Oliver Bimber*
obimber@crcg.edu
Fraunhofer Center for Research in Computer Graphics (CRCG), Inc.
321 S. Main St.
Providence, RI 02903, USA

L. Miguel Encarnação
me@crcg.edu
Fraunhofer Center for Research in Computer Graphics (CRCG), Inc.
321 S. Main St.
Providence, RI 02903, USA

Pedro Branco
pbranco@crcg.edu
Fraunhofer Center for Research in Computer Graphics (CRCG), Inc.
321 S. Main St.
Providence, RI 02903, USA

The Extended Virtual Table: An Optical Extension for Table-Like Projection Systems

Abstract

A prototype of an optical extension for table-like rear-projection systems is described. A large, half-silvered mirror beam splitter is used as the optical combiner to unify a virtual and a real workbench. The virtual workbench has been enabled to display computer graphics beyond its projection boundaries and to combine virtual environments with the adjacent real world. A variety of techniques are described and referred to that allow indirect interaction with virtual objects through the mirror. Furthermore, the optical distortion that is caused by the half-silvered mirror combiner is analyzed, and techniques are presented to compensate for this distortion.

1 Introduction and Motivation

Virtual reality (VR) attempts to provide to the user a sense of spatial presence (visual, auditory, or tactile) inside computer-generated synthetic environments. Opaque, head-mounted displays (HMDs) and surround-screen (spatially immersive) displays such as CAVEs, (Cruz-Neira, Sandin, & DeFanti, 1993) and domed displays (Bennett, 2000) are VR devices that surround the viewer with graphics by filling the user's field of view. To achieve this kind of immersion, however, these devices encapsulate the user from the real world, thus making it difficult or even impossible in many cases to combine them with habitual work environments.[1]

Other, less immersive display technology is more promising to support seamless integration of VR into everyday workplaces. Table-like display devices such as Barco's Virtual Tables or Responsive Workbenches (Krüger & Fröhlich, 1994; Krüger et al., 1995) and wall-like projection systems such as Silicon Graphics' Powerwalls allow the user to simultaneously perceive the surrounding real world while working with a virtual environment.

University of North Carolina's "Office of the Future Vision" (Raskar et al., 1998) is a consequent extension of this concept. Here, in contrast to embedding special display devices into the real work environment, an office is envisioned in which the ceiling lights are replaced by cameras and projectors that

*This work was done when the first author was at the Fraunhofer Institute for Computer Graphics in Rostock, Germany.

1. Common workplaces such as offices.

Presence, *Vol. 10, No. 6, December 2001*, 613–631

continuously scan the office environment and project computer graphics to spatially immersive displays. These displays could, in effect, be almost anything (such as, walls, tables, and cupboards) or anywhere in the office. While the cameras acquire the geometry of the office items (irregular surfaces), the rendering is modified to project graphics onto these surfaces in a way that looks correct and undistorted to an observer. This concept can offer both a high degree of immersion and the integration of VR into the workspace. Due to currently employed display technology, a main limitation of VR is that virtual environments cannot be optically mixed with the real world. If rear-projection systems are employed, real-world objects are always located between the observer and the projection plane, thus occluding the projected graphics and consequently the virtual environment. If front projection is used, physical models can be augmented with graphics by seamlessly projecting directly onto the surface of those objects instead of displaying them in the viewer's visual field (Raskar, Welch, & Chen, 1999; Raskar, Welch, & Fuchs, 1998). However, this spatially augmented reality (SAR) concept is mostly limited to visualization and is not suitable for advanced interaction with virtual and augmented real objects. Moreover, shadows that are cast by the physical objects or by the user and restrictions of the display area (size, shape, and color of the surface) represent fundamental problems in SAR systems.

In general, augmented reality (AR) superimposes computer-generated graphics onto the user's view of the real world. In contrast to VR, AR allows virtual and real objects to coexist within the same space. Opaque HMDs that display a video stream of the real world, which is premixed with graphics—or see-through HMDs (Sutherland, 1965; Bajura, Fuchs, & Ohbuchi, 1992) that make use of optical combiners (essentially half-silvered mirrors)—are currently the two main display devices for AR.[2] Similar to VR, the display technology that is employed for AR introduces a number of drawbacks: for currently available HMDs, display characteristics (resolution, field of view, focal length, depth of field, and so on) and ergonomic factors usually interfere. Although the resolution of both HMD types (opaque and see-through) is generally low (lower than projection-based VR display devices), today's optical, see-through systems additionally lack image brilliance. This is because the brightness of the displayed graphics depends strongly upon the lighting conditions of the surrounding real environment. Although higher-resolution, see-through HMDs do exist, they are often heavy and expensive. However, more ergonomic HMDs lack their optical properties.

Head-mounted projective displays (Parsons & Rolland, 1998; Inami et al., 2000) or projective head-mounted displays (Kijima & Ojika, 1997) are projection-based alternatives that employ head-mounted miniature projectors instead of miniature displays. Such devices tend to combine the advantages of large projection displays with those of HMDs. Similar to SAR, head-mounted projective displays decrease the effect of inconsistency of accommodation and convergence that is related to HMDs. Both, head-mounted projective displays and projective HMDs also address other problems that are related to HMDs: they provide a larger field of view without the application of additional lenses that introduce distorting arbitrations. They also prevent incorrect parallax distortions caused by interpupil distance (IPD) mismatch that occurs if HMDs are worn incorrectly (such as if they slip slightly from their designed position). However, corresponding with HMDs, they seriously suffer from the imbalanced ratio between heavy optics (or projectors) that results in cumbersome and uncomfortable devices or ergonomic devices with a poor image quality.

Although some researchers refer to AR as a variation of VR (such as Azuma (1997)), a strong separation between AR and VR applications does exist, and, in our opinion, is caused mainly by the technologically constrained usage of different display devices.

In this article, we introduce a prototype of a cost-effective and simple-to-realize optical extension for single-sided or multiple-sided (that is, L-shaped), table-like projection systems. A large half-silvered mirror beam splitter is applied to extend both viewing and in-

2. See Rolland, Rich, and Fuchs (1994) for additional discussion on the pros and cons between the video mixing and optical combination.

teraction space beyond the projection boundaries of such devices. The beam splitter allows an extension of exclusively virtual environments and enables these VR display devices to support AR tasks. Consequently, the presented prototype features a combination of VR and AR. Because table-like display devices can easily be integrated into habitual work environments, the extension allows the linkage of a virtual with a real workplace, such as a table-like projection system with a neighboring real workbench.

Compared to current HMDs, the application of a spatial projection displays (such as the prototype described in this article) for AR tasks features an improved ergonomics, a large field of view, a high and scalable resolution, and an easier eye accommodation (Raskar, Welch, & Fuchs, 1998). In contrast to Raskar's SAR concept, however, our optical see-through approach prevents shadow casting and does not restrict the display area to the real environment's surface shape.

2 Outline of the Article

In section 3, a description of prior work is presented. Different table-like projection systems are provided, and related devices that apply mirrors are discussed. In section 4, our current prototype is introduced from different points of view: its physical setup, its general functioning, and its interaction possibilities. Section 5 analyzes different sources of optical distortion that are typical for our mirror beam splitter extension and presents possible compensation and correction techniques. The article concludes in section 6 with an evaluation of the prototype's optical characteristics, a view of its pros and cons, a discussion of possible applications, and a presentation of future work.

3 Prior Work

Because the Extended Virtual Table prototype represents a combination of a table-like display and a mirror beam splitter, this section discusses previous and related works from two areas: table-like projection systems and related mirror displays.

3.1 Table-Like Projection Systems

Krüger's Responsive Workbench (Krüger & Fröhlich, 1994; Krüger et al., 1995) is one of the pioneering table-like projection systems. The Responsive Workbench consists of a video projector that projects high-resolution stereoscopic images onto a mirror located under the table, which, in turn, reflects it in the direction of the table top (a ground glass screen). Analyzing the daily work situation of different types of computer users, Krüger et al. chose a workbench-like system as an adaptation to the human living and working environment.

Using the Responsive Workbench metaphor, a rich palette of similar rear-projection devices are available today that differ in size, mobility, and applied projection technology. Among these systems are Wavefront's ActiveDesk, Barco's BARON, Fakespace's ImmersaDesk Series, and the Responsive Workbench itself, which is sold by TAN Projectiontechnologies.

Although all of these systems are single-sided projection devices, a few two-sided (L-shaped) systems have been developed to offer a larger and (by the normally limited projection area) less constrained viewing space. TAN's Holobench, for instance, is an extension of the Responsive Workbench, and Barco's Consul has been developed based on the BARON Virtual Table.

Within the previous six years, an enormous variety of applications (concerning almost all VR areas) that involve table-like projection systems have been described. To mention all of these developments would be beyond the scope of this article.

