# Real Mirrors Reflecting Virtual Worlds

Oliver Bimber
*Fraunhofer Institute for*
*Computer Graphics*
*Joachim-Jungius-Strasse 11,*
*18059 Rostock, Germany*
obimber@egd.igd.fhg.de

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

Dieter Schmalstieg
*Vienna University*
*of Technology*
*Karlsplatz 13/186/2,*
*A-1040 Vienna, Austria*
dieter@cg.tuwien.ac.at

## Abstract

*This paper introduces the idea of using real mirrors in combination with rear-projection systems for the purpose of interacting with and navigating through the displayed information. Subsequently a derived application is described. For this, we use a hand-held planar mirror and address two fundamental problems of applying head tracking with rear-projection planes: the limited viewing volume of these environments and their incapability of simultaneously supporting multiple observers. Furthermore, we describe the possibility of combining a reflective pad with a transparent one, thus introducing a complementary tool for interaction and navigation.*

## 1. Motivation

Head tracking represents one of the most common and most intuitive methods for navigating within immersive or semi-immersive virtual environments.

Rear-projection planes are widely employed in industry and the R&D community in the form of virtual tables [4] or responsive workbenches [10], virtual walls or powerwalls [14], or even surround-screen projection systems or CAVEs [7]. Applying head tracking while working with such devices might, however, lead to an unnatural clipping of large or extreme-located objects, of which parts are cut off by the edges of the viewing frustum (Figure 1). This situation destroys the sense of immersion into the virtual scene and represents one fundamental problem of these environments. Overcoming this problem [17] requires panning and scaling techniques (triggered by pinch gestures) to transform the projected scene to a manageable size. A synchronous transformation of the scene to a size suitable for a continuously changing viewpoint of the observer, however, would be difficult to achieve using these techniques.

In addition, using a combination of rear-projection systems and shutter-glasses makes head tracking, which simultaneously supports perspectively correct stereoscopic projection for multiple users, a technologically difficult and computationally intensive task. To our knowledge, this has not been achieved for more than two users [2], mainly due to hardware restrictions of current display technology. An alternative to a simultaneous support of multiple observers is to trade tracking between the collaborators, as stated in [17].

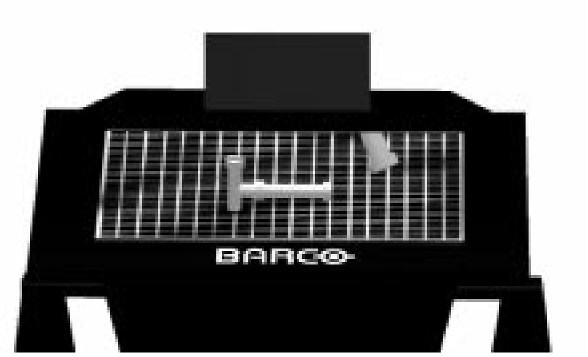

Figure 1. Head Tracking—unnatural clipping of objects.

To address these problems we propose a navigation method[1] that is complementary to single-user head tracking. It applies a planar mirror that can be used to increase the perceived viewing volume of the environment, and allows multiple observers to simultaneously gain a (perspectively and stereoscopically) less distorted impression of the projected scene than by applying regular head tracking.

With the thought in mind that nothing other than the two-dimensional images that are projected onto the projection

---

[1] Subsequently, we apply the term *navigation* using the mirror to refer to a kind of camera-controlled mechanism, rather than the motion of the observer, in terms of gaining a new perspective of the scene.

plane are physically reflected by the mirror, several interesting questions emerge. For instance: What happens if one observes the reflection of a stereoscopically projected virtual scene in a mirror? What has to be done to perceive the virtual reflection of the scene as three-dimensionally, stereoscopically and perspectively correct? Does the virtual reflection follow the same physical principles as the reflection of real objects?

