Vermeer:Direct Interaction with a 360° Viewable 3D Display

Alex Butler[1], Otmar Hilliges[1], Shahram Izadi[1], Steve Hodges[1], David Molyneaux[1,2], David Kim[1,3], Danny Kong[1]

[1]Microsoft Research, 7 J J Thomson Avenue, Cambridge, CB3 0FB, UK

[2]Computing and Communications, Lancaster University, Lancaster, LA1 4WA, UK

[3]Computing Sciences, Newcastle University, Newcastle, NE1 7RU, UK

{dab,otmarh,shahrami,shodges,a-davmo,b-davidk}@microsoft.com, d.l.kong.01@cantab.net

Figure 1: Vermeer allows direct interaction within the volume of a 360° viewable 3D display. Far left: 3D model rendered using our 3D display. Middle sequence: user places finger in the volume to interact directly with 3D model. Note how the perspective of the 3D model changes, depending on the viewer position (without requiring any head tracking or specialized eyewear). Far right: the user interacts with a spherical display using the same optical arrangement.

ABSTRACT

We present Vermeer, a novel interactive 360° viewable 3D display. Like prior systems in this area, Vermeer provides viewpoint-corrected, stereoscopic 3D graphics to simultaneous users, 360° around the display, without the need for eyewear or other user instrumentation. Our goal is to overcome an issue inherent in these prior systems which – typically due to moving parts – restrict interactions to *outside* the display volume. Our system leverages a known optical illusion to demonstrate, for the first time, how users can reach into and directly touch 3D objects *inside* the display volume. Vermeer is intended to be a new enabling technology for interaction, and we therefore describe our hardware implementation in full, focusing on the challenges of combining this optical configuration with an existing approach for creating a 360° viewable 3D display. Initially we demonstrate direct in-volume interaction by sensing user input with a Kinect camera placed above the display. However, by exploiting the properties of the optical configuration, we also demonstrate novel prototypes for fully integrated input sensing alongside simultaneous display. We conclude by discussing limitations, implications for interaction, and ideas for future work.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General terms: Design, Human Factors, Algorithms

Keywords: 360° Viewable 3D Display, In-Volume Interaction, Input Sensing, Optical Illusion, Tabletop.

INTRODUCTION

There is a growing body of research into 360° viewable 3D displays [5,8,10,13,19,20,29,32]. Such systems provide viewpoint-corrected, stereoscopic 3D graphics to simultaneous users, 360° around the display, without the need for eyewear or other user instrumentation. These technologies also tend to provide depth cues such as vergence, correct eye-accommodation and horizontal motion parallax – the latter without the need for head tracking – which provide a perception of depth even for those without stereopsis. They have been shown to improve viewer perception of 3D when compared with stereo-only displays [14].

Most implementations of such displays restrict the user from interacting within the display volume. Indeed, users are typically separated from the displayed 3D objects by a glass or acrylic shield (due to moving parts [10,19] or hazardous [20] or fragile optical elements [29]). Current user interface research within this domain has therefore focused on enabling *outside* of the volume interaction [15,16].

We introduce a novel interactive display system called Vermeer, which allows users to reach into and directly touch 3D objects within the display volume. The system builds on the 3D display research of [19] and therefore has similar benefits for viewing 3D graphics. However, Vermeer also supports *direct* interactions *within* the display volume, where users can now place their fingers (or indeed any object) to interact, as shown in Figure 1. This has not been demonstrated previously by other 360° viewable 3D displays because of their physical constraints.

To achieve this, the system exploits a known optical illusion using two parabolic mirrors [11,24]. This optical configuration allows a 360° viewable 3D display to be *reimaged* [1] such that the graphics appear to be floating in a user-accessible region above the unit. We first show how a Kinect camera, placed above the display, can be used to detect fingers inside the display volume, and demonstrate simple proof-of-concept in-volume interactions. Next we leverage the unique properties of the parabolic mirror configuration and demonstrate novel prototypes that support fully integrated input sensing alongside simultaneous display. We conclude by discussing limitations, implications for interaction, and ideas for future work.

RELATED WORK

Detailed surveys of 3D display technologies are provided by [5,6,9,17,19,30]. Our aim is to enable direct interaction with a particular class of 3D display. Such technologies are referred to as volumetric [5,10,13,29], multiview [8,12,22] or 360° light field displays [12,19,32] within the literature. Whilst there are clearly distinctions between these types of displays, for our purposes we characterize these displays as exhibiting the following properties:

- *A 360° horizontal field-of-view and omnidirectional viewing,* where users receive differing views of the 3D scene depending on their position around the display.
- *Horizontal parallax without head tracking.*
- *Simultaneous support for multiple viewers.*
- *Stereoscopic imagery* without the need for eyewear or other user instrumentation.
- *Non-planar characteristics* where the display is bound to an entire 3D volume rather than a single plane.