3.2 Related Mirror Displays

As for stereoscopic screen-based desktop systems, occlusion caused by the user's hand or handheld input devices is a major drawback of table-like rear-projection systems. This disadvantage makes a visually undistorted direct interaction with the presented virtual scene difficult, especially if force-feedback devices such as a

PHANToM (Massie & Salisbury, 1994) are applied to superimpose virtual visual and virtual haptic spaces.

A number of devices have been developed during recent years that allow the user to reach into a virtual scene without causing any occlusion. These so-called "reach-in systems" apply a horizontally arranged small mirror to reflect the graphics that are displayed on a CRT screen mounted above the mirror. While the user is looking at the mirror, he or she can simultaneously operate a spatial input device below the mirror that, in most cases, provides force feedback in relation to the stereoscopically displayed visual information. Because usually neither the input device nor the user's hands are visible by looking at the mirror, the virtual environment can be visually perceived in accordance with the corresponding haptic information without causing visual conflicts created by occlusion.

Knowlton (1977), for instance, overlaid monoscopic 2-D keycap graphics on the user's view of an otherwise conventional keyboard by using a half-silvered mirror that reflected a CRT screen. This allowed the graphics to annotate the user's fingers within the illuminated workspace below the mirror instead of being blocked by them.

Schmandt's Stereoscopic Computer Graphic Workstation (Schmandt, 1983) is another early example of a reach-in arrangement that applies an electromagnetic tracking device for input in combination with a CRT screen and a half-silvered mirror. He superimposed 3-D graphics over the transmitted image of the working area below the mirror.

Poston and Serra (1994) developed the Virtual Workbench, but used a mechanical input device to overcome the magnetic field distortion problems of Schmandt's setup, which were caused by the interference between the CRT screen and the electromagnetic tracking device. A more recent development is the apparatus by Wiegand, von Schloerb, and Sachtler (1999), which is also named Virtual Workbench. Their system offers a trackball for input, a PHANToM for input and additional force feedback, and stereo speakers for auditory feedback.

Due to the small working volume of these devices, their applications are limited to near-field operations. Although some of these systems employ half-silvered mirrors instead of full mirrors for calibration purposes, only a few support AR tasks. The maturity of systems, however, renders exclusively virtual information. Several of these devices are commercially available (such as, the Reach-In Display by Reach-In Technologies or the Dextroscope by the Medical Imaging Group) and are mainly used for medical/industrial simulation and training or psychophysics and training research (Wiegand et al., 1999).

Bimber, Encarnação, and Schmalstieg (2000a) introduced the Transflective Pad, a handheld half-silvered mirror, employed in combination with a table-like rear-projection device. The six-DOF tracked mirror supported a dynamic extension of the limited viewing volume, which is provided by such semi-immersive projection devices. It was used, for instance, to view stereoscopically projected volumetric data on a Virtual Table (Wohlfahrter, Encarnação & Schmalstieg, 2000).

Bimber et al. (2000b) later extended the concept of the Transflective Pad toward AR. In this case, the Transflective Pad was applied as an interactive image plane that folded the viewer's optical path and merged the reflected graphics with the image of the real world. Consequently, it represented a possible solution to the occlusion problem that is related to rear-projection systems. The core idea of the Transflective Pad will serve as the basis for the optical extension that is described in this article.

4 Description of the Prototype

In this section, a description is provided that details how the optical extension is built and how it is used in conjunction with a table-like projection system from both a physical and a conceptual point of view.

4.1 Physical Arrangement

The Extended Virtual Table (xVT) prototype consists of a virtual and a real workbench (cf. figure 1).

A Barco BARON serves as the display device that projects 54 in. by 40 in. stereoscopic images with a res-

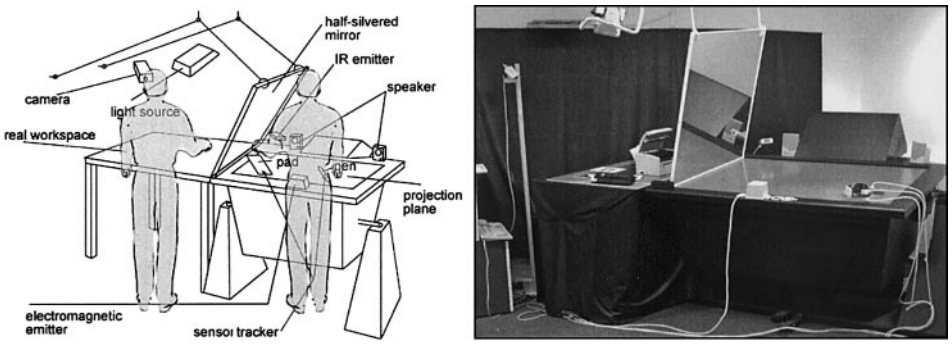

Figure 1. *Conceptual sketch (left) and photograph (right) of the xVT prototype.*

olution of 1280 × 1024 (or optionally 1600 × 1200/2) pixels on the backside of a horizontally arranged ground-glass screen. Shutter glasses, such as Stereographics' CrystalEyes or NuVision3D's 60GX, are used to separate the stereoimages for both eyes and make stereoscopic viewing possible. In addition, an electromagnetic tracking device, Ascension's Flock of Birds is used to support head tracking and tracking of spatial input devices (a pen and a pad). An Onyx InfiniteReality, which renders the graphics, is connected (via a TCP/IP intranet) to three additional PCs that perform speech recognition, speech synthesis, gesture recognition, and optical tracking.

A 40 in. by 40 in. and 10 mm thick pane of glass separates the virtual workbench (the Virtual Table) from the real workspace. It has been laminated with a half-silvered mirror foil, 3M's Scotchtint P-18, on the side that faces the projection plane, making it function like a front-surface mirror that reflects the displayed graphics. We have chosen a thick float-glass material (10 mm) to minimize the optical distortion caused by any bending of the mirror or irregularities in the glass. The half-silvered mirror foil, which is normally applied to reduce window glare, reflects 38% and transmits 40% light.[3] Note that this mirror extension costs less than $100. However, more-expensive half-silvered mirrors with better optical characteristics could be used instead (Edmund Industrial Optics, for example). With the bottom leaning onto the projection plane, the mirror is held by two strings that are attached to the ceiling. The length of the strings can be adjusted to change the angle between the mirror and the projection plane, or to allow an adaptation to the Virtual Table's slope. A light source is adjusted in such a way that it illuminates the real workbench, but does not shine at the projection plane.

In addition, the real workbench and the walls behind it were covered with a black awning to absorb light that otherwise would be diffused by the wallpaper beneath it and would cause visual conflicts if the mirror was used in a see-through mode.

Finally, a camera (a Videum VO) is applied to continuously capture a video stream of the real workspace, supporting an optical tracking of papermarkers that are placed on top of the real workbench.

4.2 General Functioning

Users can either work with real objects above the real workbench or with virtual objects above the virtual workbench.

Elements of the virtual environment, which are displayed on the projection plane, are spatially defined within a single world coordinate system that exceeds the boundaries of the projection plane, covering also the real workspace.[4]

The mirror plane splits this virtual environment into

3. Values valid if used with 6 mm thick regular glass.

4. In our case, the origin of the coordinate system is located at the center of the projection plane.

two parts that cannot be simultaneously visible to the user. This is due to the fact that only one part can be displayed on the projection plane. We determine the user's viewing direction to support an intuitive visual extension of the visible virtual environment. If, on the one hand, the user is looking at the projection plane, the part of the environment that is located over the virtual workbench is displayed. If, on the other hand, the user is looking at the mirror, the part of the environment located over the real workbench is transformed, displayed, and reflected in such a way that it appears as a continuation of the other part in the mirror.

Using the information from the head tracker, the user's viewing direction is approximated by computing the single line of sight that originates at her point of view and points towards her viewing direction. The plane the user is looking at (that is, the projection plane or mirror plane) is the one that is first intersected by this line of sight. If the user is looking at neither plane, no intersection can be determined, and nothing needs to be rendered at all.

In case the user is looking at the mirror, the part of the virtual environment behind the mirror has to be transformed in such a way that, if displayed and reflected, it appears stereoscopically and perspectively correct at the right place behind the mirror. As for the handheld Transflective Pad described by (Bimber et al., 2000b), we use an affine transformation matrix to reflect the user's viewpoint (that is, both eye positions that are required to render the stereoimages) and to inversely reflect the virtual environment over the mirror plane.

If we inversely reflect the graphical content from the viewer's averting side of the mirror to the opposite side and render it from the viewpoint that is reflected vice versa, the projected virtual environment will not appear as a reflection in the mirror. The user then sees the same scene that she would perceive without the mirror if the projection plane were large enough to visualize the entire environment. This is due to the neutralization of the computed inverse reflection by the physical reflection of the mirror.