Besides navigation, a multitude of interaction possibilities can be derived by introducing the mirror device. For example, combining the mirror device with a transparent pad leads to a particularly powerful tool that offers the application of everyday items (e.g. a pad and a pen). Thus, the fusion of the tool with the virtual world, as well as an intuitive handling, is supported.

The environment used for the following experiments is a Barco BARON [4] Virtual Table (virtual workbench, virtual plane). To gain the correct three-dimensional impression of the projected scene, we have applied head tracking and stereoscopic viewing in combination with shutter-glasses (e.g., Stereographics' CrystalEyes [15] or NuVision3D's 60GX [11]).

## 2. Reflecting virtual worlds with real mirrors

Planar mirrors enable us to perceive the reflection of stereoscopically projected virtual scenes three-dimensionally:

Instead of computing the stereo images based on the positions of the user's physical eyes (as it is usually done for head tracking), the corresponding reflection of the user's eyes within the reflection space (i.e. the space behind the mirror plane) has to be employed.

Because of the symmetry between the real world and its reflected image, the physical eyes perceive the same perspective by looking from the physical space through the mirror plane into the reflection space, as the reflected eyes perceive by looking from the reflection space through the mirror plane into the physical space (Figure 2).

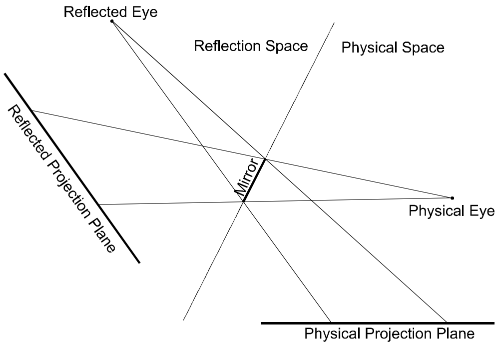

Figure 2. Perspective symmetry.

Thus, by applying the reflected eyes, the projection of the scene appears as a reflection in the mirror and will be perceived three-dimensionally by the user's physical eyes with the correct mirrored perspective. Thereby, the real mirror reflects the virtual world and, in doing so, follows the same physical principles as in the real world.

If the mirror plane is represented as:

$$f(x,y,z) = ax + by + cz + d = 0, \qquad (1)$$

with its normal vector $\vec{N} = [a,b,c]$ (2)

then the reflection of a point (in physical space coordinates) can be calculated as follows:

$$\vec{P}' = \vec{P} - \frac{2}{(\vec{N}^2)}(\vec{N}\vec{P} + d)\vec{N}, \qquad (3)$$

where $\vec{P}$ is the physical point and $\vec{P}'$ its reflection. To make use of the binocular parallax, the reflections of both eyes have to be determined. In contrast to head tracking, the positions of the reflected eyes, rather than the physical eyes, are used to compute the stereo images – even though the same algorithm is employed.

We can apply the reflection theorem to compute a vector's reflector:

$$\vec{L}' = 2(\vec{N}\vec{L})\vec{N} - \vec{L}, \qquad (4)$$

where $\vec{L}'$ is the reflector of $\vec{L}$.

If $\vec{E}$ is the user's generalized physical eye position and $\vec{X}$ a visible point on the mirror plane, then

$$\vec{L} = \frac{\vec{E} - \vec{X}}{|\vec{E} - \vec{X}|}. \qquad (5)$$

Hence, we can compute the visible points that are projected onto the projection plane ($g(x,y,z) = 0$) and are reflected by the mirror plane ($f(x,y,z) = 0$) as follows:

$$R = \{\vec{Y} | \vec{Y} = \vec{X} + \lambda\vec{L}, g(\vec{Y}) = 0, f(\vec{X}) = 0\}, \qquad (6)$$

where $\vec{X}$ is the point on the mirror plane that is visible to the user, and $\vec{Y}$ is the point on the projection plane that is reflected towards the user at $\vec{X}$.

