

Through the Combining Glass

Diego Martinez Plasencia¹, Florent Berthaut¹, Abhijit Karnik², Sriram Subramanian¹

¹Department of Computer Science, University of Bristol, UK

²School of Computing and Communications, Lancaster University, UK

{diego.martinez-plasencia, sriram.subramanian}@bristol.ac.uk, a.karnik@lancs.ac.uk




Figure 1: Interaction opportunities of our design space: (A) Reflections allow users to interact with artifacts inside a museum cabinet. (B) Reflections can reveal internal details of real objects. (C) A Digital Musical Instrument augmented with projection mapping and a volumetric display. (D) This augmentation is visible from any point of view and from both sides of the mirror.

ABSTRACT

Reflective optical combiners like beam splitters and two way mirrors are used in AR to overlap digital contents on the users' hands or bodies. Augmentations are usually unidirectional, either reflecting virtual contents on the user's body (Situating Augmented Reality) or augmenting user's reflections with digital contents (AR mirrors). But many other novel possibilities remain unexplored. For example, users' hands, reflected inside a museum AR cabinet, can allow visitors to interact with the artifacts exhibited. Projecting on the user's hands as their reflection cuts through the objects can be used to reveal objects' internals. Augmentations from both sides are blended by the combiner, so they are consistently seen by any number of users, independently of their location or, even, the side of the combiner through which they are looking. This paper explores the potential of optical combiners to merge the space in front and behind them. We present this design space, identify novel augmentations/interaction opportunities and explore the design space using three prototypes.

Author Keywords

Optical Combiners; Augmented Reality; Multi-user

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI):

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive version of Record was published in UIST '14 which can be found at:

<http://dx.doi.org/10.1145/2642918.2647351>

INTRODUCTION

Reflective transparent surfaces (beam splitters, half-silvered mirrors, shopping windows or, more formally, *planar reflective optical combiners*), show our reflection overlapped with the objects behind them. The planar nature of the surface ensures that each point of the reflection produces a *stigmatic pair* [19]. The reflection of any point is seen at a unique 3D position behind the combiner, independent of the observer position and reflections are perceptibly indistinguishable from a real object behind the mirror [3]. Also, the partially transparent nature of the combiner implies that the objects behind the combiner are merged with the reflections of the objects in front of it.

Augmented mirrors [1, 9, 11, 21, 27] use reflections as high fidelity avatars that can be augmented with virtual clothes or visual effects. Situated AR [13, 14, 24] exploits planar reflections by overlapping 3D virtual objects with the physical space and enabling direct interaction with those objects. In all these applications, augmentations remain unidirectional – either augmenting our real self or reflected body with virtual contents.

However, *planar reflective optical combiners* have a bigger potential. They visually merge the physical spaces in front and behind them, which enables many more novel augmentations and interaction possibilities.

Firstly, users' reflections can overlap unreachable content placed behind the combiner, enabling interaction with this content. For instance, when looking at objects inside a museum cabinet (as in Figure 1.A) we can use the fingers' reflection to create a 3D cursor that is visible inside a museum cabinet, to any number of users without specialised eyewear. In the *Augmented Musical Instrument* prototype the user can interact with content inside a




Figure 2: Approaches related to interacting through the mirror: (A) Video see-through augmented mirrors, (B and C) different implementations of augmented mirrors. (D) Situated Augmented Reality. (E) Shadow-based interaction.

volumetric display, which in most cases cannot be reached into (Figure 1.C).

Secondly, augmentations can be bidirectional. As both spaces are blended, both can contribute to create a unified experience. Simple projection mapping on the artifacts inside the museum cabinet can easily augment their surface. However, as illustrated in Figure 1.B, by projecting on users' hands in front of the combiner we can allow the user to slice through the museum artifacts to reveal their internal structure. The reflection is actually mapped inside the object and this augmentation is visible from any viewpoint.