Note that the transformation matrix can simply be added to a matrix stack or integrated into a scene graph without increasing the computational rendering cost. However, because its application reverses also the polygon order (which might be important for correct front-face determination, lighting, culling, and so forth), appropriate steps have to be taken in advance (such as, explicitly reversing the polygon order before reflecting the scene).

The mirror-plane parameters (a,b,c,d) can be determined within the world coordinate system in different ways:

- The electromagnetic tracking device can be used to support a three-point calibration of the mirror plane.
- The optical tracking system can be applied to recognize markers that are temporarily or permanently attached to the mirror.
- Because the resting points of the mirror on the projection plane are known and do not change, its angle can be measured using a simple ruler.

Note that all three methods can introduce calibration errors caused by either tracking distortion (electromagnetic or optical) or human inaccuracy. Our experiments have shown that the optical method is the most precise and less vulnerable to errors. Intersense's tracking system, for instance, is based on inertial navigation and ultrasound time-of-flight. It constitutes an alternative to magnetic tracking because it doesn't suffer from distortions caused by presence of metal in the environment (Foxlin, Harrington, & Pfeifer, 1998). To avoid visual conflicts between the projection and its corresponding reflection—especially for areas of the virtual environment whose projections are close to the mirror—we optionally render a clipping plane that exactly matches the mirror plane (that is, with the same plane parameters a,b,c,d). Visual conflicts arise if virtual objects spatially intersect the side of the user's viewing frustum that is adjacent to the mirror, because, in this case, the object's projection optically merges into its reflection in the mirror. The clipping plane culls away the part of the virtual environment that the user is not looking at. (We reverse the direction of the clipping plane, depending on the viewer's viewing direction while maintaining its position.) The result is a small gap between the mirror

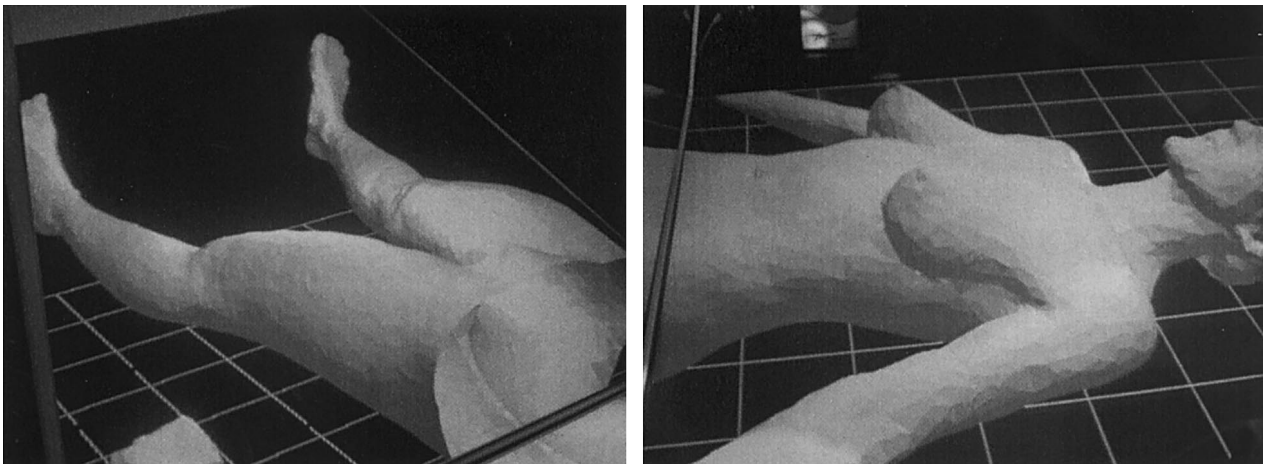

Figure 2. *A large coherent virtual content (a life-size human body for medical training) viewed in the mirror (left), or on the projection plane (right). The real workspace behind the mirror is not illuminated.*

and the outer edges of the viewing frustum in which no graphics are visualized. This gap helps to differentiate between projection and reflection and, consequently, avoids visual conflicts. Yet, it does not allow virtual objects that are located over the real workbench to reach through the mirror. We can optionally activate or deactivate the clipping plane for situations in which no or minor visual conflicts between reflection and projection occur to support a seamless transition between both spaces.

If the real workspace behind the mirror beam splitter is not illuminated, the mirror acts like a full mirror and supports a nonsimultaneous visual extension of an exclusively virtual environment (that is, both parts of the environment cannot be seen at the same time). Figure 2 shows a large coherent virtual scene whose parts can be separately observed by either looking at the mirror or at the projection plane. Note that none of the photographs shown in this article are embellished; they were taken as seen from the viewer's perspective (rendered monoscopically). However, the printouts may appear darker and with less luminance than in reality, mainly due to the camera response.

Figure 3 shows a simple example in which the mirror beam splitter is used as an optical combiner. If the real workspace is illuminated, both the real and the virtual environment are visible to the user, and real and virtual objects can be combined in the AR manner. Note that the ratio of intensity of the transmitted light and the reflected light depends upon the angle between the beam splitter and the projection plane. Acute angles highlight the virtual content, and obtuse angles let the physical objects shine through brighter.[5]

4.3 Interacting Through the Mirror

A large variety of interaction techniques has been explored for table-like projection systems. Van de Pol, Ribarsky, and Hodges (1999) present a good classification and evaluation of interaction techniques for such devices.

A tracked stylus is the main input device that we apply to support direct interaction in front of the mirror and indirect interaction with objects behind the mirror. In addition, a transparent pad or tablet (which is also tracked) is applied to feature two-handed interaction as described in several studies (Coquillart & Wesche, 1999; Encarnação, Bimber, Schmalstieg, & Chandler, 1999; Schmalstieg, Encarnação, & Szalavári, 1999).

5. This is similar to the electronic shuttering of HMDs.

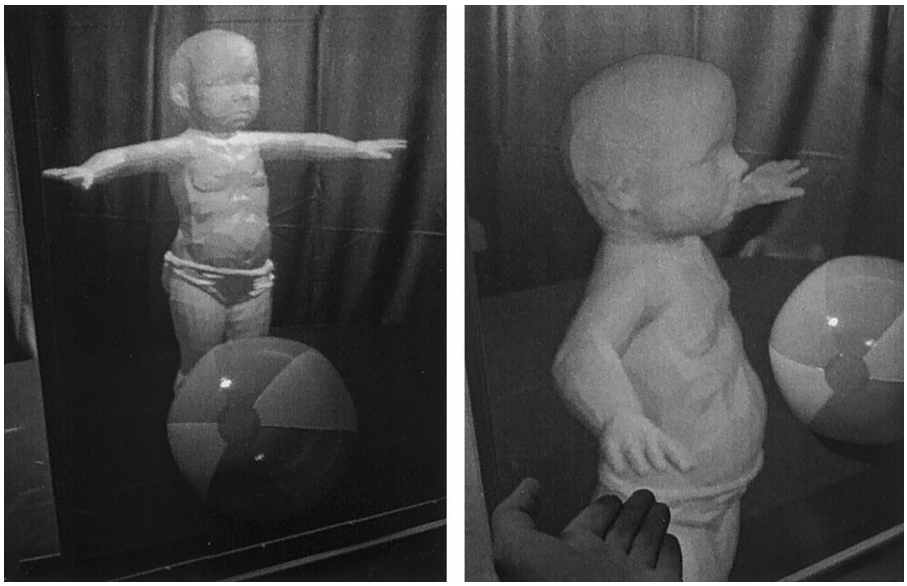

Figure 3. *Left: Real objects behind the mirror (the ball) are illuminated and augmented with virtual objects (the baby). The angle between mirror and projection plane is 60 deg. Right: Without attaching a clipping plane to the mirror, the baby can reach her arm through the mirror. The angle between mirror and projection plane is 80 deg.*

Virtual objects can be exchanged between both sides of the mirror in different ways, in the see-through mode as well as in the opaque mode. For example, they can be picked with the stylus, either directly or from a distance over the virtual workbench, or indirectly over the real workbench. Virtual objects can then be pushed or pulled through the mirror (see figure 4), either directly or indirectly, by using nonlinear arm-extension methods, such as Go-Go (Poupyrev, Billinghurst, Weghorst, & Ichikawa, 1996).

As illustrated in figure 5, a virtual laser beam that is cast from the stylus through the mirror is used to select and manipulate (that is, to move and place) virtual objects behind the mirror plane. This ray tool allows for interaction with virtual objects on a remote basis and offers an indirect object placement by "beaming" the object back and forth along the ray.

Virtual-virtual and real-virtual object collision detection is applied over the real and the virtual workbench to simulate a realistic interaction behavior between objects. This fundamental collision detection capability enables us to implement gravity-based automatic object placement methods, as described by Breen, Whitaker, Rose, and Tuceryan (1996).

Real objects also can occlude virtual ones, as figures 3 and 5 show. This is achieved by implementing a method for see-through optical combination introduced by Breen et al. (1996) and an object-based blurring technique as described by Fuhrmann, Hesina, Faure, and Gervautz (1999).