# 3. Reflective pad

## 3.1. Increasing the viewing volume

A hand-held 6DOF-tracked mirror is used to complement single-user head tracking. We can employ (3) to calculate the positions of the user's reflected eyes. Instead of using the physical eye positions, their reflections are used to compute the two stereo images that enable the user to perceive a stereoscopically and perspectively correct reflection of the projected scene in the mirror. Navigation is supported by reorienting and moving the mirror, or by looking at it from a different point of view. This allows for taking up *difficult-to-reach* inspection points[2], so that normally clipped scenes can still be observed (Figure 3a, 3b).

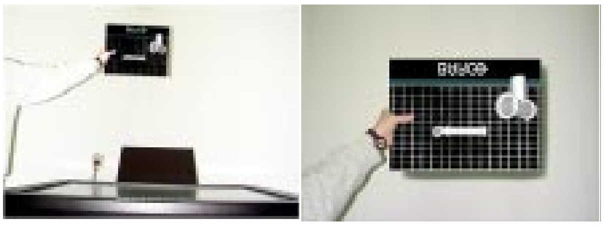

Figure 3a, 3b. Mirror tracking—difficult-to-reach inspection points.

In contrast to a manual transformation of the scene (as done in [17]), mirror tracking allows us to observe unnaturally clipped areas intuitively, even when the observer's viewpoint is changing continuously.

The mirror itself can also be used as a clipping plane that enables us to investigate the interiors of objects:

$$f(x, y, z) + \Delta , \qquad (7)$$

where $\Delta$ is the clipping plane offset. The offset is particularly useful for reflecting the intersection in the mirror.

In accordance with the term *head tracking*, we call this technique *mirror tracking*. In terms of providing both navigation possibilities in a complementary manor, the user can switch between mirror tracking and head tracking at will (for example, by picking up and laying down the mirror).

Note that mirror tracking is not a substitute for head tracking, but instead complements it.

## 3.2. Supporting multiple observers

By presuming that the physical and the reflected eyes are always located at a constant position relative to the mirror plane (e.g. on the mirror plane), we are able to navigate trough the scene without applying head tracking, just by tracking the mirror. Even though this approximation does not result in a correct mirrored perspective, it does support multiple observers in navigating through the scene by moving the mirror. The perceived perspective, however, is not the one of a single observer (as it would be the case with head tracking), but it can be thought of as the perspective two stereo-cameras would capture from the eye positions that are constant relative to the mirror-plane. Everyone looking at the mirror can then observe this perspective – although (due to the individually visible reflections of the projection plane) a slightly different portion of the scene is visible to each observer (Figure 5). This technique can be better illustrated by comparing it with a portable camera-display combination as shown in Figure 4.

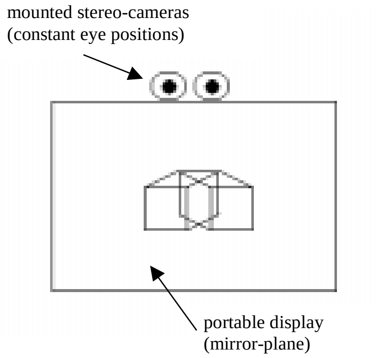

Figure 4. Illustration—portable camera-display combination.

To perceive an undistorted stereoscopic impression, the participants would have to look at the mirror roughly orthogonal (matching the constant eye positions); we found, however, that a satisfactory stereoscopic impression can still be perceived by taking up obtuse-angled lines of vision with relatively small mirrors.

Techniques that compute the best average viewpoint (BAV) for a group of people create similar problems. The closer a single observer is to the average viewpoint, the less distorted is his view; the more compact the group is clustered, the lower the stereoscopic and perspective error for every single observer. However, in contrast to BAV techniques, where large individual stereoscopic and perspective distortions can appear if the group is scattered, our approach offers an approximated correct collective perspective to the individually visible portion of the scene.

---

[2] Furthermore, the mirror (eventually in combination with other, fixed installed mirrors) offers the ability to distribute an infrared-signal (commonly used to synchronize the shutter-glasses with the rendering process) into different directions. This can prevent the signal from being interrupted, even if a continuous line of sight from the receiver to the sender is not given, as long as its reflection is visible in any of the mirrors.

