeat the

above procedure to capture additional lines. In practice, because of measurement errors, these captured lines may not pass exactly through the eye, and may not even intersect each other at all. Additional lines help improve the accuracy of the estimated eye position.

We use n captured lines to determine the position of the eye in the sensor coordinate frame as the 3D point that has the shortest distance to all the lines. First, let l_i be the i^{th} line passing through the points with coordinates $\mathbf{P}_i = (P_{ix}, P_{iy}, P_{iz})^T$ and $\mathbf{Q}_i = (Q_{ix}, Q_{iy}, Q_{iz})^T$, for $1 \leq i \leq n$. We further let $\mathbf{U}_i = (U_{ix}, U_{iy}, U_{iz})^T = \mathbf{P}_i - \mathbf{Q}_i$. Let $\mathbf{M} = (M_x, M_y, M_z)^T$ be the coordinates of the point that has the shortest total distance to all the n lines of sight. If all the lines intersect *exactly* at this common point, then the following is true for all $1 \leq i \leq n$: $\mathbf{M} + t_i \mathbf{U}_i = \mathbf{Q}_i$, where each t_i is some scalar whose value is yet to be determined.

By combining all instances of this equation for all $1 \leq i \leq n$, we can write them in the form of $\mathbf{Ax} = \mathbf{b}$ as:

$$\begin{bmatrix} 1 & 0 & 0 & U_{1x} & 0 & \cdots & 0 \\ 0 & 1 & 0 & U_{1y} & 0 & \cdots & 0 \\ 0 & 0 & 1 & U_{1z} & 0 & \cdots & 0 \\ 1 & 0 & 0 & 0 & U_{2x} & \cdots & 0 \\ 0 & 1 & 0 & 0 & U_{2y} & \cdots & 0 \\ 0 & 0 & 1 & 0 & U_{2z} & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 0 & 0 & 0 & 0 & \cdots & U_{nx} \\ 0 & 1 & 0 & 0 & 0 & \cdots & U_{ny} \\ 0 & 0 & 1 & 0 & 0 & \cdots & U_{nz} \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z \\ t_1 \\ t_2 \\ \vdots \\ t_n \end{bmatrix} = \begin{bmatrix} Q_{1x} \\ Q_{1y} \\ Q_{1z} \\ Q_{2x} \\ Q_{2y} \\ Q_{2z} \\ \vdots \\ Q_{nx} \\ Q_{ny} \\ Q_{nz} \end{bmatrix}$$

where \mathbf{A} is a $3n \times (n + 3)$ matrix of the known U elements, \mathbf{x} is an $(n + 3)$ -element column vector of the known M and t_i elements, and \mathbf{b} is a $3n$ -element column vector of the known Q elements. In practice, because of errors in the measurements, $\mathbf{Ax} = \mathbf{b}$ is almost always an inconsistent system; that is, \mathbf{b} is not in the *range* of \mathbf{A} . Generally, the columns of \mathbf{A} are independent; therefore \mathbf{A} has a *rank* of $n + 3$. So, the *least squares solution* of $\mathbf{Ax} = \mathbf{b}$ is just $\mathbf{x}' = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{b}$. The symmetric matrix $\mathbf{A}^T \mathbf{A}$ is invertible because it is an $(n + 3) \times (n + 3)$ matrix and it has the same rank as \mathbf{A} . For more information about linear least-squares solutions, see Strang 1988.

The last step of the viewpoint calibration is just to extract M_x , M_y , and M_z from \mathbf{x}' . The column vector $\mathbf{M} = (M_x, M_y, M_z)^T$ is the estimate of the right eye's position with respect to the target coordinate frame. Using MATLAB (MathWorks, Inc., 2001), the least-squares solution \mathbf{x}' can be computed as $\mathbf{x}' = \mathbf{A} \backslash \mathbf{b}$.

In our current implementation we use only two lines. In this case, the eyepoint is at the middle of the shortest segment connecting the two lines. We display the computed shortest distance between the two lines as part of the calibration process, and offer the user the choice of whether to continue or repeat the calibration. We also display the computed interpupillary distance (IPD) after calibrating both eyepoints, and ask the user to repeat the process if it differs significantly from the measured IPD. Because, typically, the virtual objects in our application are closer to the projection surfaces than to the eyepoint, the errors in estimating the eyepoint result in displacement errors of smaller magnitude in the displayed imagery. During the use of our system, we observed that eyepoint estimation errors of less than 1 cm are accurate enough for our purposes. Since the projector-based component of the hybrid display uses the point situated at the middle of the segment connecting the two eyepoints as its rendering viewpoint, the impact of these errors on the perceived continuity of the two components of the imagery presented to the users is minimal.