In addition, optical tracking is applied above the real workbench. A camera captures a single video stream of the real workspace and tracks movable paper markers over the real workbench (Kato, Billinghurst, Blanding, & May, 1999). We use the markers to track real world objects, for calibrating the setup, or—as illustrated in figure 5—as placeholders for multimedia information (images, video or textual information, and so on).

A pointing method that is similar to the ones described by Whitaker, Crampton, Breen, Tuceryan, and Rose (1995) and Fuhrmann, Schmalstieg, and Purgathofer (1999) is used to register stationary real objects,

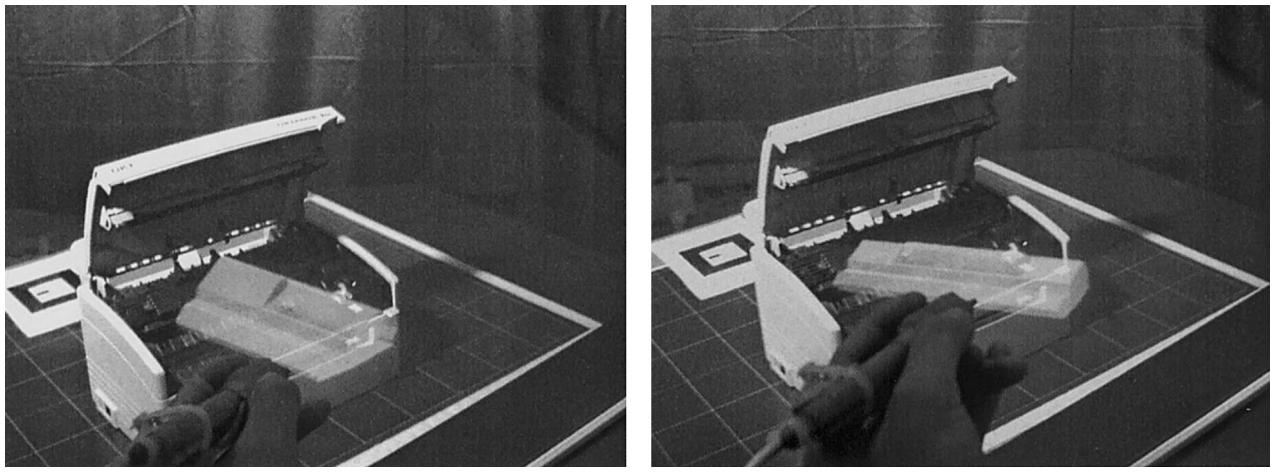

Figure 4. *A virtual object is pushed through the mirror, a clipping plane is attached to the mirror.*

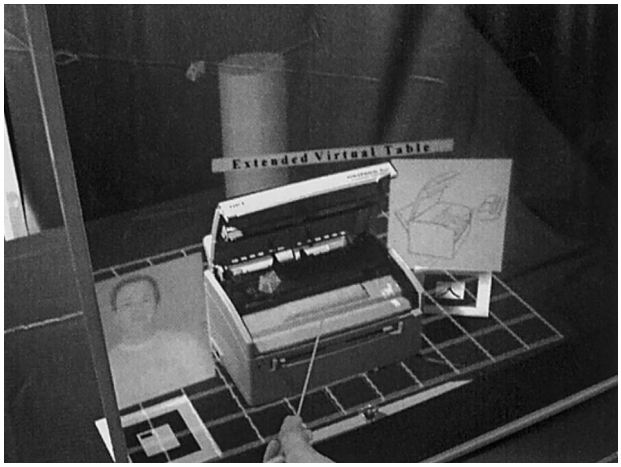

Figure 5. *Ray casting and optical tracking within an augmented real environment. A virtual cartridge is mounted into a real printer.*

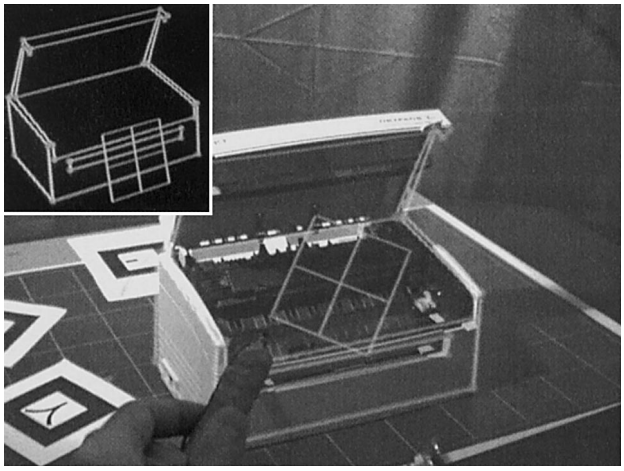

Figure 6. *Registering a real object using pointing.*

whose locations are needed for collision detection and occlusion.

Virtual cross wires are attached to the stylus tip and rendered monoscopically to allow a precise adjustment of landmark points on the real object's surface (as illustrated in figure 6). The spatial position of surface points can be determined by adjusting them with the cross wires from at least two different perspectives. The computed coordinates are used to minimize the distance function between measurements and predefined corresponding points by applying Powell's direction set method (Press, Teukolsky, Vetterling, & Flannery, 1992).

The disadvantage of pointing and ray casting for indirect interaction is that they cannot be used for directly defining spatial points. To support an immediate definition of points within the 3-D free space, we offer a remote usage of the same tools (the pen and the pad tools) that are applied directly above the virtual workbench. To allow an ergonomic usage of the remote tools, they can be frozen in their current position and

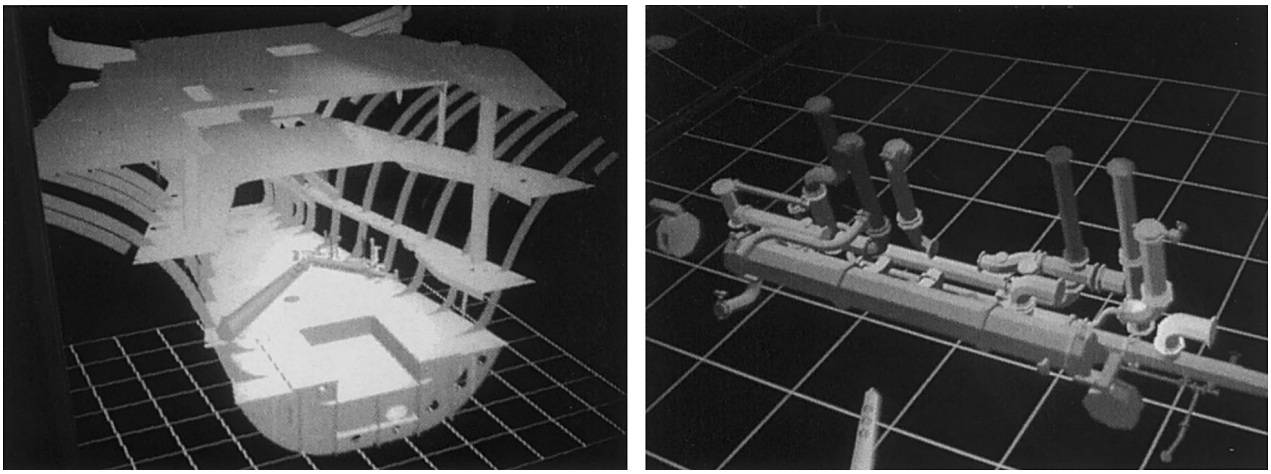

Figure 7. *Distance manipulation with remote tools behind the mirror (left) and close manipulation above the virtual workbench with direct tools (right).*

orientation until the input devices are relocated to a more convenient posture. After unfreezing the remote tools, the new position and orientation offset is computed and used.

Sketching and drafting are examples where spatial input tools are better suited than the ray or pointing techniques.

Although remote tools allow for distance manipulation, direct tools support close manipulation. Figure 7 illustrates an example of an intuitive and smooth transition between distance and close manipulation. Using the remote pen, distance manipulation of ship components (pipeline clusters in figure 7) within a ship section, which is visualized behind the mirror, is supported as long as the user looks at the mirror. If the user picks a component within the ship section and then looks at the virtual workbench, the component is automatically transformed to fill out the entire space of the projection area at the virtual workbench. It is translated, rotated, and scaled in such a way that a convenient close manipulation with the direct pen tool is supported. If, however, the user picks the component above the virtual workbench and then looks at the mirror, the component is automatically transformed back to the ship section and is downscaled to fit into the ship coordinate system.

5 Distortion Compensation and Correction

5.1 Optical Distortion

Optical distortion is caused by the elements of an optical system. It does not affect the sharpness of a perceived image, but rather its geometry. The distortion can be corrected optically (for example, by applying additional optical elements that physically rescind the effect of other optical elements) or computationally (such as, by pre-distorting generated images). Although optical correction may result in heavy optics and non-ergonomic devices, computational correction methods might require high computational performance.