Five observations can be made from our approach:

- The *collective perspective* of the images that are visible in the mirror is always approximately correct for all observers (although the perceived images slightly differ for every observer), since it depends only on the mirror (camera-display paradigm) and not on a single user.
- An individual *stereoscopic error* exists for every observer. However, relatively small mirrors force the observers to cluster together in order to look simultaneously at the mirror-plane. As with BAV techniques, this keeps the individual stereoscopic error low.
- Coordination problems that arise out of simultaneous navigation with the whole group are handled by reorienting the mirror instead of the moving the group. This keeps the group clustered and the individual stereoscopic error remains low.
- Navigation control can be intuitively delegated to another person by simply handing over the mirror.
- As it is the case for single users, mirror tracking also enables a group to increase their viewing volume in the environment.

Figure 5. Navigation with multiple observers.

### 3.3. Interacting with the reflection space

Our mirror-interaction paradigm supports indirect manipulative metaphors using ray pointers (e.g. ray casting, gaze-directed interaction and pointing [17]): Instead of using the original pointing selector, the ray pointer's reflector must be applied to identify objects or locations in the physical space, while pointing at their reflection behind the mirror plane (Figure 6).

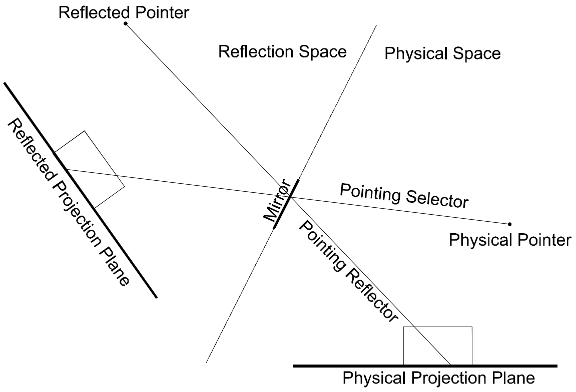

Figure 6. Pointing interaction with the reflection space.

To compute the pointing reflector we can apply (3) or (4). This simple illustration represents the basis for all ray-pointing interactions with the reflection space. Note that a direct manipulation (such as virtual hands, direct picking [17]) of the reflections is not possible because of the physical constraints of the mirror. Indirect interaction metaphors, however, require additional pointing devices (e.g. a tracked pen or a tracked forearm (Figure 5), etc.).

## 4. Transparent-reflective pad

Inspired by the work of Szalavári and Gervautz [16] who use opaque pen and pad props in combination with see-through head-mounted displays, Schmalstieg, Encarnação, and Szalavári [13] apply a tracked hand-held Plexiglas pad and a pen to support multiple two-handed interaction techniques. The transparent pad is augmented with 3D graphics from the Virtual Table's display and offers a wide variety of interaction possibilities. It can serve, for instance, as a palette for tools, controls and objects as well as a window-like see-through interface.

Since we have found that the application of 6DOF-tracked, hand-held transparent and reflective pads in combination with rear-projection are complementary to each other in application possibilities and interaction range, we propose their combination.

Our approach uses a Plexiglas pad that is simultaneously transparent and reflective. To achieve this, we use a special semitransparent foil (such as 3M's Scotchprint P18 [1]) that is normally used on windows to control sunlight. The foil's surface is identical on both sides and either reflects or transmits light, depending on its orientation to the light source (i.e. the projection plane in our case) (Figure 7a, 7b).

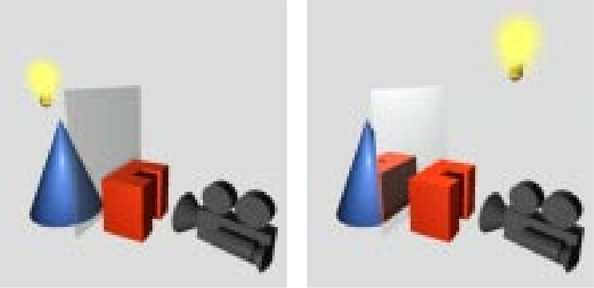

Figure 7a, 7b. Functioning of the foil (transparent / reflective mode).