Such systems clearly contrast with *stereo* only displays which require specialized glasses and support a single stereoscopic depth cue. They also contrast with *autostereo* only displays (e.g. based on lenticular arrays or parallax barriers [9,17]), which remove the need for glasses, but again only deliver stereoscopic-cues and typically support limited viewing positions.

Many different techniques have been proposed to enable such 360° viewable 3D displays (see [5,13,19] for a detailed overview). These include swept volume approaches which use moving parts, such as rotating display elements to render slices of volumetric data at high speed, and form full 3D imagery through persistence of vision [6,13]. An extension to these approaches are systems that rotate a high speed LCD [22], LED [12,32] or projected display [19], and restrict viewing angles to create perspective-corrected renderings of a 3D scene in different directions. Solid state displays that render slices of 3D data onto a stack of dynamically switching diffuser planes [29] have been proposed as well as more unconventional approaches that use multiple lasers to stimulate points or regions within photonically excitable mediums [20].

Whilst these systems offer steps towards realistic 3D display, they do not naturally allow a user to interact directly within the display volume, either because of rapidly moving mechanical elements, physical encasing of optical components, or hazardous or fragile materials. Typically a transparent glass or acrylic dome or other barrier separates the physical display volume from the user.

Despite these limitations, previous work has reported different approaches to interact with virtual 3D content rendered by such displays from *outside* of the volume. In [16] the 6DOF pose of fingers are tracked using a Vicon motion capture system to support interactions outside and on the dome of an Actuality volumetric display [10]. [15] presents an evaluation of targeting tasks using a Vicon tracked stylus to point into the volumetric display from the outside.

Our goal is to remove this limitation on 360° viewable 3D displays, and allow direct interaction within the volume.

Aside from spatial 3D displays, there is also considerable literature on head-worn optical or video see-through 3D display systems (see [3] for a detailed overview). Whilst enabling greater user mobility and novel AR and VR scenarios, many head-worn displays have a number of drawbacks including small field-of-view, continually forcing the user to refocus between virtual graphics close to the eye and real-world objects far away (an issue for optical see though displays), inherent latency, and inaccuracies in tracking, which can result in sickness and fatigue during use [7,18]. Perhaps the biggest drawback of head-worn solutions is ergonomics; many available devices are still heavy and require tethered hardware to be worn.

Another approach is to present stereoscopic imagery to a limited number of users is to leverage a horizontal projection screen and tracked shutter glasses, as in [2]. In addition to head-worn glasses, the system also tracks styli for input. While less intrusive this approach still suffers from missing depth cues and possible distortions in depth perception [18]. Also it is non-trivial to scale such systems to many users due to frame rate limitations of projection and glasses hardware [2].

In contrast to these previous approaches, we wish to support even more lightweight 'walk up and use' tabletop scenarios, without any user instrumentation, eyewear or otherwise, or the use of tracked input devices.

Another means of co-locating 3D graphics and input without glasses is to use an optical combiner such as a beam-splitter [4,25,26,31]. In these arrangements 3D graphics produced on a CRT or LCD screen are reflected as a virtual image displayed to the user through a beam splitter. This co-locates the virtual reflected scene with the user's interaction space. Input is typically provided using a stylus [25,31] or force-feedback device [26].

While these setups provide many depth cues including full motion parallax when employing head-tracking they are inherently single user and cannot easily provide omnidirectional views of the scene. The virtual showcase [4] circumvents some of these limitations by utilizing several or a single curved beam-splitter. This provides a more omnidirectional viewing experience but fully encases the display preventing users from reaching into the volume.

VERMEER PROTOTYPE

In order to enable direct, unencumbered interaction with a 360˚ viewable 3D display we exploit a known optical illusion used in the commercially available Opti-Gone Mirage product [24]. When two equally-sized parabolic mirrors are placed in a 'clamshell' setup (facing each other with the center of each mirror coinciding with the focal point of the other) any object that can be placed in the central region of the lower mirror is *reimaged* so that an image of that object can be observed 'floating' just above the upper mirror, as if it were real. This well-known phenomenon (described fully in [1,11]), forms the basis of our display.