Thirdly, the combination of spaces achieved by placing a planar optical combiner and looking through it can be consistently seen by any number of users, independent of their location or the side of the combiner through which they are looking. This for example, allows spectators of our *Augmented Musical Instrument* to freely move around the two-way mirror, and yet always see the musician immersed with the digital contents (Figure 1.C and . Figure 1.D)

This paper takes a new perspective on *planar reflective optical combiners*, focusing on their ability to merge the spaces in front and behind them. We describe the design space, identify novel interaction and augmentation possibilities. We describe relevant features that designers need to address and demonstrate/exemplify them using three prototype systems. We also report interviews with practitioners, which illustrate the potential of our approach and suggest even more examples of how our techniques can be used in practice.

RELATED WORK

Reflections have captivated our imagination and drawn the attention of many pieces of literary and cinematographic fiction. Their ability to create a reflected version of the real world has been used to portray them as portals to alternative realities (e.g. Through the looking glass¹). Other stories use mirrors that not only reflect the visible world, but also show other elements beyond our perception. In the *Never Ending Story*², Atreyu faces a mirror that reflects the viewer's true

self, which becomes one of the most revealing points in the story. The *Mirror of Erised*³ reveals Harry Potter's deepest desires. Exploiting this aura of mystery, magicians and theatre performances have used mirrors to create different special effects, such as the well known Pepper's ghost⁴.

This idea of augmenting reflections with non-real (virtual) elements has been explored in practice. Approaches like *i-Mirror* [12], *augmented mirror box* [26] or the virtual mirrors by Adidas [9] and Bloomingdale's⁵ use video see-through AR (see Figure 2.A), recording the user and adding virtual elements to the digital image. These approaches lose the affordances of real mirrors, replacing volumetric reflections with digitized 2D images of the users, having a fixed perspective (i.e. the actual position of the camera) and inhibiting mutual eye contact between two observers.

Other systems combine a planar optical combiner with 2D displays. In some of them, the image of the augmentations is formed on the combiner's surface (Figure 2.B), either because a flat display is right in contact with it [11, 12, 21, 27] or they use projectors pointing at the back of the mirror [1]. Given the limited depth of field (DOF) of human eye, in most of these systems (with the exception of *YouMove* [1]), users will be able to focus on either their reflections or the augmentations, but not both.

Acknowledging this constraint, systems like *i-Mirror* [12], *Holoflector*⁶ and *Holocubtile* [8] place the display at a distance from the optical combiner. This allows the augmentations to appear at the correct depth (Figure 2.C). The user's eye is able to correctly focus on both the reflection and the augmentations. It also allows the augmentations to correctly map to the location of the user's body (*Holoflector*) or around/inside the *Holocubtile*. Also several observers can see virtual contents correctly aligned to the real ones. Disney's *Haunted Mansion*⁷ uses a similar

¹ http://en.wikipedia.org/wiki/Through_the_Looking-Glass

² [http://en.wikipedia.org/wiki/The_NeverEnding_Story_\(film\)](http://en.wikipedia.org/wiki/The_NeverEnding_Story_(film))

³ http://harrypotter.wikia.com/wiki/Mirror_of_Erised

⁴ http://en.wikipedia.org/wiki/Pepper%27s_ghost

⁵ <http://gizmodo.com/245716/interactive-mirror-debuting-at-bloomingdale>

⁶ <http://research.microsoft.com/apps/video/default.aspx?id=159487&l=i>

⁷ <https://disneyworld.disney.go.com/attractions/magic-kingdom/haunted-mansion/>

approach to enable interaction of virtual ghosts with users' reflections (e.g. swapping visitor's heads). However, the resulting systems have a large footprint; the contents are flat and can only be presented at a fixed depth.

Situated Augmented Reality systems [5], like the *Virtual Workbench* [24], *Toucheo* [13] and *HoloDesk* [14] use the same principle in an inverse manner. These systems reflect virtual contents that seem to overlap a user's real hands (see Figure 2.D). In these devices, the display's reflection appears at the approximate depth where the user's hands will be located. The smaller working area decreases the footprint and active or passive stereo can be used to display 3D contents. However, they usually support only one user and place constraints on the location of the user's eyes and hands. *Vermeer* [6] and *RePro3D* [37] combine other arrangements (parabolic mirrors and two-way mirrors with retro-reflective materials, respectively) with a volumetric display to collocate the user's finger into the 3D volume. No eyewear is required in this situation, but the footprint limitations and constraints on observer locations still apply.