The integrated sun protection, however, blocks a fraction of the transmitted light. Thus, graphics that users observe through the pad appear to be darker than graphics that they see without the pad. We can reduce the resulting color difference by setting up adequate light sources with respect to the pad that illuminate the affected objects and allow them to be projected more brightly before they are transmitted through the pad.

### 4.1. Active mode selection

The two modes of the transparent-reflective (transflective) pad are complementary and can be switched intuitively by holding it with respect to the specific task (i.e. see-through (Figure 12) or mirror (Figure 10)).

Whether the transparent or the reflective mode is active can be determined as follows (Figure 8):

If $\vec{E}$ is the user's generalized physical eye position and $\vec{P}_m$ a point on the pad plane that is projected onto the projection plane, then (with respect to (1)) the transparent mode is active, if

$$sign(f(\vec{E})) \neq sign(f(\vec{P}_m)) \qquad (8)$$

(i.e., the points are on the opposite side of the pad plane). The reflective mode is active, if

$$sign(f(\vec{E})) = sign(f(\vec{P}_m)) \qquad (9)$$

(i.e., the points are on the same side of the pad plane).

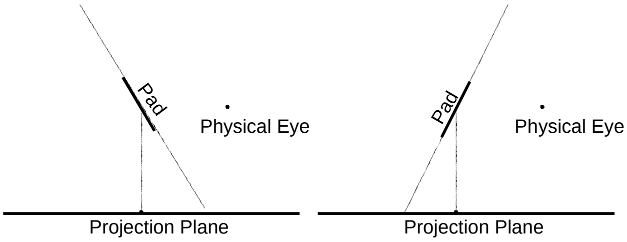

Figure 8. Division by the pad plane.

Simply comparing whether the two points are divided by the pad plane, however, is insufficient since the reflective mode is commonly favored while holding the pad orthogonal to the projection plane, and tracking distortion makes an exact detection of the active mode impossible.

Therefore we also evaluate the solid angle between the two normal vectors of the planes and assign mode specific function zones (Figure 9):

If the solid angle between $\vec{N}$ (normal vector of the pad) and $\vec{Z}$ (normal vector of the projection plane) is defined as follows:

$$c = \cos(\alpha) = \frac{\vec{N}\vec{Z}}{|\vec{N}||\vec{Z}|}, \qquad (10)$$

then we propose the following algorithm to determine the active mode (with respect to (8) and (9)):

$$M = \begin{cases} reflective, & ((c < T_1) AND (c > T_2)) \quad OR \\ & (sign(f(\vec{E})) = sign(f(\vec{P}_m))) \\ transparent, & ((c \geq T_1) OR (c \leq T_2)) \quad AND \\ & (sign(f(\vec{E})) \neq sign(f(\vec{P}_m))) \end{cases} \qquad (11)$$

$T_1$ and $T_2$ are specific threshold values that define the function zones. We found that $T_1 = 0.5$ and $T_2 = -0.5$ are suitable values that support an intuitive switching between the modes.

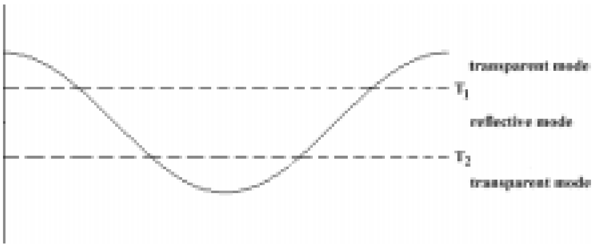

Figure 9. Function zones.