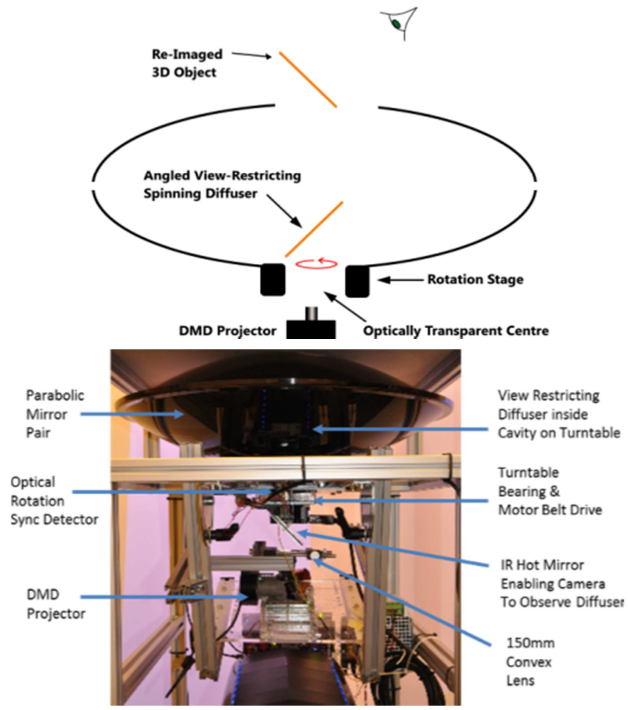

Figure 2: The Vermeer Display. Top: shows basic principle of operation. A 360° viewable 3D display based on spinning optics is placed at the bottom of a pair of parabolic mirrors. The display is reimaged above the mirrors, where it can be observed by multiple users who can then place their fingers inside the view volume. Bottom: hardware configuration showing main components.

The basic architecture of Vermeer is a 360° viewable 3D display sited between these two parabolic mirrors as depicted in Figure 2. Users can view content 360° around the reimaged display and can also touch anywhere within the perceived display volume to interact with the 3D content directly. Users must view the image from around 45° relative to the center axis above the top aperture in order to see into the mirror cavity.

Building this configuration is non-trivial, particularly as it requires us to build and extend the 360° viewable 3D display described in [19] to support both projection and imaging from below. The previous system used top-down projection onto an anisotropic mirror [19]. However, top-down projection is not feasible for our requirements, as fingers interacting within the volume will occlude the projection. In the remainder of this section we describe the salient parts of our hardware and software implementation.

We sourced the Opti-Gone Mirage Model 22 consumer product [24], which provides 56cm diameter parabolic mirrors with a 15cm diameter opening in both mirrors. This setup provides a cylindrical view volume of about 50 mm diameter and 55 mm height for the 3D display to appear. Inside this volume, a projection screen (35 mm wide and 70 mm high) is mounted at a 45° angle upon a rotating stage, with circular cut-out at its center. This is driven by belt using an Animatics SM2316D SmartMotor. This allows us to control the orientation of the projection screen without obstructing the optical path from the projector mounted underneath the diffuser.

A high speed DMD projector is utilized to display multiple renderings of the 3D scene at very high frame rates – the projection is synchronized with the SmartMotor to ensure that the correct view is displayed at the correct time (or rather when the diffuser is positioned in the corresponding orientation). The projector is mounted so that additional, optical elements can be positioned in the optical path. For example, lenses to focus the projection for a shorter throw distance. Additionally a hot mirror enables an infrared (IR) camera to sense user input (as described later).

View Restrictive Diffuser

A narrow view angle is required to present sufficiently different views to multiple users around the display and to ensure inter-ocular separation of images for stereopsis. Previous work [19] used a reflective mirror with anisotropic characteristics (to provide a narrow viewing angle in the horizontal direction) coupled with projection from above. Our current implementation uses a back-projection diffuser with a stack of off-the-shelf privacy films [27] on top (we use 7 in our current implementation). The stacked arrangement simulates deeper micro-louvres and generates the narrow viewing angle required.

This solution greatly minimizes the visibility of vertical dead banding, as apparent in other parallax barrier techniques such as macro scale venetian blinds. Our stacked arrangement can lead to small Moiré artifacts in the image although these are not as apparent when the diffuser is spinning. We have also recently sourced a custom single substrate *view control film* (from Shin-Etsu) with a narrower field-of-view and no visible Moiré. However, the stacked arrangement is a more readily available solution.

High-speed Projection and Rendering Pipeline

To display multiple viewpoints per full rotation of the spinning diffuser, we use a Texas Instruments DMD Discovery 0.7" XGA D4100 high speed projector. We target a horizontal view angle of ~1.875° such that a full 360° rotation is made up of 192 independent viewpoints. We use a DLi D4100 DVI interface daughterboard in combination with a custom FPGA to split each outgoing 24-bit DVI frame into 24 binary DMD images on the projector. Each single 24-bit RGB DVI frame transmitted from the PC to the projector therefore can contain 24 binary images each containing 24 different projections of the 3D scene. Our DVI connection runs at 120Hz, totaling 2880 binary projector frames per second, achieving a total refresh rate of 15Hz for the entire display.

Our focus is on the novel use of the mirrors for input so for our display we employ a much simpler 3D perspective rendering than that proposed by [19]. A virtual camera circles the scene looking down at a 45° angle to compensate for the diffuser mounted at the same angle. The virtual camera also rotates about its own up vector to correct for the rotation of the diffuser stage. A custom rendering technique writes 24 consecutive views into an off-screen render target. A final post-processing pass combines these views into a single 24-bit RGB image.