4.4.2 View Frustum Calibration. In the second stage, we use a pen to mark single horizontal and vertical 3D lines on a wall. We make the lines long enough to extend beyond the FOV of the HMD. We measure the 3D positions of the two physical lines in the tracker coordinate frame, which is fixed with respect to the physical world. We display two horizontal and two vertical 2D lines in the HMD (see Figure 11). These lines intersect at four known screen positions. One at a time, the user uses the peripheral vision in each eye to line up each 2D horizontal (vertical) line in the HMD with the 3D horizontal (vertical) line in the real world. This procedure allows us to determine four planes in the sensor coordinate frame that intersect to form four rays originating from the viewpoint. Taking the coordinates of any 3D point on a ray and pairing it with the corre-




Figure 11. View frustum calibration.

sponding 2D pixel coordinates, we get a pair correspondence. With four pair correspondences, one for each ray, we use a variant of the camera calibration technique presented in Trucco and Verri (1998) and Faugeras (1993) to solve for the view parameters of the HMD display and use these parameters to construct OpenGL viewing matrices at run time. The difference from the method presented by Trucco and Verri and by Faugeras consists of the fact that we already know the position of the viewpoint, so we do not need to solve for its coordinates. In Low and Ilie (2003) we present a detailed description of our method in the context of a different problem: computing a view frustum to maximize an object's image area in order to improve the resolution of shadow and texture maps.

This procedure has worked well for us. Our experience is that a user can typically align the displayed 2D and real 3D lines with what appears to be less than one degree of error, which, for our setup, results in relatively good registration between the HMD and projected imagery. However, the procedure can be time consuming, and difficult to perform for inexperienced users. This is the reason why, for the human subject experiment described in section 5, we precalibrated the system. Subjects were instructed to perform mechanical adjustments to the HMD so that the perceived imagery looked continuous to them.

4.5 System Configuration

The system configuration of our hybrid prototype is shown in Figure 12. Solid arrows represent the transfer of data from input devices, dashed arrows represent




Figure 12. The configuration of the system used for the hybrid display prototype.

the transfer of tracking data to clients, and dotted arrows represent the transfer of images to display devices. Each projector is driven by a separate renderer. To ensure coherence between the two displays of the HMD, we use one renderer with two viewports. In our first implementation, each renderer was a separate graphics pipe on an SGI Reality Monster. Since then, we have moved to a cluster of PCs, and we use a dual-head graphics card for the HMD.

Global system state coherence is enforced using a message-passing library that implements the “shared memory” metaphor. We did not enforce frame coherence between the renderers because it would have dramatically decreased the performance of the system. Mismatches between renderers are not noticeable because the frame rates are quite high (minimum 25 frames/second), and the report interval for tracking data is relatively small. All the renderers connect to a VRPN server (Taylor, Hudson, Seeger, Weber, Jeffrey, et al., 2001) to get the most recent tracking data from our wide-area HiBall optical tracker (Welch, Bishop, Vicci, Brumback, Keller, et al., 2001). The same VRPN server controls the button that we used as a user interface.

5 Experiment and Results

For an objective evaluation of the hybrid display, we conducted a formal human subject experiment involving 25 subjects randomly recruited from the UNC-Chapel Hill student population. The applicants were screened for



Figure 13. The virtual environment used in our experiment.

(among other things) normal vision and prior experience with virtual reality systems (we wanted none).

While we could not test the subjects on visual tasks that could only be performed in an immersive surgical training system, we have designed simpler tasks that we believe are analogous to the actual tasks. In keeping with the visual training requirements described in section 1, our primary criterion for the experimental tasks was that they should force the subject to concentrate on a central up-close task while simultaneously maintaining awareness of peripheral events. The two different displays we considered were (1) our hybrid display (both HMD and projectors), and (2) an HMD alone. For the hybrid display, we used the same physical setup as our prototype, and for the HMD-only setup, we just switched off the projectors and kept our laboratory dark.

We decided not to test a projector-only setup for two reasons. Given the available equipment, we were unable to produce stereo-projected images of sufficient resolution. In addition, we wanted to keep the user study manageable in scope. We also did not compare our hybrid display against general-purpose visualization systems such as the CAVE because we were specifically targeting a surgical training application for which we judged such visualization systems to be less suitable.