### 4.2. Functionality

The modes' major functions are:

- interaction in the transparent mode (window-controls, through-the-plane tools, magic-lenses [5,18], etc. as stated in [13])
- complementary navigation in the reflective mode (difficult-to-reach viewing, multiple observer viewing, indirect interaction, clipping-plane-in-hand, etc.)

Even though the modes are complementary in most cases, a certain overlap exists. On the one hand, a two-handed indirect interaction in combination with a tracked pen would also be supported in the reflective mode (interaction with the reflection space), and seems to be an interesting possibility for interactions from *difficult-to-reach* positions (e.g. in the inside of objects, etc.). Furthermore, the application of different window-tools (such as the ones used by Schmalstieg et al. [13]) is also imaginable in the reflective mode. On the other hand, navigation (clipping-plane-in-hand, etc.) can also be realized in the transparent mode. Although this method requires overlapping application possibilities, it is still complementary in the interaction range.

Besides the clipping plane defined in (7), an additional one can be found that might be set up to support rear-plane clipping from both sides of the pad plane, in both the transparent and the reflective mode. The proper clipping plane must be activated with respect to both the active mode and the side of the pad plane on which the user is located (Table 1).

Table 1. Active clipping plane selection.

| **Transparent mode** | | **Reflective mode** | |
|---|---|---|---|
| $f(\bar{E}) \geq 0$ | $f(\bar{E}) < 0$ | $f(\bar{E}) \geq 0$ | $f(\bar{E}) < 0$ |
| $-f(x,y,z) - \Delta$ | $f(x,y,z) + \Delta$ | $f(x,y,z) + \Delta$ | $-f(x,y,z) - \Delta$ |

Furthermore, we can overload the transparent and the reflective modes with a multitude of different functionalities. As stated in [13], the user can activate different modes that are supported by the two different sides of the transparent pad. Thus, for instance, window-controls (such as buttons, sliders, etc.) are offered on one side, and through-the-plane tools (such as magic lenses [5,18], etc.) are provided on the other side. The user can switch between them anytime by turning over the pad.

The transflective pad, however, offers users the ability to make use of four sides (two transparent and two reflective ones), and supports a variety of application-specific interaction potentialities. Encarnação, Bimber, Schmalstieg and Chandler [8], for instance, describe the possibility of using two-dimensional free-hand sketches to control and create objects by sketching on the pad.

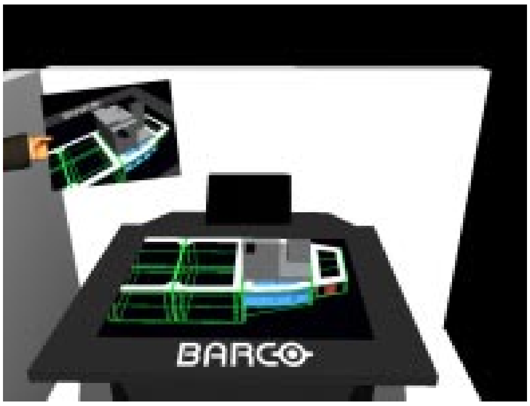

Figure 10. Reflective mode—head tracking (projection plane) and difficult-to-reach viewing (mirror).

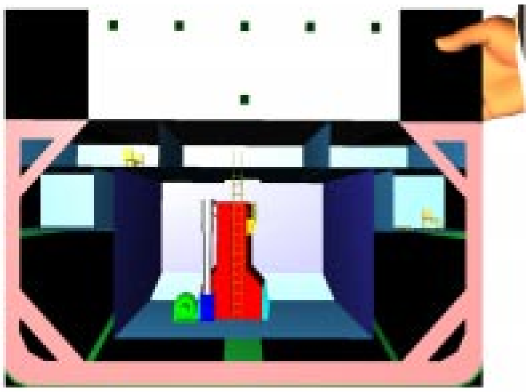

Figure 11. Clipping-plane-in-hand (transparent / reflective mode).