Optical distortion is critical in AR applications because it prevents the precise registration of the virtual and real environment.

The purpose of the optics used in HMDs, for instance, is to project two equally magnified images in front of the user's eyes in such a way that they fill out a wide field of view (FOV) and fall within the range of accommodation (focus). To achieve this, however, lenses are used in front of the miniature displays or in front of mirrors that reflect the displays within see-through HMDs. The lenses, as well as the curved display surfaces of the miniature screens, may introduce

optical distortion, which is normally corrected computationally to avoid heavy optics that would result from optical approaches.

Because, for HMDs, the applied optics form a centered (on-axis) optical systems, pre-computation methods can be used to efficiently correct geometrical aberrations during rendering. Rolland and Hopkins (1993) describe a polygon wrapping technique as a possible correction method for HMDs. The optical distortion for HMDs is constant (because the applied optics is centered[6]), and so a two-dimensional lookup table is precomputed that maps projected vertices of the virtual objects' polygons to their pre-distorted location on the image plane. Note that this requires subdividing polygons that cover large areas on the image plane. Instead of pre-distorting the polygons of projected virtual objects, the projected image itself can be pre-distorted, as described by Watson and Hodges (1995), to achieve a higher rendering performance.

Correcting optical distortion is more complex for the mirror beam splitter extension, because, in contrast to HMDs, the image plane that is reflected by the mirror is not centered with respect to the optical axes of the user, but is off axis in most cases. In fact, the alignment of the reflected image plane dynamically changes with respect to the moving viewer, whereas the image plane itself remains at a constant spatial position in the environment. In the case of the xVT, there are mainly three sources of optical distortion projector calibration, mirror flexion, and refraction. Note that we correct optical distortion only while the user is working in the see-through mode (that is, while looking through the half-silvered mirror at an illuminated real environment). For exclusive VR applications, optical distortion is not corrected even if the mirror is used as an extension.

5.1.1 Projector Calibration. The projector that is integrated into the Virtual Table can be calibrated in a way that it projects distorted images onto the ground-glass screen. Projector-specific parameters (such as geometry, focus, and convergence) can usually be adjusted manually or automatically using camera-based calibration devices. Whereas a precise manual calibration is very time consuming, an automatic calibration is normally imprecise, and most systems do not offer a geometry calibration (only calibration routines for convergence and focus).

For exclusive VR purposes, however, we can make use of the fact that small geometric deviations are ignored by the human-visual system. In AR scenarios, on the other hand, even slight misregistrations can be sensed.

To precisely and easily calibrate geometry, we apply a two-pass method and render a regular planar grid (U) that largely covers the projection plane. The distorted displayed grid is then sampled with a device that is able to measure 2-D points on the table top. After a transformation of the sampled grid (D) into the world coordinate system, it can be used to pre-distort the projected image, because with D, the geometrical deviation ($U - D$), which is caused by the miscalibrated projector can be expressed. A pre-distorted grid (P) can then be computed with $P = U + (U - D)$. If we project P instead of U, the pre-distortion is rescinded by the physical distortion of the projector, and the visible grid appears undistorted.

To pre-distort the projected images, however, we first render the virtual environment into the frame buffer, then map the frame buffer's content as texture onto P (while retaining the texture indices of U and applying a bilinear texture filter), and render P into the beforehand cleaned frame buffer as described by Watson and Hodges (1995) for HMDs. Note that this is done for both stereoimages at each frame.

To sample grid points, we apply Dukane Corporation's Mimio, a device that is usually used to track pens on a whiteboard. The Mimio is a hybrid (ultrasonic and infrared) tracking system for planar surfaces which is more precise and less susceptible to distortion than the applied electromagnetic tracking device. It provides a stable positional resolution of 0.3 mm with an update rate of 87 measurements/second. As illustrated in figure 8, its receiver has been attached to a corner of the Virtual Table. (Note the area where the Mimio cannot receive correct data from the sender, due to distortion. This area has been specified by the manufacturer.)

6. Rotations of the eyeball are not considered.

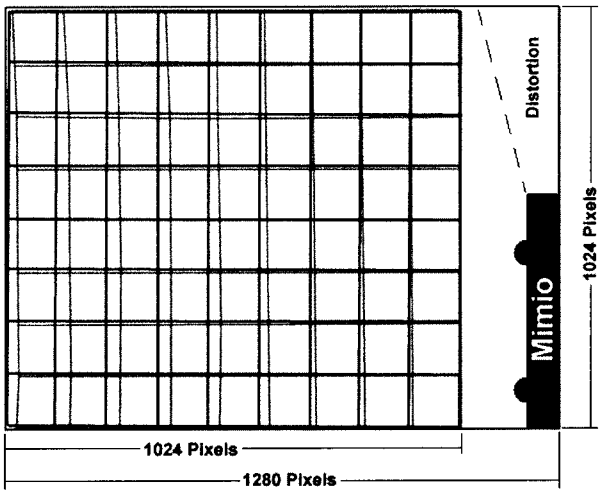

Figure 8. *Sampled distorted grid (gray) and pre-distorted grid (black) after projection and resampling. The physical arrangement of the sampling device on the Virtual Table.*

Because the supported maximal texture size of the used rendering package is 1024×1024 pixels, U is rendered within the area (of this size) that adjoins to the mirror. We found that 10×9 sample points for an area of 40 in. by 40 in. on the projection plane is an appropriate grid resolution that avoids oversampling but is sufficient enough to capture the distortion.

Figure 8 illustrates the sampled distorted grid D (gray), and the pre-distorted grid P (black) after it has been rendered and resampled. Note, that figure 8 shows real data from one of the calibration experiments. Other experiments delivered similar results.

The calibration procedure has to be done once or once in a while, because the distortion behavior of the projector can change over time.

5.1.2 Mirror Flexion. For the mirror beam splitter, a thick float-glass material has been selected to keep optical distortion (caused by bending) small. Due to gravity, a deformation of the reflected image cannot be avoided as a slight flexion of the mirror affects the first-order imaging properties of our system (that is, the magnification and location of the image).

Figure 9 (left) illustrates the optical distortion caused by flexion. A bent mirror does not reflect the same projected pixel for a specific line of sight than an unbent mirror.

Correction of the resulting distortion can be realized by transforming the pixels from the position where they should be seen (reflected by an ideal unbent mirror) to the position where they can be seen (reflected by the bent mirror) for the same line of sight.

Because a transformation of every single pixel would be inefficient, the correction of mirror flexion can be combined using the method described in subsection 5.1.1.

For every point $\vec{U}$ of the undistorted grid U, the corresponding point of reflection $\vec{R}$ on the bent mirror has to be determined with respect to the current eye position of the viewer $\vec{E}$. Note that this requires knowledge of the mirror's curved geometry. If the surface of the mirror is known, $\vec{R}$ can simply be calculated by reflecting $\vec{U}$ over the known (unbent) mirror plane. (The reflection matrix, described by Bimber et al., (2000b) can be used for this.) Then, find the intersection between the bent mirror's surface and the straight line that is spanned by $\vec{E}$ and the reflection of $\vec{U}$. Note that, if the mirror's entire surface is not known, an interpolation between sample points (taken from the mirror's surface) can be done to find an appropriate $\vec{R}$. If $\vec{R}$ has been determined, the normal vector at $\vec{R}$ has to be computed. This is also possible with the known mirror geometry. The normal vector usually differs from the normal vector (a,b,c) of the unbent mirror, which is the same for every point on the unbent mirror's surface. With the computed $\vec{R}$ and its normal, the equation parameters (a',b',c',d') for a plane that is tangential to $\vec{R}$ are identified. To compute the position to which $\vec{U}$ has to be moved on the projection plane to be visible for the same line of sight in the bent mirror, $\vec{E}$ has to be reflected over (a',b',c',d'). The intersection between the projection plane and the straight line that is spanned by the reflection of $\vec{E}$ and $\vec{R}$ is $\vec{U}'$.