A 2D approach with some resemblance to our approach is shadow-based interaction (Figure 2.E) as demonstrated by Myron Kreuger's *VideoPlace* [20] or Benko and Wilson [2]. In these examples, the users' shadows become 2D embodiments that allow them to interact in a remote 2D surface (i.e. to play games or explore the night sky). In the case of *VideoPlace*, these "shadow embodiments" can also be augmented/modified, but the reverse possibility (shadows to augment projected contents) was not considered.

COMBINED SPACES THROUGH THE GLASS

Planar reflective surfaces create *stigmatic pairs* [19]. This means that they create a bijective mapping, where each point in front of the mirror is reflected as exactly one point behind the mirror ($f(x,y,z)=(x,y,-z)$, assuming the mirror plane as $z=0$). This mapping is independent of the optical path (observer location). If thin optical combiners are used (to minimize the effects of refraction), this results in reflections: a) that are exactly similar to the original object; b) whose location is consistently overlapping a unique point behind the combiner and c) are visible at that location independent of the observer position.

The ability of optical combiners to also transmit a part of the light incident on them allows reflections to overlap the objects behind it. A second benefit is that the bijective mapping is also effective from the other side of the mirror. As a result, any observer looking through the combiner will see the spaces in front and behind the mirror blended, irrespective of his/her position or the side of the mirror he/she is looking through. It is even possible for users to establish mutual eye contact (reinforcing agency and immersion). However, this illusion of two collocated spaces only happens while users look through the combining glass.

A common theme within previous approaches using *planar reflective optical combiners* is that these systems aim to augment our real or reflected body with virtual contents. One space contains the virtual contents only (displayed through some display technology), while the user is constrained to the other side.

We acknowledge the potential of such an approach, but consider it as a limited use-case of the ability of planar optical combiners to merge the space in front and behind the mirror. In the following subsections we describe novel interaction/augmentation possibilities that result from applying our perspective: a) merging our (real or reflected) body with elements on the other side of the combiner; b) the potential of objects in any of these spaces to augment the other space (bidirectional augmentations); or c) the possibility to experience them from both sides of the combiner.




Figure 3: A. Musician's reflection can reach through the panels of our volumetric display. (B and C) Reflected hands serve as multiuser 3D pointers inside the cabinet.

Using reflections as interactive avatars

The combiner enables an environment where the user's reflection can co-habit a physical space behind the combiner. Such a reflection can act as a high-resolution and responsive virtual avatar that can interact with objects on the other side of the combiner. This has obvious benefits when the objects need to be beyond our reach, such as in a museum or a jewellery shop.

While the user perceives the reflected avatar as a physical embodiment collocated with objects behind the mirror, the avatar has no corporeal form. Thus it can pass through solid objects and is not affected by toxic or dangerous objects. In an operating theatre scenario, an instructing surgeon's reflection can be mapped by the operating table without affecting the sterile environment. As discussed later in our Musical Instrument example, the combiner allows the reflection of the musician to reach into a volumetric display (Figure 3.A) which is inaccessible in many cases.

The reflected avatar can also solve technical constraints that cannot be solved by using real physical objects. For example, it would be very difficult to create a 3D cursor that was visible inside a museum cabinet, to any number of users without specialised eyewear. However, our reflected fingertip can easily act as a 3D cursor floating inside the cabinet. For instance, we use the reflection as a virtual

camera to explore 3D reconstructions of museum exhibits. This can reveal close details of the artifacts that are hardly visible from outside the cabinet (Figure 3.B). We also use it as a little lamp, to illuminate the artifacts and facilitate inspection (Figure 3.C).

It is important to note that interacting using your reflection can raise challenges. While moving in the plane parallel to the mirror is simple, others interactions might be more difficult. We discuss these issues later in this paper.