Figure 14. A word search puzzle.

5.1 Experiment Description

The purpose of the experiment was to study the effect of two different displays on the performance in tasks that simultaneously required (1) the ability to be visually aware of changes in the surroundings in the virtual environment and (2) the ability to visually and mentally concentrate on a central, static virtual object.

We developed two specific hypotheses about the hybrid display compared to the HMD alone: (H1) Hybrid display users are more visually aware of changes in their virtual surroundings and (H2) Users of the hybrid display can visually and mentally concentrate better on the central static virtual object. To test these hypotheses, we created a virtual environment that consisted of a simple room with brick walls and a wooden table in the middle as shown in Figure 13.

To test hypothesis H1, four times every minute, the system displayed a spinning, cube-like “token” that entered from the far left or right side of the subject and floated toward the center (see upper left in Figure 14 and upper right in Figure 15). When near the center, the token would reverse its course and move back to where it came from, eventually disappearing from the room. The lifetime of each token was 8 seconds, chosen randomly as part of each of the four 15-second intervals of every minute. We used the same precomputed se-



Figure 15: A “Where’s Waldo?” puzzle.

quence of tokens for all the subjects. There were two types of tokens: bonus tokens and penalty tokens. The bonus tokens had a single number (a point value) on all sides, while a penalty token had the letter *X* on all sides. Once a bonus token appeared, the digit shown was decremented every second. The subjects were instructed to press a button as soon as they identified a bonus token, and do nothing if they noticed a penalty token. We did not actually penalize the subjects for pressing the button when seeing a penalty token, but we wanted them to think they would be. The idea was to discourage them from pressing the button without actually looking at the tokens. While we did not display a score, the users had visual feedback (colored text messages) when they pressed the button.

To test hypothesis H2, we made the users simultaneously solve some puzzles that appeared on a small podium-like *virtual screen* on the wooden table shown in Figure 13. We used two different sets of puzzles.

The first set consisted of ten word-search puzzles (one is shown in Figure 14). For each word-search puzzle, the subjects were shown ten words in the right column of the virtual screen and asked to find any eight of them in the grid of letters. When a word was found, the subjects were required to speak the word and describe its orientation (top-down, bottom-up, left-to-right, or right-to-left), so that one of the experimenters could




Figure 16: A subject during the experiment.

cross out the word on the virtual screen by pressing a key on a keyboard. When the subjects found eight words in a puzzle, the experimenter advanced to the next puzzle.

The second set of puzzles (see Figure 15) consisted of 20 colorful and detailed pictures from various “Where’s Waldo?” books (Handford, 1997). We overlaid an annotated grid on the pictures and scaled them so that the fictional cartoon character, Waldo, would occupy approximately $\frac{1}{4}$ of a grid cell. For each puzzle, the subjects were required to look for Waldo in the picture. Once Waldo was found, the subjects were asked to speak the coordinates of Waldo’s location. If the coordinates were correct, the experimenter advanced to the next puzzle; otherwise, the subjects were asked to search again.

In each trial of the experiment, the subjects were given 10 minutes to solve either the set of word-search puzzles or the set of Waldo puzzles. We made the choice in advance, with a uniform random distribution. The subjects were told that their objective was to correctly identify as many bonus tokens and simultaneously solve as many puzzles as they could within 10 minutes. We specifically emphasized that the two tasks were equally important. The subjects were given about one minute to practice before the actual trial began.

During a session, each subject was required to do two

Table 1. Summary of Experiment Data

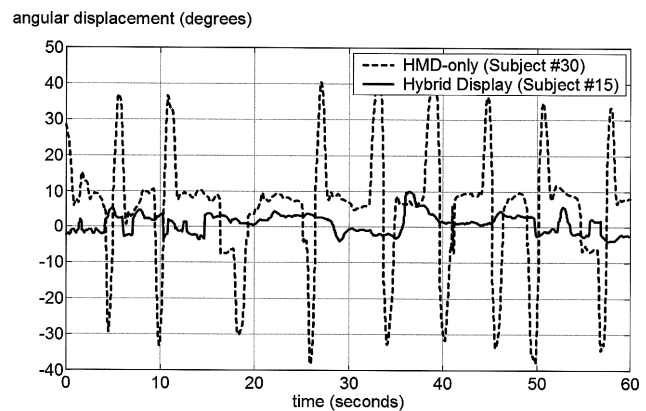
Puzzles	Where's Waldo				Word search			
	1 (HMD only)		2 (hybrid display)		3 (HMD only)		4 (hybrid display)	
Combinations	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Bonus tokens %	54.4	16.3	99.1	2.1	55.2	21.4	97.7	8.1
Puzzles (words)/min	2.0	.5	2.1	.7	4.4	1.4	5.4	2.3
Avg. angular speed (°/s)	23.8	5.6	9.5	3.1	18.4	5.8	5.3	2.3
Head turns/min	9.0	2.2	.5	.9	9.1	3.1	.4	.9