Synchronization

Synchronization between the mechanically rotating element and the projected image stream is critical to ensure that the 360° generated viewpoints remain consistent. The FPGA decoding incoming DVI frames was designed to detect a pattern embedded in the first 15 lines of the very first binary image of a 360˚ cycle. A 2μs pulse is generated at each start of a full rotation. A simple LED and photodiode optical detector generates a second pulse whenever the angled diffuser passes a reference (0˚) position. A TI MSP430 embedded microprocessor monitors the difference in time between the incoming pulses. A simple control loop adjusts the speed of rotation of the motor to ensure that the first frame of the series is projected at the diffuser's zero angle position.

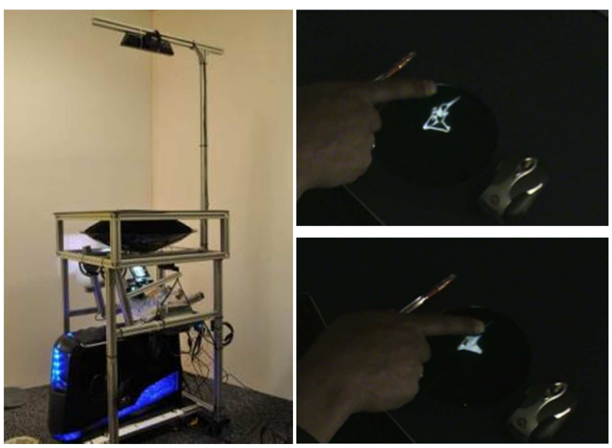

Figure 3 Left: The Vermeer display with Kinect depth sensor mounted above. Right: the user's finger is sensed within the volume using the Kinect sensor and causes the ship to switch from wireframe to dithered rendering.

IN-VOLUME INTERACTION

Figure 1 shows interactive graphics rendered on Vermeer as seen from different viewpoints around the setup. Note how the virtual 3D model shares the same perspective foreshortening with real objects in the scene such as the mouse and a user's finger. Also note how the position of the 3D model stays steady in relation to the fingertip illustrating correct spatial registration. The graphics are reimaged from the bottom of the parabolic mirror pair, allowing users to directly touch the 3D model.

To first explore the possibilities for interaction with dynamic 3D scenes we mounted a Kinect camera directly above the mirror pair (Figure 3, left).

We implemented a fingertip tracking technique using the depth data from the Kinect camera. Using the depth data, it is relatively straightforward to segment the user's hand from the background once it enters the volume, trace the contour of the user's hand and detect fingertips using a peak-and-valley image processing algorithm [23] at 30Hz (Figure 4). Detected fingertip positions are Kalman filtered and tracked over time to be used as spatial input for our demonstrators. Once fingertips are found, we average the depth values of a patch of pixels around the fingertip to infer its 3D position. To avoid depth distortions only pixels lying within the contour are considered.

Figure 4 left to right: depth map real-world converted; foreground extracted; contour traced and fingertips detected using peak-and-valley algorithm.

In the most simplistic scenario a user may place a finger within the display volume in order to animate a 3D character or change the rendering used (Figure 3, right). However, it is possible to implement more fine-grained interactions than this simplistic binary touch/no-touch scenario. The actual spatial location of the fingertip can be used to make better use of the 3D nature of our system. When a user touches the center of the 3D character the rendered model will start to walk following a pre-defined animation. In contrast, if the user touches the top of the model the animated character waves his arm and lowers it back once the finger retracts from the volume.

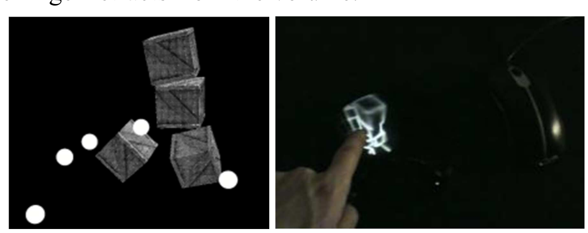

Figure 5 Left: On screen physics-enabled scene with dithered gray-scale rendering; white spheres serve as proxy for fingers and interact with virtual objects. Right: Scene as viewed by the user.

A final example demonstrates more fine-grained interactive capabilities of our system. Rather than interacting with pre-defined animations it is possible to use a physics simulation (in our case NVIDIA PhysX) to control the position of virtual objects via real-world concepts such as forces, collisions and friction. In this instance we use the 3D fingertip position data and model the users' fingers as spherical proxy objects within the simulation. This allows users to knock over stacks of virtual spheres or boxes by directly intersecting their fingers with the virtual objects inside the volume (see Figure 5).

INTEGRATED SENSING

So far we have demonstrated how the reimaged graphics of a 360° viewable 3D display can be directly interacted with using fingertips sensed using a Kinect camera. This arrangement, while enabling some interesting interaction possibilities, increases the complexity of the setup, and suffers from many of the issues of top-down camera systems, in particular additional bulk and occlusions. In this section, we demonstrate two novel prototypes that further leverage the optical properties of the mirrors to sense user input from *within* the display arrangement – making it more appealing as a self-contained tabletop form-factor.