Bidirectional augmentation

For any user looking through the combiner, both spaces appear blended together. Augmentations can be used on either side to create a unified space. For example, we can place a volumetric display like DepthCube [34] on one side of a planar optical combiner overlapping it with musical controller itself augmented through projection mappings. The volumetric display at the other side of the mirror shows augmentations that seem to float over the controller, together with the augmentation on the controller's side (see Figure 4.C).

In addition to augmenting reflections with virtual contents, optical combiners can allow our reflection to slice through solid objects placed behind them. Using projection mapping on a handheld physical prop, one can see it overlapped inside the object (Figure 4.B). The projected image is reflected off the planar combiner and aligned with the intersection of the slicer and the object when viewed through the combiner. This can be consistently perceived by any number of observers when seen through the optical combiner. An example scenario where this is useful is airport security checks, where the screening agent can visually reveal where the passenger may be hiding something while both of them are separated by a glass wall.

DESIGNING WITH REFLECTED SPACES

This section describes relevant aspects that designers need to take into account when creating systems that merge the spaces in front and behind the mirror. The prototypes, described later in this paper, are used here to explore the design space, demonstrating each of its aspects (e.g. materials, illumination, display technologies). They are also used to illustrate interesting uses of our concepts, such as the ability to use our reflection to overcome constraints that apply behind the combiner, using augmentations on either side to build a common interactive experience or the possibility to see them from any point around the mirror.

Controlled combination of spaces

Pepper's ghost⁴ is one of the earliest techniques exploiting the combination of contents in front of and behind the mirror. In this theatre technique, a half silvered mirror located at 45 degrees to the stage maps a dark, hidden room on top of the visible stage. By increasing the lighting of the hidden room, the objects within seemed to appear, like ghosts, in the middle of the stage. This is a very powerful




Figure 4 A. Augmentations on either side of the combiner are merged. (B and C) Projector based illumination allows interactions between objects (i.e. internals and intersection).

feature that designers need to consider when creating interactive experiences through a combiner. The perceived visibility of each space depends on two factors: the Transmission to Reflectance (TR) ratio of the reflecting surface and the illumination of each space.

Mirror fabrication techniques allow optical combiners with a wide range of TR ratios. The TR ratio determines the static visibility/blending of both spaces and its choice will be influenced by the requirements of the system being designed. For instance, our museum cabinet uses a low reflectance surface, as users are more interested in seeing the objects within the cabinet than their own reflection. In the *Augmented Musical Instrument (AMI)* and *AR mirrors*, we use a 50% TR two way mirror, as the goal is to equally merge both the user/musician and the virtual contents in a common 3D space.

The second relevant factor is illumination, which allows us to control the dynamic visibility of both spaces. Apart from a global control, like the used in Pepper's ghost, the use of projector based illumination [25] and advanced sensing technologies allows us to exert a fine grained control over the visibility of elements.

Reproducing effects such as consistent occlusions requires observer tracking and a different optical combiner for each user [5], which would break our unified combined space. However, other viewpoint independent effects, like intersections (hiding the parts of the objects inside other objects), can be reproduced consistently. Figure 4.A shows how only the parts of the hand that are outside the earth are illuminated. Non-illuminated parts seem to disappear inside the objects.

Reflections, depth of focus and display technologies

One can use a wide range of technologies to create the virtual 3D contents that augment the combined space. Many taxonomies have been proposed to classify these 3D capable displays, for e.g. Holliman et al. [17]. The classification is primarily based on the depth perception cues that they afford, identifying two broad categories *true depth 3D displays* and *flat panel 3D displays*.

As mentioned earlier, reflections in a mirror provide all perception cues and are indistinguishable from a real object placed behind the mirror (although occlusions of objects




Figure 5: Parameters to consider to assure reflections fall within the zone of comfort of the planar display.

behind the mirror will be missing). This allows reflections to be naturally blended with contents in *true depth 3D* displays. The museum cabinet and AMI exemplify the usage of projection mapping and volumetric displays (two broad categories of true depth 3D displays).