However, it is not sufficient to transform the undistorted grid with respect to the mirror's flexion and the observer's viewpoint only, because the projector distortion (described in subsection 5.1.1) is not taken into account. To imply projector distortion, every $\vec{U}'$ has to be pre-distorted, as described in subsection 5.1.1. Be-

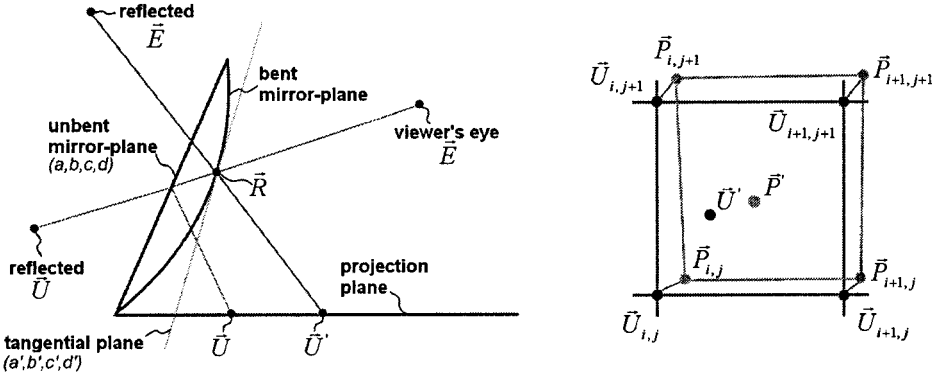

Figure 9. *Optical distortion caused by flexion (left). Bilinear interpolation within an undistorted/pre-distorted grid cell (right).*

cause the $\vec{U}'$s normally do not match their corresponding $\vec{U}$s, and a measured distortion $\vec{D}'$ for each $\vec{U}'$ does not exist, an appropriate pre-distortion offset can be interpolated from the measured (distorted) grid D (as illustrated in figure 9 (right)). This can be done by bilinear interpolating between the corresponding points of the pre-distorted grid P that belongs to the neighboring undistorted grid points of U, which form the cell that encloses $\vec{U}'$.

In summary, we have to compute a new pre-distorted grid (P'), depending on the mirror's flexion (R), the current eye-positions of the viewer ($\vec{E}$), and the projector distortion (D). The resulting P' can then be textured as described in subsection 5.1.1 (for both stereo-images at each frame).

Note that finding an exact method that precisely determines the mirror's flexion is a topic for our future research. Using the electromagnetic tracking device to sample the mirror's surface turned out to be insufficient due to the nonlinear tracking distortion over the extensive area.

5.1.3 Refraction. On the one hand, a thick pane of glass stabilizes the half-silvered mirror and consequently minimizes optical distortion caused by flexion. On the other hand, however, it causes another optical distortion that results from refraction. Because the transmitted light that is perceived through the half-silvered mirror is refracted but the light that is reflected by the front surface mirror foil is not, the transmitted image of the real environment cannot be precisely registered to the reflected virtual environment, even if their geometry and alignment match exactly within the world coordinate system.

All optical systems that use any kind of see-through elements have to deal with similar problems. Although for HMDs, aberrations caused by refraction of the lenses are mostly assumed to be static[7] (as stated by Azuma (1997)), they can be corrected with paraxial analysis approaches. For other setups, such as the reach-in systems that were mentioned in section 1 or our mirror extension, aberrations caused by refraction are dynamic because the optical distortion changes with a moving viewpoint. Wiegand et al. (1999), for instance, estimated the displacement caused by refraction for their setup to be less than 1.5 mm, predominantly in the +y direction of their coordinate system. Although an estimation of a constant refraction might be sufficient for their apparatus (a near-field virtual environment system with fixed viewpoint that applies a relatively thin (3 mm) half-silvered mirror), our setup requires a more precise definition because it is not a near-field VE system but rather a mid-field VR/AR system, it considers a head-tracked viewpoint, and it applies a relatively thick half-silvered mirror (10 mm).

7. We do not consider rotations of the eyeballs.

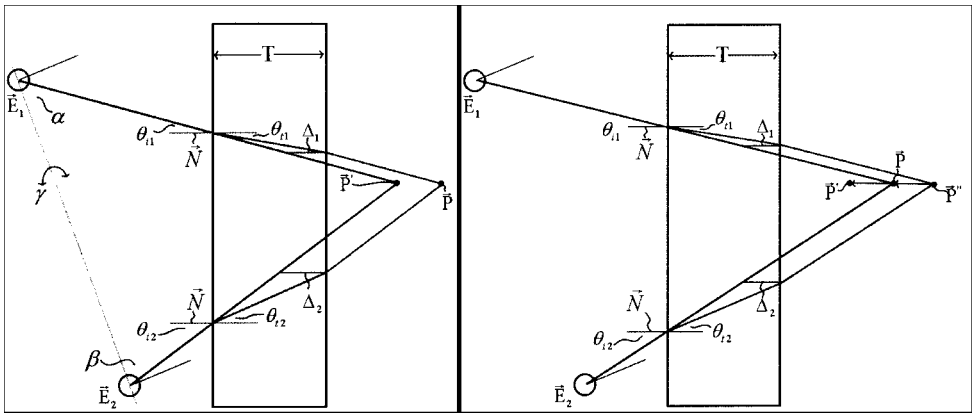

Figure 10. *Precise refraction method (left) and refraction approximation (right).*

Because we cannot pre-distort the refracted transmitted image of the real world, we artificially refract the reflected virtual world instead to make both images match. Figure 10 illustrates our approaches.

With reference to figure 10 (left), the observer's eyes ($\vec{E}_1$, $\vec{E}_2$) have to converge to see a point in space ($\vec{P}'$) in such a way that the geometric lines of sight (in black) intersect in $\vec{P}'$. If the observer sees through a medium whose density is higher than the density of air, the geometric lines of sight are bent by the medium, and she perceives the point in space ($\vec{P}$) where the resulting optical lines of sight (colored in dark gray) intersect; that is, she perceives $\vec{P}$ instead of $\vec{P}'$ if refraction bends her geometric lines of sight.

To artificially refract the virtual environment, our goal is to translate every point $\vec{P}$ of the virtual environment to its corresponding point $\vec{P}'$, following the physical rules of refraction. Note that all points $\vec{P}$ are virtual points that are not physically located behind the mirror beam splitter and consequently are not physically refracted by the pane of glass, but are reflected by the front surface mirror. The resulting transformation is curvilinear rather than affine. Thus, a simple transformation matrix cannot be applied.

Using Snell's law for refraction, we can compute the optical line of sight for a corresponding geometric line of sight. Note that, in the case of planar plates, both lines of sight are simply shifted parallel along the plate's normal vector ($\vec{N}$)[8] by an amount (Δ) that depends on the entrance angle (θ_i) between the geometric line of sight and $\vec{N}$, its thickness (T), and the refraction index (η), a material-dependent ratio that expresses the refraction behavior compared to vacuum (as an approximation to air).

The amount of translation (Δ) can be computed as follows:

$$\theta_t = \sin^{-1}\left(\frac{\sin\,\theta_i}{\eta}\right) \tag{1}$$

The refraction-dependent amount of displacement along the plate's normal vector can be computed as follows.

$$\Delta = T\left(1 - \frac{\tan\,\theta_t}{\tan\,\theta_i}\right), \tag{2}$$

with

$$\lim\,\theta_i \rightarrow \frac{\pi}{2} \Rightarrow \Delta = T \text{ and } \lim\,\theta_i \rightarrow 0 \Rightarrow$$
$$\Delta = T\left(1 - \frac{\sin\,\theta_t}{\sin\,\theta_i}\right) = T\left(1 - \frac{1}{\eta}\right) = const.$$

With constant T (that is, 10 mm) and constant η (1.5 for regular glass), the refractor of a ray that is spanned

8. If the mirror is bent, as described in subsection 5.1.2, the normal vector of the mirror plane is not constant, and the corresponding normals of the points on the mirror surface that are intersected by the actual lines of sight have to be applied.

by the two points ($\vec{P}_1$, $\vec{P}_2$) depends upon the entrance angle (θ_i) and can be computed as follows (in parameter representation):

$$\vec{R} = \vec{P}_1 + \Delta\frac{\vec{N}}{|\vec{N}|} + \lambda(\vec{P}_2 - \vec{P}_1) \qquad (3)$$

Note that the optical line of sight is the refractor that results from the geometric line of sight that is spanned by the viewer's eye ($\vec{E}$) and the point in space ($\vec{P}$) at which she is looking. Because no analytical correction method exists, we apply a numerical minimization to precisely refract virtual objects that are located behind the mirror beam splitter by transforming their vertices within the world coordinate system. Note that, similar to Rolland's approach (Rolland & Hopkins, 1993), our method also requires subdividing large polygons of virtual objects to sufficiently express the refraction's curvilinearity.

The goal is to find the coordinate $\vec{P}'$ where the virtual vertex $\vec{P}$ has to be translated in such a way that $\vec{P}$ appears spatially at the same position as it would appear as a real point, observed through the half-silvered mirror (that is, refracted). To find $\vec{P}'$, we first compute the geometric lines of sight from each eye ($\vec{E}_1$, $\vec{E}_2$) to $\vec{P}$. We then compute the two corresponding optical lines of sight using equation 3 and their intersection ($\vec{P}''$). During a minimization procedure (Powell's direction set method (Press et al., 1992)) we minimize the distance between $\vec{P}$ and $\vec{P}''$ while continuously changing the angles α,β (simulating the eyes' side-to-side shifts and convergence) and γ (simulating the eyes' up-and-down movements), and use them to rotate the geometric lines of sight over the eyes' horizontal and vertical axes. (The axes can be determined from the head tracker.) The rotated geometric lines of sight result in new optical lines of sight and consequently in a new $\vec{P}''$.