Infrared Reimaging

One optical property of the parabolic mirrors, particularly interesting for finger sensing, is the ability to reimage light in the near-IR range. This can be confirmed by placing an IR diffusing object inside the mirror-pair and illuminating the object using IR LEDs. The example in Figure 6 left shows a dome (hemispherical object) being illuminated and reimaged in this way. Here a handheld IR camera observes a reimaged 3D IR object floating above the mirror.

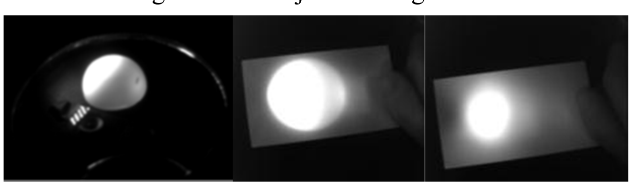

Figure 6 Left: A reimaged IR illuminated dome, as seen from a handheld IR camera above the mirrors. Middle and right: A diffuser sheet is intersected with this IR dome at different heights, and imaged from below.

One interesting aspect is that the floating IR object is *optically real*. We illustrate this in Figure 6 right, by moving a diffuser sheet (tracing paper) above the top mirror aperture, and using an IR camera to image this area from below. The figure clearly shows a circular spot where the diffuser intersects with a cross-section of the reimaged IR dome. The IR light is scattered downwards and imaged by the camera. Moreover, the size of the spot changes as the diffuser is moved up and down, intersecting with different slices of the reimaged 3D IR dome. This optical property can be leveraged for in-volume touch sensing. An arbitrary 3D surface, reimaged using IR, can float in the display area. Fingertips intersecting with this reimaged IR object will be clearly illuminated and detected with an IR camera.

Figure 7 (top left) shows our first novel prototype exploiting this optical property. Figure 7 (top right) shows the setup more clearly. A dome object is placed at the base of the mirror cavity. A video projector is used to render visible graphical output onto the dome. Graphics are pre-distorted to compensate for the non-planar projection screen. To sense user input, an IR camera (located at the base of the mirror pairs) images through the aperture in the top mirror half.

From above, in visible light, the user can see the reimaged dome object and projected graphics. However the IR camera below is directly imaging this same space (without observing the reimaged visible object). When the user's fingers are placed inside the display volume, they intersect the reimaged IR dome, and IR light will be scattered downwards, where the camera can image a bright reflection of the fingertip, as illustrated in Figure 7.

A user interacting with the system is thus presented with interactive full-color graphics floating above the mirror. When touching (or rather intersecting) the surface of the reimaged projection surface we can detect and track the resulting bright IR blobs in the camera image for user input. In our proof-of-concept implementation a map of the earth is projected, as shown in Figure 7, bottom.

A finger moving along the surface of the reimaged hemisphere spins the globe about an imagined axis perpendicular to the motion trajectory. It is worthwhile noting that only the section of the finger actually intersecting the reimaged sphere is IR illuminated so that fingers hovering above the sphere's surface can be distinguished from ones touching. This prototype demonstrates how we can sense intersections between physical objects above the mirror and IR illuminated objects inside. These objects do not need to be planar as demonstrated here.

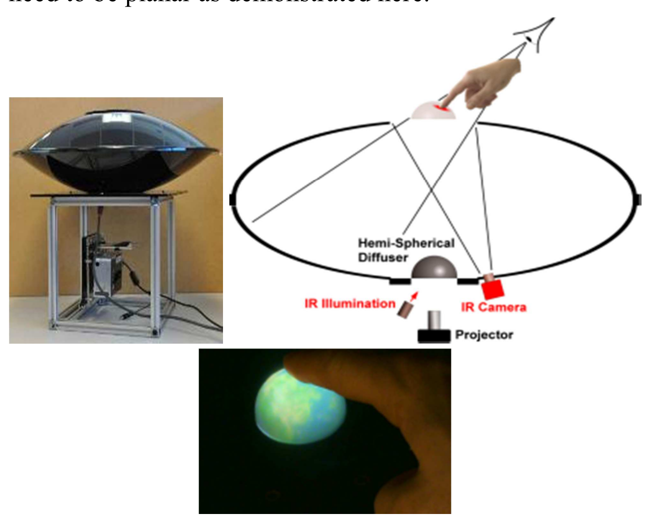

Figure 7: Top left: Prototype of integrated fingertip sensing. Top right: Schematic overview. A rear-projected hemispherical diffuse surface is also illuminated from below with IR. Bottom: A user spins the rendered map of the earth by dragging a finger across the surface of the reimaged globe.

Fully Integrated Touch Sensing and 3D Display

The previous prototype shows how the optical properties of the parabolic mirror pair extend to IR based fingertip sensing. A more well-known but rarely exploited property of the mirror arrangement is that the floating reimaged object can in fact be re-illuminated (e.g. using a small torch or laser) under certain circumstances [1, 28]. If a laser is shone onto the floating reimaged object, but aimed to ensure that the light actually passes down into the aperture of the upper mirror, the bright laser spot will appear on the reimaged object as if it were real (this is described in [28], page 28).