Flat panel 3D displays (e.g. parallax barriers or multi-view displays), on the other hand, can suffer from vergence-accommodation (VA) conflicts. This conflict arises when these displays show an object that appears at a different depth than the plane of the display. The user's eyes verge towards the location where the object is supposed to be, but they still need to focus on the display screen. There exists a "zone of comfort" where the conflict causes less severe symptoms (see [29] and [16] for an interesting discussion). Several estimations have been proposed [22, 28], and it is commonly assumed that human eye has an average *depth of field* (DOF) of ± 0.3 dioptres (DPT), allowing objects at that distance from the display to remain in focus [7].

The AR mirrors following the approach in Fig 2.C. (i.e. the display is located right behind the mirror) can illustrate the relevant role that DOF plays in our systems. In these systems, our reflection is located behind the mirror, but augmentations appear on the mirror's surface. If we focus on our reflected face, the screen can fall outside the depth of focus of our eyes, so we cannot fixate on the virtual content. If we focus on the screen, we might not be able to fixate on the reflection. Both our reflection and the augmentations should fall inside the DOF of our eye, but in this approach (i.e. the display is located right behind the mirror) this only happens when the observer is further apart than 1.66 meters of the mirror (i.e. the display depth is $1/1.66\text{m} = 0.60$ DPT and the reflection $1/3.33 = 0.30$ DPT; the difference matches the average human DOF, so both remain in focus).

Thus, some consideration is needed when placing the display, so that its contents are merged with the observer's reflection. We describe four steps required to avoid these issues and we will use our prototype of AR mirror as an example to illustrate these steps (Figure 5).

Determine size of the interactive space around the user

Determine the space that will be mapped behind the mirror and immersed with the contents shown in the display (working volume). In our AR mirror the user's arm length (L) is used as the baseline for our design. Users must be able to use their reflection to interact comfortably. While fully stretching the arms can be uncomfortable, a distance between $0.4L$ and $0.8L$ allows for comfortable interaction [35]. We also wanted to be able to modify the user's facial appearance. The size $L=70$ cm (average arm's length in a male adult) was found enough to accommodate the user's arms during interaction and his face/body, as to augment his appearance.

Determine location of the planar 3D display

The size of the zone of comfort of a planar display is determined by the distance between the observer and each point of the display, which results in a parabolic shape. However, from a design point of view it is usually more practical to identify the closest point and approximate the shape by a parallelepiped.

Considering the observer distance to the display (o), the limits of the region (i.e. near (n) and far (f) planes) need to satisfy some conditions. First, their distances to the display in dioptres must equal the estimated DOF used ($R=0.3$ DPT). Also, they must accommodate the size of the working volume (Eq 3).

$$\frac{1}{o} - \frac{1}{n} = R \quad (1); \quad \frac{1}{f} - \frac{1}{o} = R \quad (2); \quad f - n = L \quad (3)$$

Solving this set of equations, one can determine the required observer distance (Eq 4) and the distances of the near and far planes (Eq5 and Eq3). Given the size of the working volume in our example ($L=70$ cm), this means that the display must be located at $o=1.03$ m, to provide a zone of comfort of the desired size ($n=0.79$ m (1.27 DPT); $f=1.49$ m (0.67 DPT)).

$$o = \sqrt{\frac{L}{(2+R \cdot L) \cdot R}} \quad (4); \quad n = \frac{o}{1+R \cdot o} \quad (5);$$

Determine location of the mirror

Once we determined the space around the user (working volume) and the display's zone of comfort, the designer needs to position the mirror to make sure that the working volume is mapped within the display's zone of comfort.