Finally, $\vec{P}'$ is the intersection of the (by some α,β,γ) rotated geometric lines of sight where $|\vec{P} - \vec{P}''|$ is minimal (that is, below some threshold ε). This final state is illustrated in figure 10.

In summary, we have to find the geometric lines of sight whose refractors (the corresponding optical lines of sight) intersect in $\vec{P}$ and then calculate the precise coordinate of $\vec{P}'$ as intersections of the determined geometric lines of sight. Because $\vec{P}'$ is unknown, the resulting minimization problem is computationally expensive and cannot be solved in real time. To achieve a high performance on an interactive level, we implemented an approximation of the presented precise method.

Table 1. *Comparison Between Precise Refraction and Approximated Refraction*

Displacement caused by refraction (mm)	Minimal	Maximal	Average
Precise Method	3.75	10.34	6.08
Approximation Method	3.53	9.78	5.95

With reference to figure 10 (right), we compute the refractors of the geometric lines of sight to the vertex $\vec{P}$ and their intersection $\vec{P}''$. Because the angular difference between the unknown geometric lines of sight to the unknown $\vec{P}'$ and the geometric lines of sight to $\vec{P}''$ is small, the deviations of the corresponding refractors are also small. We approximate $\vec{P}'$ with $\vec{P}' = \vec{P} + (\vec{P} - \vec{P}'')$.

To compare the effectiveness of the outlined analytical approximation with the precise numerical method, we refracted vertices that covered the entire working volume behind the mirror beam splitter over time (that is, from different points of view) with both the approximation and the precise method. The results are shown in table 1. (The minimization procedure was executed with a threshold of $\varepsilon = 0.01$ mm.)

The spatial distance between the approximately refracted points and their corresponding precisely refracted points serves as error function. The results are shown in table 2.

Note that the average deviation between precise method and approximation is far below the average positional accuracy of the electromagnetic tracking device as described in subsection 5.2. Thus, a higher optical distortion is caused by the inaccurate head tracker than by applying the approximation to correct refraction misalignments. However, if refraction is not handled at all, the resulting optical distortion is higher than the one caused by tracking errors. Also note that the presented

Table 2. *Average Deviation Between Precise Method and Approximation*

	Minimal	Maximal	Average
Deviation (mm)	0.03	1.38	0.19

approximation is correct only for plane parallel plates. If the mirror is bent, the normals at the intersections of the in-refractor and the out-refractor differ. However, we approximated this by assuming that the mirror's flexion is small, and the two normals are roughly equivalent. Determining both normals is computationally too expensive for interactive applications and does not (in case our system) result in major visual differences.

5.2 Nonoptical Distortion

Accurate registration requires accurate tracking. In addition to the nonlinear tracking distortion, end-to-end system delay (the time difference between the moment that the tracking system measures a position/orientation and the moment the system reflects this measurement in the displayed image) or lag causes a "swimming" effect, in which virtual objects appear to float around real objects. However, because ideal tracking devices do not yet exist, we apply smoothing filters (sliding average windows) to filter high-frequency subbands (noise) from the tracking samples and prediction filters (Kalman filters (Azuma, 1995)) for orientation information and linear prediction for position information to reduce the swimming effect.

The applied tracking device (Ascension's Flock of Birds) provides a static positional accuracy of 2.5 mm (by 0.75 mm positional resolution) and a static angular accuracy of 0.5 deg. (by 0.1 deg. angular resolution). The highest update rate (without system delay) is 100 measurements/second.

6 Discussion and Future Work

We have presented an optical extension of table-like projection systems. Using a large optical combiner, we can enrich the scope of such display devices by allowing a viewing and interaction beyond the display's boundaries. In addition, the extension enables such projection systems to feature AR tasks for which, according to the term *projection-based VR,* we coined the term *projection-based AR*. In contrast to immersive display technology, semi-immersive workbenches can easily be integrated into habitual work environments. Furthermore, the extension offers the combination of a virtual and a real workbench, representing a hybrid workplace.

The projection-based AR concept attempts to detach the display device from the user and consequently addresses some of the drawbacks that are related to HMDs by benefitting from advantages of the well-established projection-based VR approach.

Our experiments have shown that the optical characteristics (such as field of view, resolution, and image brilliance) of our current setup can catch up with the optical characteristics of high-end optical see-through HMDs. However, in contrast to HMDs, the utilized projection technology is scalable, both in resolution and field of view. Addressing the ergonomic factor, glasses that have to be worn if active or passive shuttering is applied are much lighter and less cumbersome than HMDs. Furthermore, because the reflected image plane can be spatially better aligned with the real environment that has to be augmented, the fixed focal length problem related to HMDs (wherein the image plane is attached to the viewer) is reduced. Consequently, an easier eye accommodation is supported.

Patrick et al. (2000) and Johnson and Steward (1999) indicate that statistically no significant difference in acquiring spatial knowledge can be found between closed-view HMDs and large projection-screen conditions. However, they also state that the lower-cost projection screens are an attractive alternative to expensive and uncomfortable HMDs (that is, discomfort due to the poor ergonomics and simulation sickness that is due to fast head motions). These findings can also be applied to relate optical see-through HMDs to projection-based AR devices.

Beside individual drawbacks, head-attached displays in general (HMDs, head-mounted projective displays (Parsons & Rolland, 1998; Inami et al., 2000) and pro-

jective HMDs (Kijima & Ojika, 1997), suffer from an imbalanced ratio between heavy optics (lenses, displays, and projectors) that results in cumbersome and uncomfortable devices and ergonomic devices with a low image quality. This ratio can be better balanced by introducing devices that detach the display technology and the supporting optics from the user.

However, the xVT lacks in three major factors: mobility, direct interaction with augmented real objects, and single-user application.

We believe that stable and precise long-range tracking will exist in the near future, enabling AR applications using HMDs and head-mounted projector displays to be highly mobile. Nevertheless, the intention of the xVT is to combine two table-like workplaces in which the users focus on the workspace above the workbenches. For this, we neither require long-range tracking nor a high degree of mobility.

HMDs also offer direct interaction with augmented real objects that are in arm's reach of the user. In the case of the xVT, the mirror represents a physical barrier for the user's hands and the input devices and, to a large degree, prevents direct interaction with superimposed objects. We can either directly interact with real objects on the real workbench and with virtual objects on the virtual workbench, or indirectly interact with virtual objects above the real workbench through the mirror. To find additional interaction metaphors that support a more convenient and more realistic handling of the superimposed real workspace will be a goal of our future research. Additional mechanical installation on the real workbench (such as a turntable or robot arms) can be useful to remotely interact with real objects. Input devices that are better suited for indirect interaction (such as the Cubic Mouse (Fröhlich & Plate, 2000)) can be used in addition to traditional input tools. Those devices will provide force feedback in upcoming versions, which makes a direct interaction more convincing.

Although HMDs provide an individual image plane for each participant of a multiple-user session, users of large projection systems have to share the same image plane. Thus, multiple-user scenarios are difficult to realize with such technology. Because the xVT also applies a single large projection plane, it faces the same problem. Although some solutions exist that simultaneously support two users (for example, Agrawala et al. (1997)), they are not widely applied because they require special hardware.

The xVTs nonsimultaneous viewing of the two projection spaces (in front of the mirror and behind it) can be either interpreted as a disadvantage because the application of a second projector and another diffuse projection plane would support a simultaneous viewing (as is the case for L-shaped workbenches) or as an advantage, because no second projector is required. However, an additional opaque projection plane would make a see-through mode impossible. Instead, a front-projected or rear-projected semitransparent (nonreflective) or a rear-projected holographic projection plane (such as Pronova's HoloPro) could be applied, but at the cost of image quality or the viewing range (compared to the much better optical see-through characteristic of half-silvered mirrors).

We envision the possible application areas of the xVT to range from visualization of scientific data or simulation results in combination with real-world objects, over semi-immersive telecooperation to hybrid modeling and assembly.

In hybrid modeling and assembly applications, virtual mockups (VMUs) can be modeled above the virtual workbench and then be assembled to corresponding physical mockups (PMUs) that are located on the real workbench. A combination of VMUs and PMUs to hybrid mockups (HMUs) supports a more realistic early design review and possible early refinements within the conceptual product design phase.

The xVT can also be used to support a shared design review and modeling between several distributed parties within a semi-immersive telecooperation session. The telepresence factor is limited in most telecooperation systems that apply desktop screens because communicating through a screen barely gives an impression of presence. A stereoscopic and life-sized video stream of a remote collaborator can be visualized in the mirror of the local user, giving her the impression that the collaborator is standing right beside her virtual workbench. Two or more participants that could see each other (and each other's local design space) within their local mir-

rors could bring together and discuss different virtual components that could be exchanged intuitively by passing them through their mirrors. In this way, they make use of an efficient, location-spanning form of teamwork.