different trials. There was a 10-minute break between the end of the first trial and the start of the second. Each trial tested one of the following four combinations: (1) HMD-only and Waldo puzzles, (2) hybrid display and Waldo puzzles, (3) HMD-only and word-search puzzles, and (4) hybrid display and word-search puzzles. The subjects were randomly assigned combinations 1 & 4, 4 & 1, 2 & 3, or 3 & 2, for the session. Figure 16 shows a subject doing a combination 2 trial.

5.2 Results and Analyses

We conjectured that hypothesis H1 would be supported if the number of bonus tokens identified by the hybrid display users was significantly higher than those identified by the HMD-only users, whereas hypothesis H2 would be supported if the number of puzzles solved and the number of words found by the hybrid display users was significantly higher.

Table 1 shows a summary of some of the data we obtained from the experiment. “Bonus Tokens %” is the percentage of bonus tokens that were successfully identified during the duration of a trial. The large differences between the mean bonus tokens percentage of the hybrid display and the mean bonus tokens percentage of the HMD (highlighted in dark gray) clearly support hypothesis H1. However, if we look at the “Puzzles (Words)/min” row, we do not see any significant differences (highlighted in light gray) between the hybrid display and the HMD. Therefore, hypothesis H2 is not

**Figure 17.** Angular displacement plot of head rotation.

supported by our experimental data. This was quite unexpected because during the experiment, we saw the HMD-only users making large and frequent head turns to scan for tokens. We assumed this would have been very distracting and would have largely reduced the subjects’ ability to solve puzzles. The frequent head turns by the HMD users can be observed in the last row of Table 1. Figure 17 shows an angular-displacement plot that compares the amplitude and frequency of head rotations made by an HMD user and a hybrid display user.

We believe that the failure of the experimental data to support H2 may lie in the inappropriate choice of tasks for our experiment. The ability to solve word-search or Waldo puzzles may vary significantly across the general

Table 2. *Mutual Correlations in Combination 3*

	Correlation coefficient	<i>p</i> -value
Bonus tokens % vs. head turns/min	.85	.0005 < <i>p</i> < .0025
Words/min vs. head turn/min	.56	.025 < <i>p</i> < .05
Bonus tokens % vs. words/min	.63	.02 < <i>p</i> < .015

population. In retrospect, we realize that a simpler central task, such as monitoring some changing symbols for a “magic letter,” would likely have less variance. For a few users, the puzzles proved so involving that they were completely absorbed by the central task and had to be reminded to give equal importance to the peripheral task.

We also performed correlation analyses on the experimental data and discovered some surprises in the results of combination 3 (see Table 2). While we expected the number of bonus tokens identified and the frequency of head turns to be correlated (row 1), we did not expect to see the positive correlations between words found and head turn frequency (row 2), and bonus tokens and words found (row 3). In fact, it appears that words found, bonus tokens, and head turns were mutually correlated. One possible explanation is that subjects who had difficulty with the word-search puzzles may have had less leftover mental capacity to maintain peripheral awareness. This explanation is supported by the fact that the same users who had difficulties with the central task had to be reminded to give equal importance to the peripheral task.

6 Conclusions and Future Work

We began this work with basically only an intuition about the benefits of a hybrid approach, believing that it should be possible to realize the combined benefits of tracked HMD and projector-based displays. At this point, we have a concrete prototype, and some objective assessment of the usability and benefits of the approach. While the analysis of our user study does not

completely support our original hypotheses, we are encouraged enough by the results to invest significant further effort in exploring larger prototypes, issues related to mono–stereo transitions, and variations, such as the use of a see-through HMD.