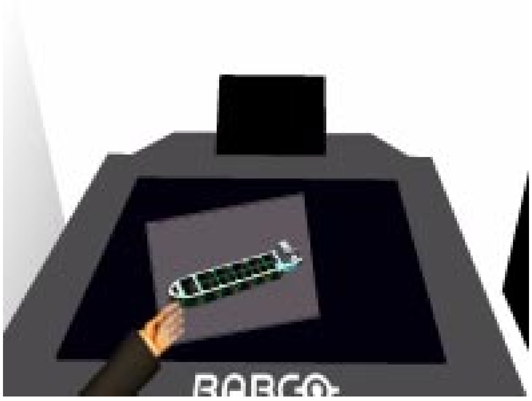

Figure 12. Transparent mode—object palette.

## 5. Informal User Study

A transflective pad of the size 15” x 11” has been employed in combination with the ship scenario to investigate its usability (Fig. 10, 11, 12). Eight unbiased participants, who had never worked with or even seen the system setup were asked to explore the interior of the ship using the pad in both the transparent and the reflective mode. The pad’s functionality was explained to them and the task specified: Within two minutes, each participant had to “clip” inside the ship, navigate through its engine-room by moving the pad, and observe the asymmetric engine from all directions. After the time had expired, they were asked to sketch the engine from memory.

The goal was to find out whether or not the two different types of view (the real-world type of view supported by the transparent mode, and the reflection type of view supported by the reflective mode) can be cognitively combined to achieve the correct three-dimensional plasticity, and whether the handling of the pad is intuitive. The fact that all of the participants were able to sketch the engine rather accurately—even with most of its details at the right places—leads to the following assumption (which was verbally confirmed by the participants): Using the pad in the reflective mode while knowing that one is looking into a mirror, makes one expect to see a reflection in it. The reverse prediction can be made for the transparent mode. Thus users can cognitively combine the two types of views correctly and perceive a three-dimensional picture of the scene.

Even though the participants found the handling of the pad and the switching between transparent and reflective mode intuitive, some of them experienced difficulties with the clipping-plane-in-hand concept. This, however, can be attributed to the distortion of our electromagnetic tracking-device (Polhemus Isotrack), which increases if one moves the pad closer to the engine and consequently (in our case) further away from the transmitter.

In addition to this initial informal user test, CRCG [6] is preparing a more comprehensive usability study that will help to evaluate the proposed techniques more thoroughly.

## 6. Conclusions

We have introduced the idea of using real mirrors to reflect virtual worlds and addressed two fundamental problems of applying head tracking in combination with rear-projection planes:

First, we have shown how to use a hand-held tracked mirror to increase the limited viewing volume of these environments; second, we have given an approximation to the approach of supporting multiple observers in navigating through the scene.

Moving the mirror to navigate[3] through an information space that is essentially larger than the display device (i.e. the mirror) supports a visual perception phenomenon that is known as *Parks Effect* [12]. That is, moving a scene on a fixed display is not the same as moving a display over a stationary scene because of the persistence of the image on the viewer’s retina. Thus, if the display can be moved, the effective size of the virtual display can be larger than its physical size, and a larger image of the scene can be left on the retina. This effect can also be observed with other handheld devices, such as Fitzmaurice’s Chameleon [9] or the Art+Com VR Display [3], etc.

Since the transparent and the reflective pad complement each other, we proposed their combination to build a more powerful interaction and navigation tool—the transflective pad.

We found three general advantages of the transflective pad:

- Users are able to handle the pad intuitively because of their familiarity with everyday physical items, such as trays, mirrors, pads and pens, or even window controls (i.e. buttons, sliders, etc.).
- When applied with large rear-projection systems, the transflective pad provides a cost-efficient and easily applicable solution to combine the presented techniques with traditional (semi-)immersive VR tasks (in contrast to other approaches that make use of additional electronic tools, such as the Chameleon [9], etc.).
- The transflective pad represents a handheld device that, in combination with rear-projection systems, offers stereoscopic viewing. To our knowledge, this cannot be realized with today’s portable devices that, for instance, apply Plasma LCD screens. This is due to their low update rates.