Embedding the display technology into the real environment potentiality opens new application areas for AR, as it did for VR. However, such projection-based AR/VR displays will not substitute head-attached displays, but rather present an application-specific alternative.

References

Agrawala, M., Beers, A. C., Frölich, B., McDowall, I., Hanrahan, P., & Bolas, M. (1997). The two-user responsive workbench: Support for collaboration through individual views of a shared space. *Computer Graphics Proceedings, Annual Conference Series SIGGRAPH '97,* 327–332.

Azuma, R. (1995). *Predictive tracking for augmented reality.* Unpublished doctoral dissertation, University of North Carolina at Chapel Hill, TR95-007.

———. (1997). A survey of augmented reality. *Presence: Teleoperators and Virtual Environments, 6*(4), 355–385.

Bajura, M., Fuchs, H., & Ohbuchi, R. (1992). Merging virtual objects with the real world: Seeing ultrasound imagery within the patient. *Computer Graphics (Proceedings of the ACM SIGGRAPH '92),* 203–210.

Bennett, D. (2000). Elumens [on-line]. Available: http://www.virtual-reality.com.

Bimber, O., Encarnação, L. M., & Schmalstieg, D. (2000a). Real mirrors reflecting virtual worlds. *Proceedings of IEEE Virtual Reality 2000 (VR '00) Conference, 1,* 21–28.

———. (2000b). Augmented reality with back-projection systems using transflective surfaces. *Computer and Graphics Forum (Proceedings of EUROGRAPHICS 2000), 19*(3), 161–168, 530.

Breen, D. E., Whitaker, R. T., Rose, E., & Tuceryan, M. (1996). Interactive occlusion and automatic object placement for augmented reality. *Computer and Graphics Forum (Proceedings of EUROGRAPHICS '96), 15*(3), C11–C22.

Coquillart, S., & Wesche, G. (1999). The virtual palette and the virtual remote control panel: A device and an interaction paradigm for the Responsive Workbench. *In Proc. of the 1999 IEEE Virtual Reality Conference,* 213–216.

Cruz-Neira, C., Sandin, D., & DeFanti, T. (1993). Surround-screen projection-based virtual reality: The design and implementation of the CAVE. *Computer Graphics, Proceedings of SIGGRAPH '93,* 135–142.

Encarnação, L. M., Bimber, O., Schmalstieg, D., & Chandler, S. D. (1999). A translucent sketchpad for the Virtual Table Exploring Motion-based Gesture Recognition. *Computer Graphics Forum (Proceedings of EUROGRAPHICS '99), 19*(3), C-277–C-285.

Foxlin, E., Harrington, M., & Pfeifer, G. (1998). Constellation: A wide-range wireless motion-tracking system for augmented reality and virtual set applications. *Proceedings of SIGGRAPH '98,* 371–378.

Fuhrmann, A., Hesina, G., Faure, F., & Gervautz, M. (1999). Occlusion in collaborative augmented environments. *Proceedings of 5th Eurographics Workshop on Virtual Environments,* 179–190.

Fuhrmann, A., Schmalstieg, D., & Purgathofer, W. (1999). Fast calibration for augmented reality. *Proceedings of ACM Symposium on Virtual Reality Software and Technology 1999 (VRST '99),* 166–167.

Fröhlich, B., & Plate, J. (2000). The Cubic Mouse: A new device for three-dimensional input. *Proceedings of ACM CHI 2000,* 526–531.

Haber, R. N., & Hershenson, M. (1973). *The psychology of visual perception.* New York: Holt, Rinehart & Winston.

Inami, M., Kawakami, N., Sekiguchi, D., Yanagida, Y., Maeda, T., & Tachi, S. (2000). Visuo-haptic display using head-mounted projector. *Proceedings IEEE Virtual Reality '00,* 233–240.

Johnson, D. M., & Steward, J. E. (1999). Use of virtual environments for the acquisition of spatial knowledge: Comparison among different visual displays. *Military Psychology, 11*(2), 129–148.

Kato, H., Billinghurst, M., Blanding, R., & May, R. (1999). The ARToolkit [on-line]. Available: http://www.hitl.washington.edu/research/shared_space.

Kijima, R., & Ojika, T. (1997). Transition between virtual environment and workstation environment with projective head-mounted display. *Proceedings of IEEE Virtual Reality Annual International Symposium,* 130–137.

Knowlton, K. C. (1977). Computer displays optically superimpose on input devices. *Bell System Technical Journal, 56*(3), 367–383.

Krüger, W., Bohn, C. A., Fröhlich, B., Schüth, H., Strauss, W., & Wesche, G. (1995). The Responsive Workbench: A virtual work environment. *IEEE Computer, 28*(7), 42–48.

Krüger, W., & Fröhlich, B. (1994). The Responsive Workbench. *IEEE Computer Graphics and Applications, 14*(3), 12–15.

Massie, T. H., & Salisbury, J. K. (1994). The PHANToM haptic interface: A device for probing virtual objects. *Proceedings of the ASME Dynamic Systems and Control Division, 55*(1), 295–301.

Parsons, J., & Rolland, J. P. (1998). A non-intrusive display technique for providing real-time data within a surgeon's critical area of interest. *Proceedings of Medicine Meets Virtual Reality '98,* 246–251.

Patrick, E., Cosgrove, D., Slavkovic, A., Rode, J. A., Verrtti, T., & Chiselko, G. (2000). Using a large projection screen as an alternative to head-mounted displays for virtual environments. *Proceedings of CHI 2000,* 478–485.

Poston, T., & Serra, L. (1994). The virtual workbench: Dextrous VR. *Proceedings of Virtual Reality Software and Technology (VRST '94),* 111–121.

Poupyrev, I., Billinghurst, M., Weghorst, S., & Ichikawa, T. (1996). Go-Go interaction technique: Non-linear mapping for direct manipulation in VR. *Proceedings of UIST '96,* 79–80.

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (Eds). (1992). *Numerical Recipes in C—The art of scientific computing* (2nd ed.) (pp. 412–420). Cambridge, UK: Cambridge University Press.

Raskar, R., Welch, G., Cutts, M., Lake, A., Stesin, L., & Fuchs, H. (1998). The office of the future: A unified approach to image-based modeling and spatially immersive displays. *Computer Graphics, Proceedings of SIGGRAPH '98,* 179–188.

Raskar, R., Welch, G., & Chen, W-C. (1999). Table-top spatially-augmented reality: Bringing physical models to life with projected imagery. *Proceedings of Second International Workshop on Augmented Reality (IWAR '99),* 64–71.

Raskar, R., Welch, G., & Fuchs, H. (1998). Spatially augmented reality. *Proceedings of IEEE Workshop on Augmented Reality (IWAR '98),* 63–72.

Rolland, J. P., & Hopkins, T. (1993). A method of computational correction for optical distortion in head-mounted displays. (Tech. Rep. No. TR93-045). UNC Chapel Hill, Department of Computer Science.

Rolland, J., Rich, H., & Fuchs, H. (1994). A comparison of optical and video see-through head-mounted displays. *Proceedings of SPIE: Telemanipulator and Telepresence Technologies,* 293–307.

Schikore, D. R., Fischer, R. A., Frank, R., Gaunt, R., & Whitlock, B. (2000). High-resolution multiprojector display walls. *IEEE Computer Graphics and Applications, 20*(4), 38–44.

Schmalstieg, D., Encarnação, L. M., & Szalavári, Zs. (1999). Using transparent props for interacting with the virtual table. *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics (I3DG '99),* 147–153.

Schmandt, C. (1983). Spatial input/display correspondence in a stereoscopic computer graphics workstation. *Computer Graphics, 17*(3), 253–261.

Sutherland, I. (1965). The ultimate display. *Proceedings of IFIP '65,* 506–508.

Van de Pol, R., Ribarsky, W., & Hodges, L. (1999). Interaction techniques on the virtual workbench. *Proceedings of Virtual Environments '99,* 157–168.

Watson, B., & Hodges, L. (1995). Using texture maps to correct for optical distortion in head-mounted displays. *Proceedings of IEEE VRAIS '95,* 172–178.

Wiegand, T. E., von Schloerb, D. W., & Sachtler, W. L. (1999). Virtual workbench: near-field virtual environment system with applications. *Presence: Teleoperators and Virtual Environments, 8*(5), 492–519.

Whitaker, R., Crampton, C., Breen, D., Tuceryan, M., & Rose, E. (1995). Object calibration for augmented reality. *Computer and Graphics Forum (Proceedings of EUROGRAPHICS '95), 14*(3), 15–27.

Wohlfahrter, W., Encarnação, L. M., & Schmalstieg, D. (2000). Interactive volume exploration on the StudyDesk. [CD-Rom]. *Proceedings of the Fourth International Immersive Projection Technology Workshop.*