While we would not claim that our hybrid approach is appropriate for general-purpose visualization, we believe it could extend to other applications. Examples include HMD-based applications where surrounding imagery would provide useful context or awareness, or projector-based applications where high-fidelity, close-up 3D imagery would improve the user’s insight into the data or task. In particular, we believe the approach might be useful for early evaluation and simulation of vehicle interiors in automobiles, tractors, submarines, and ships.

Acknowledgments

We thank Tammy Pike for help in managing experiment participants; Haolong Ma for both the use of his 3D operating room models and his operational help with the experiment; Hsin-Yee Teh for help in the statistical analyses; and Andrew Raj for his insightful suggestions. Gennette Gill was the subject in some of the photographs that appear in this paper. We thank the members of the UNC EVE group, as well as Jim Mahaney, Herman Towles, David Harrison, and Kurtis Keller for their support. We would also like to thank our reviewers, our audience at the IEEE Virtual Reality 2003 conference, and especially Steven Feiner for their feedback and suggestions. This research was supported by NSF ITR Grant no. 0121657: “Electronic Books for the Tele-Immersion Age” and NSF Grant no. 0205925, “Image-based Rendering in Forensic Science, Education, and Historical Presentation.”

References

- Arthur, K. (2000). *Effects of Field of View on Performance with Head-mounted Displays*. Unpublished doctoral dissertation, Department of Computer Science, University of North Carolina at Chapel Hill.
- Arun, K. S., Huang, T. S., & Blostein, S. D. (1987). Least-squares fitting of two 3-D point sets. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(5), 698–700.
- Ashdown, M., & Robinson, P. (2003). The Escritoire: A personal projected display. *Journal of the Winter School of Computer Graphics*, 11(1), 33–40.
- Baudisch, P., Good, N., & Stewart, P. (2001). Focus Plus Context screens: Combining display technology with visualization techniques. *Proceedings of the ACM Symposium on User Interface Software and Technology 2001*, 31–40.
- Cruz-Neira, C., Sandin, D. J., & DeFanti, T. A. (1993). Surround-screen projection-based Virtual Reality: The design and implementation of the CAVE. *Proceedings of ACM SIGGRAPH 1993*, 135–142.
- Davson, H. (1990). *Physiology of the eye* (5th edition). New York: Pergamon.
- Dugdale, J., Fortin, M., & Turner, J. *Current Developments in Visual Display Technology for Fighter Type Simulators*. Available from Link Simulation and Training Web site: <http://www.link.com/pdfs/itsec.pdf>
- Faugeras, O. (1993). *Three-Dimensional Computer Vision*. Cambridge, MA: MIT Press.
- Feiner, S., & Shamash, A. (1991). Hybrid user interfaces: Breeding virtually bigger interfaces for physically smaller computers. *Proceedings of the ACM Symposium on User Interface Software and Technology 1991*, 9–17.
- Handford, M. (1997). *Where's Waldo?* Cambridge, MA: Candlewick.
- Hurault, F. (1993). A head slaved visual system for flight simulators. *Proceedings of the International Training Equipment Conference and Exhibition*, 37–42.
- Kaiser Electro-Optics, Inc. (2000). *HiDef family of head-mounted Displays*. Available at <http://www.keo.com>
- Kruger, W., Bohn, C.-A., Frohlich, B., Schuth, H., Strauss, W., & Wesche, G. (1995). The responsive workbench: A virtual work environment. *IEEE Computer*, 28(7), 42–48.
- LEEP Systems, Inc. *New from LEEP Systems: A Virtual Reality Adventure for Trade Show Visitors*. Archived news page at <http://world.std.com/~leep/oldindex.htm>, June 1996.
- Li, K., Chen, H., Chen, Y., Clark, D. W., Cook, P., Damianakis, S., et al. (2000). Early experiences and challenges in building and using a scalable display wall system. *IEEE Computer Graphics and Applications*, 20(4), 671–680.
- Lok, B. (2001). Online model reconstruction for interactive virtual environments. *Proceedings of the ACM Symposium on Interactive 3D Graphics 2001*, 69–72.
- Low, K.-L., & Ilie, A. (2003). Computing a View Frustum to Maximize an Object's Image Area. *Journal of Graphics Tools*, 8(1), 3–15.
- Low, K.-L., Welch, G., Lastra, A., & Fuchs, H. (2001). Life-sized projector-based dioramas. *Proceedings of the ACM Symposium on Virtual Reality and Software Technology 2001*, 93–101.
- MathWorks, Inc. (2001). MATLAB Version 6.1.0.450 R12.1 [Computer software]. Available at <http://www.mathworks.com>. Released May 18, 2001.
- Monterey Technologies (1999). *Presentation of Aircraft State Information in an Ambient Head-Mounted Display*. Available at <http://www.montereytechnologies.com/hmd.htm>
- Pastoor, S., & Liu, J. (2002). 3-D display and interaction technologies for desktop computing. In B. Javidi & F. Okano (Eds.), *Three-Dimensional Television, Video, and Display Technologies* (pp. 315–356). Berlin: Springer Verlag.
- Raij, A., Gill, G., Majumder, A., Towles, H., & Fuchs, H. (2003). PixelFlex2: A comprehensive, automatic, casually-aligned multi-projector display. *Proceedings of the IEEE International Workshop on Projector-Camera Systems PRO-CAMS-2003*.
- Raskar, R., Brown, M., Yang, R., Chen, W.-C., Welch, G., Towles, H., et al. (1999). Multi-projector displays using camera-based registration. *Proceedings of the IEEE Conference on Visualization 1999*, 161–168.
- Slater, M., & Usoh, M. (1993). Simulating peripheral vision in immersive virtual environments. *Computers & Graphics*, 17(6), 643–653.
- Strang, G. (1988). *Linear Algebra and Its Applications* (3rd edition). San Diego: Saunders Publishing/Harcourt Brace Jovanovich College Publishers.
- Sutherland, I. E. (1968). A head-mounted three dimensional display. *Proceedings of the 1968 Fall Joint Computer Conference, AFIPS Conference*, 33(1), 757–764.
- Tachakra, S., Jaye, P., Bak, J., Hayes, J., & Sivakumar, A. (2000). Supervising trauma life support by telemedicine. *Journal of Telemedicine and Telecare*, 6(1), 7–11.
- Taylor R. M., II, Hudson, T. C., Seeger, A., Weber, H., Jeffrey J., Helser, A. T. (2001). VRPN: A device-independent,