This demonstrates the basic reversibility of the optical setup, and highlights that light emitted by, or diffusely scattered off, any object above the top cavity can be reimaged downwards into the bottom cavity. We use both these properties of the optical arrangement to support fully integrated sensing in our 360° viewable 3D display.

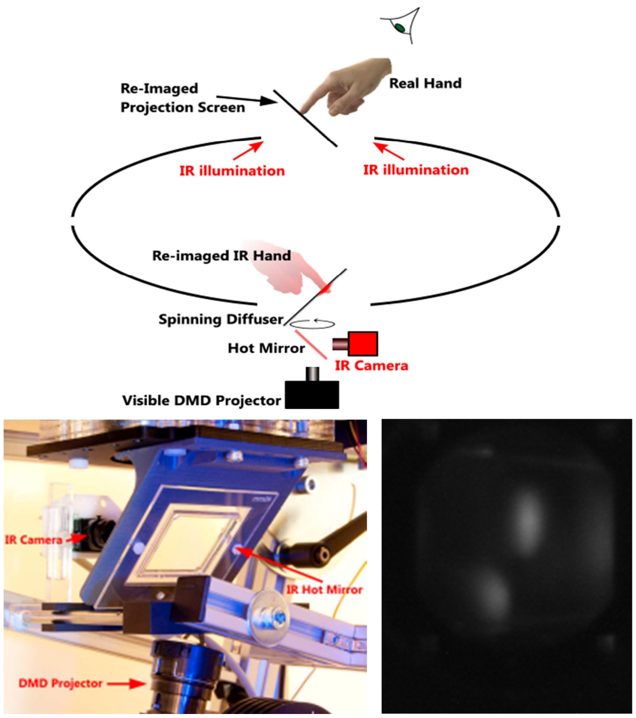

Figure 8 top: Schematic overview showing downward reimaging of the user's hand. Bottom left: Modified setup. A hot mirror allows projector and IR camera to share same optical path. Bottom right: Two finger 'touches' observed by the camera pointed at the spinning diffuser.

First an IR camera is added to our arrangement using a hot mirror, such that it shares the same optical path as the projector as depicted in Figure 8 (bottom left). We use a ring of 850nm IR LEDs mounted onto the cavity of the top parabolic mirror to flood the area above the mirror pair with diffuse IR light. When a user's finger moves into this area it will scatter some of this light into the mirror.

The illuminated part of the finger will be reimaged downwards into the mirror cavity. Much in the same way that the diffuser intersected with the IR dome in our previous system, the spinning diffuser can intersect with the optically real image of the finger at the bottom of the parabolic mirrors. This is then visible to the IR camera underneath. Fingertips are readily detected by an IR camera pointing at the diffuser. The full setup is summarized in Figure 8.

At first the fact that the diffuser is spinning at high speeds seems to pose a major problem, but in practice using a regular 30fps camera we have observed minimal motion blur and fingertips are clearly discernible as shown in Figure 8 (bottom right).

DISCUSSION

Our Vermeer prototype generates 360° views of a 3D object which can be observed by multiple people around a tabletop surface without the need for glasses. Depending on the chosen 3D volumetric display at the heart of the parabolic mirror pair such a display can provide several depth cues: vergence, accommodation, horizontal motion parallax and stereopsis. We developed this specific 3D display based on the design proposed by Jones et al. [19]. Although many of the display's features have been demonstrated in the context of light-field and volumetric displays in the past, our system, to our knowledge for the first time, demonstrates the additional feature of direct-touch interaction within the display volume. We built a fully-functional 3D display and illustrated the design changes necessary to allow for a new rear-projected optical configuration. This configuration utilizes a commercially available optical illusion to support reimaging of the displayed content.

We have demonstrated how a depth-sensing technology such as Kinect can be used to provide 'in-display' touch input. However, perhaps more interestingly, we have demonstrated how the optical configuration can itself be leveraged, to support novel means for integrated in-volume touch sensing, whilst simultaneously displaying content. This reduces the complexity and size of the setup.

Our prototypes demonstrate proof-of-concept manipulations of 3D virtual models within the display volume. However, we have to contrast this against the relatively small dimensions of the viewable volume. This small size limits the number of fingers interacting simultaneously in the volume. It is however interesting to note that the Kinect sensing mechanism is capable of sensing a much larger volume than the viewable area. This enables us to combine indirect interaction techniques that have previously been proposed in the literature [15,16] with direct in-volume interactions. For example, a bi-manual technique could be implemented where one hand is used to select objects in the volume (by touching them) and the other hand is used to manipulate parameters from further away.