---

[3] Single-user or multiple-user navigation.

We believe that the key idea—real mirrors reflecting a virtual world in all its variants—has the potential to offer many new R&D opportunities in the area of human-computer interaction. This is evidenced by the fact that several of the people who observed the development process and the creation of this paper picked up the idea to derive further applications. We hope that many others will follow.

## References

[1] 3M, CORP. P-18 is the most efficient Scotchtint Sun Control film. *URL: http://www.mmm.com/market/construction/-html/products/product73_p.html*, 1999.

[2] M. Agrawala, A. C. Beers, B. Fröhlich, I. McDowall P. Hanrahan, and M. Bolas (1997). The Two-User Responsive Workbench: Support for Collaboration Through Individual Views of a Shared Space. *Computer Graphics Proceedings, Annual Conference Series SIGGRAPH '97*, pp. 327-332, 1997.

[3] Art+Com, AG, *URL: http://www.artcom.de/projects/vrf/ welcome.en*, 1996.

[4] BARCO, INC., BARON, *URL: http://www.barco.com/-projecti/products/bsp/baron.html, 1997*.

[5] E. Bier, M. Stone, K. Pier, W. Buxton, and T. DeRose. Toolglass and Magic Lenses: The See-through Interface. *Proceedings of SIGGRAPH'93*, pp. 73-80, 1993.

[6] CRCG, INC., Center for Research in Computer Graphics. *URL: http://www.crcg.edu*, 1998.

[7] C. Cruz-Neira, D. Sandin, and T. DeFanti. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. *Proceedings of SIGGRAPH'93*, pp. 135-142, 1993.

[8] L.M. Encarnação, O. Bimber, D. Schmalstieg, S.D. Chandler. A Translucent Sketchpad for the Virtual Table Exploring Motion-based Gesture Recognition. *Computer Graphics Forum (Proceedings of EUROGRAPHICS'99)*, vol. 19(3). NCC Blackwell, 1999.

[9] G.W. Fitzmaurice. Situated Information Spaces and Spatially Aware Palmtop Computer. CACM, vol. 35(7), pp. 38-49, 1993.

[10] W. Krüger, C. Bohn, B. Fröhlich, H. Schüth, W. Strauss, and G. Wesche. The Responsive Workbench: A Virtual Work Environment. *IEEE Compute*r, vol. 28(7), pp. 42-48, 1995.

[11] NuVision3D (1997). 60GX,
*URL: http://www.nuvision3d.com/the60gx.html*, 1997.

[12] T.E. Parks. Post Retinal Visual Storage. *American Journal of Psychology*, vol. 78, pp. 145-147, 1965.

[13] D. Schmalstieg, L.M. Encarnação, Zs. Szalavári. Using Transparent Props for Interaction with the Virtual Table. *ACM SIGGRAPH Symposium on Interactive 3D Graphics (I3DG'99)*, Atlanta, GA. ACM Press, pp. 147-153, 1999.

[14] Silicon Graphics, Inc., Powerwall,
*URL: http://www.stereographics.com/html/-body_powerwall.html*, 1997.

[15] StereoGraphics Corp., CrystalEyes,
*URL: http://www.stereographics.com/html/-body_crystaleyes.html*, 1998.

[16] Zs. Szalavári, M. Gervautz. The Personal Interaction Panel – A Two Handed Interface for Augmented Reality. *Computer Graphics Forum (Proceedings of EUROGRAPHICS'97)*, vol. 16(3), pp. 335-346. NCC Blackwell, 1997.

[17] R. van de Pol, W. Ribarsky, L. Hodges. Interaction Techniques on the Virtual Workbench. *Proceedings of Virtual Environments'99*, pp. 157-168, 1999.

[18] J. Viega, M. Conway, G. Williams, and R. Pausch. 3D Magic Lenses. *Proceedings of ACM USIT'96*, pp. 51-58, ACM Press, 1996.