- network-transparent VR peripheral system. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology 2001*, 47–54.
- Toto, L., & Kundel, H. (1987). An eye motion dependent interactive display. In J. K. O'Regan & A. Levy-Schoen (Eds.), *Eye Movements: From Physiology to Cognition*. (pp. 643–649). Amsterdam: Elsevier Science Publishers.
- Trucco, E., & Verri, A. (1998). *Introductory Techniques for 3-D Computer Vision*. Edgewood Cliffs, NJ: Prentice Hall.
- Tuceryan M., & Navab, N. (2000). Single point active alignment method (SPAAM) for optical see-through HMD calibration for AR. *Proceedings of the IEEE and ACM International Symposium on Augmented Reality 2000*, 149–158.
- van Dam, A., Fuchs, H., Becker, S., Holden, L., Ilie, A., Low, K.-L., et al. (2002). Immersive electronic books for teaching surgical procedures. Proceedings of the *Pre-ICAT CREST Symposium on Telecommunication, Teleimmersion, and Telexistence*.
- Watson, B., Walker, N., Hodges, L. F., & Reddy, M. (1997). An evaluation of level of detail degradation in head-mounted display peripheries. *Presence: Teleoperators and Virtual Environments*, 6(6), 630–637.
- Welch, G., Bishop, G., Vicci, L., Brumback, S., Keller, K., & Colucci, D. (2001). The HiBall Tracker: High-performance wide-area tracking for virtual and augmented environments. *Presence: Teleoperators and Virtual Environments 10(1)*, 1–21.
- Werner, E. B. (1991). *Manual of visual fields*. New York: Churchill Livingstone.
- Whitestone, J. J., & Robinette, K. M. (1997). Fitting to maximize performance of HMD systems. In J.E. Melzer & K. Moffitt (Eds.), *Head-Mounted Displays: Designing For the User* (pp. 175–206). New York: McGraw-Hill.
- Yang, R. (2003). *View-Dependent Pixel Coloring—A Physically-Based Approach for 2D View Synthesis*. Unpublished doctoral dissertation, Department of Computer Science, University of North Carolina at Chapel Hill.
- Yoshida, A., Rolland, J. P., & Reif, J. H. (1995). Design and applications of a high-resolution insert head-mounted display. *Proceedings of the IEEE Virtual Reality Annual International Symposium 1995*, 84–93.