Another interesting opportunity that arises from the combination of input with our display setup has to do with the optical properties of the illusion. When a finger enters the image formed above the mirror it does not occlude the object as expected which leads to severe distortions in depth perception. Since we have full programmatic control of the displayed graphics and information about the users' fingers location, and spatial configuration, it is conceivable to mitigate this issue by providing interesting feedback to the user. One simple solution would be to move graphical elements out of the way so that they can never intersect with a physical object within the viewable volume, or perhaps add interesting effects causing 3D models to morph as fingers enter the display volume.

Finally the combination of a multi-view omnidirectional display and direct input could be used in a number of multi-user scenarios. For example, the system can show different views to users located at opposing sides of the setup (e.g., different decks of cards), allowing them to view personalized content privately.

As well as these implications for interaction, there are of course a number of limitations of Vermeer currently, which may be interesting to overcome in the future:

- *Display size* is small compared to the size of the parabolic mirrors currently.
- *Viewpoint constraints* exist even though the display is essentially walk-up-and-use; in particular view truncation is apparent when the viewpoint is too high or too low and the auto-stereo nature of the display works best at a certain distance from the display.
- *Distortion* of the displayed image occurs as only one point on the rotating projection surface is at mirror's focal point; the rest of the surface will not be reimaged without distortion. In addition, the mirrors we used in our prototype are not manufactured perfectly and cause additional (albeit minor) distortion.
- *Rendering fidelity* is essentially limited by the end-to-end speed of the rendering pipeline. With a higher DVI frame rate and correspondingly faster projector interface it would be possible to render at a higher frame-rate. A higher resolution and/or color projection system would also improve the displayed image.

FUTURE WORK

Perhaps the most obvious future work is in making the Vermeer display bigger and enabling higher fidelity. Clearly, in order to scale such an interactive display system to accommodate larger 3D scenes a larger parabolic mirror pair would be required. However in such a case the user may become too far away from the central hole to reach and interact with the 3D display. To mitigate this we have experimented with truncating the optics, i.e. using partial parabolic mirror sectors. We have cut a mirror pair in half diametrically and verified this provides a 180˚ viewing angle with the user situated within the truncated half - close enough to glance into the bottom mirror and touch the display volume. This is a compromise which may be acceptable to some users for the benefit of obtaining much larger 3D content displays whilst being close to the displayed objects to manipulate and interact with them.

Another area of future work which we have begun to experiment with is the possibility of extending the integrated sensing approach to image the environment around the display. We have observed an interesting effect that occurs when we place a webcam into the parabolic mirror arrangement. As depicted in Figure 9, this generates a reimaged virtual camera at the top of the mirror pair. When we view the live video from the webcam, we have observed that the image generated corresponds not to the real camera's location, but to the reimaged device (ignoring geometric distortions due to the mirror). Any object such as a playing card held in the field of view of the reimaged webcam is captured – as if the camera were physically at that reimaged location (Figure 9).

Figure 9 Left: a webcam placed at the bottom of the lower mirror will be reimaged and therefore viewable above the upper mirror. Right: any objects in the field of view of the floating webcam will actually be imaged by the sensor.

In our moving diffuser 3D display, we could imagine an arrangement where a high speed camera and mirror are used to image the scene from the revolving stage within the volume. This could allow us to capture the real lighting conditions and appropriately relight the virtual scene or sense the upper body of users interacting with the display, tracking their heads to correct for vertical perspective.

One final area of future work is to consider how a high speed camera could be used for our prototype in Figure 8 for more accurate 3D touch sensing. Here if such a camera is synchronized with the revolving diffuser, it should be possible to capture true 3D intersections between illuminated fingertips and each different diffuser orientation. This provides us with the possibility to sense more than just the presence or 2D location of fingertips in the volume, but instead accurately determine intersections between touch-points and the virtual 3D model. Other styles of volumetric display, in particular [29], with uniform layered switching diffuser planes may have advantages over a rotating angled diffuser in such a scenario.

CONCLUSIONS

We have shown that it is possible to use a pair of parabolic mirrors to reimage a 360° viewable 3D display enabling the use of direct in-volume interaction, which has not previously been possible with such displays. We have refined the design used in previously reported 3D displays to support our approach, and have presented our prototype system, Vermeer, in some detail. We have augmented the Vermeer display initially with a Kinect depth sensor to illustrate interaction with the displayed image.

In addition, we have demonstrated novel prototypes that replace the external Kinect sensor with integrated sensing configurations where cameras inside or adjacent to the parabolic mirror system can detect fingertip interactions within the floating display volume.

Our aim has been to introduce a new platform for interaction, enabling others to explore new interactive possibilities with such emerging 3D displays.

REFERENCES

1. Adhya, S. and Noé, J. A Complete Ray-trace Analysis of the 'Mirage' Toy. *In SPIE Education and Training in Optics and Photonics Conference (ETOP)*. 2007.
2. Agrawala, M., Beers, A., McDowall, I., Froehhlich, B., Bolas, M., Hanrahan, P. The two-user Responsive Workbench: support for collaboration through individual views of a shared space. *In ACM SIGGRAPH '97*.
3. Bimber, O. and Raskar, R. 2005. Spatial Augmented Reality: Merging Real and Virtual Worlds. *A. K. Peters*, Natick, MA, USA.
4. Bimber, O. Fröhlich, B., Schmalstieg, D. Encarnacao, L.M. The virtual showcase: a projection-based multi-user augmented reality display. *In ACM SIGGRAPH'02*
5. Blundell, B. and Schwarz, A.J. *Volumetric Three-Dimensional Display Systems*, Wiley-IEEE Press, 2000
6. Blundell, B. and Schwarz, A.J. *Creative 3-D Display and Interaction Interfaces: A Trans-Disciplinary Approach*, Wily-Interscience, 2005.
7. Cobb, S., Nichols, S., Ramsey, A., Wilson, J. Virtual Reality-Induced Symptoms and Effects (VRISE). *Presence: Teleoper. Virtual Environ. 8*, 2 (April 1999)
8. Cossairt, O.S., Napoli, J., Hill, S.L., Dorval, R.K., Favalora, E. Occlusion-capable multiview three-dimensional display. *In Applied Optics (46)*. (2007).
9. Dodgson, N. A. 2005. Autostereoscopic 3D displays. *Computer 38, 8, 31–36*.
10. Dorval, R.K., Thomas, M., Bareau, J. Volumetric three-dimensional display system. *US Patent #6,554,430* (issued 29 Apr 2003).
11. Elings, V.B., Landry, C.J. Optical Display Device. *US Patent #3647284*. Issued March 7, 1972.
12. Endo, T., Sato, M., Kajiki, Y., Honda, T. Cylindrical 3D Video Display Observable from All Directions. *In Proc Pacific Conference on Computer Graphics and Applications (PG '00)*.
13. Favalora, G. E. 2005. Volumetric 3D displays and application infrastructure. In *Computer 38, 8, 37–44*.
14. Grossman, T. and Balakrishnan, R., An Evaluation of Depth Perception on Volumetric Displays, In *AVI 2006*
15. Grossman, T. and Balakrishnan, R., The Design and Evaluation of Selection Techniques for 3D Volumetric Displays, In *Proceedings of UIST 2006.*
16. Grossman, T., Wigdor, D. and Balakrishnan, R., Multi-Finger Gestural Interaction with 3D Volumetric Displays, In *Proceedings of UIST 2004.*
17. Halle, M., Autostereoscopic displays and computer graphics. In Proceedings of *ACM SIGGRAPH* '97
18. Hoffman, D., Girshick, A., Akeley, K. and Banks, S. 2008. Vergence–accommodation conflicts hinder visual performance and cause visual fatigue. *In Journal of Vision* 2008 8(3): 33
19. Jones, A., Bolas, M., McDowall, I., Yamada, H., & Debevec, P. Rendering for an Interactive 360 Degree Light Field Display, In *Proceedings of ACM SIGGRAPH 2007*.
20. Kimura, H., Uchiyama, T., and Yoshikawa, H. 2006. Laser produced 3D display in the air. In *Proceedings of ACM SIGGRAPH 2006 Emerging Technologies*.
21. Kiyokawa, K. , Kurata, K., Ohno, H. An Optical See-Through Display for Mutual Occlusion of Real and Virtual Environments. *In Proc ACM ISAR 2000*.
22. Maeda, H., Hirose, K., Yamashita, J., Hirota, K., and Hirose, M. 2003. All-around display for video avatar in real world. In *Proceedings of ISMAR 2003.*
23. Malik, S., McDonald, C. and Roth, G. "Hand Tracking for Interactive Pattern-based Augmented Reality". In *Proceedings of ISMAR 2002*.
24. Opti-Gone, http://www.optigone.com/m22.htm (verified July 2011).
25. Poston, T., Serra, L. Dextrous virtual work. *Commun. ACM 39, 5 (May 1996)*, 37-45.
26. Scharver, C., Evenhouse, R., Johnson, A., Leigh, J. Designing Cranial Implants in a Haptic Augmented Reality Environment. *Comm. of the ACM, Volume 47, Issue 8*, August 2004
27. Shin-Etsu View Control Film, http://www.shinetsu.info/index.php/products/view-control-film/view-control-film
28. Suffern, K., Ray tracing from the Ground-up. AK Peters, pp. 22, 2007.
29. Sullivan, A. A Solid-State Multi-Planar Volumetric Display, SID Symposium Digest Tech Papers 34, 1531-1533 (2003).
30. Travis, A. R. L. The display of three-dimensional video images. *Proceedings of the IEEE 85, 11 (Nov).*
31. von Wiegand, T., Schloerb, D., Sachtler, W.L. Virtual Workbench: Near-Field Virtual Environment System with Applications. *Presence: Virt. Environ. 8, 5 (October 1999)*.
32. Yendo, T., Kawakami, N., and Tachi, S. 2005. Seelinder: the cylindrical lightfield display. In *Proceedings of SIGGRAPH 2005 Emerging Technologies*.