The working volume in our example covers 56cm in front of the observer's eyes (space for user to interact with objects) and 14 cm behind (to include the rest of the head and back). The plane at 56cm in front of the observer will be mapped to the near plane of the zone of comfort, at 78cm in front of the observer. As a result (assuming negligible refraction), the mirror needs to be located at $(56+78)/2=67$ cm in front of the user (almost at the limit of his reach). The display is located $103-67 \sim 35$ cm behind

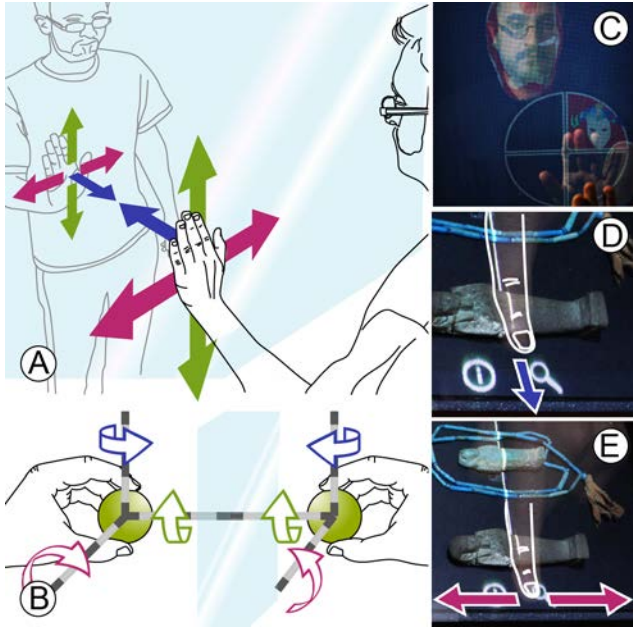


Figure 6: Analysis of translations (A) and rotations (B) through the mirror. Example interactions that streamline natural and non natural translations in the museum cabinet (C and D) and the AR Mirror (E).

the mirror. This allows us to allocate real objects between the mirror and the display, as explained later in the description of the AR mirror.

Determine the size of the display

Given that the mirror lays in the exact midpoint between us and our reflection, the size of our reflection when projected on the mirror is always, exactly, half our size [3]. Once the location of the mirror and the display technology has been determined, the interaction designer should use this feature to work out the required size of the planar display, so that the whole reflection overlaps the display, making it possible to augment the reflection. In our AR mirror example, the required size of the panels was 61x38 cm to actually cover the whole of the user's head and torso (80x50 cm).

Manipulation through a mirror

A particular property of the systems proposed in this paper is that users need to look through the optical combiner to experience the blended spaces. This implies that they will always interact with their reflected image.

Mirrors have been used to explore our ability to interact using reflected images, or even totally different mappings. Mirror tracing tasks [30] reveal that even though performance is initially significantly lower, subjects learn quickly over time [30].

George M. Stratton explored a variety of more complex mappings. For several days he wore contraptions that allowed him to only see the world upside-down [32], left-right inverted [33] or even observing his own body from

above [31]. These experiments revealed that, even though proprioceptive conflicts were initially very strong, these disappeared after three days of usage. These experiments give evidence of the ability of human brain to learn and adapt to this kind of interaction, and are encouraging as they indicate that, should mirror interaction become widespread, users could easily learn and adapt to it.

However, we still need to manage the interaction to support walk-up and use interactions where the user will not have the opportunity to learn through repeated use. As shown in Figure 6.A, translations along the mid-frontal plane remain natural (movements parallel to the mirror are mapped to similar actions in the other side of the mirror) but, translations along the Z axis are inverted.

The opposite situation arises with rotations (Figure 6.B). Rotations around the Z axis are natural, but those involving the other two axes are inverted. Our initial observations revealed that manipulations were easily understandable as far as they remained in one of these two categories (i.e. translations parallel to the mirror or purely perpendicular to them). However, trajectories involving both (i.e. movements in a horizontal plane) were more confusing.

The interaction designer should take these limitations into consideration when creating the interactive dialogue [10] with the system. Even if the task has three degrees of freedom (e.g. translation), and the finger/hand can provide them, it seems preferable to decompose it in steps that do not mix natural and inverted interactions, whenever possible.

The menus in our AR mirror follow this rationale. When the user raises his hand in front of him, a ring like menu appears. He can select a specific category using natural translations in the plane parallel to the mirror (Figure 6.C).

The museum cabinet poses a more challenging scenario. Users first select artifacts with their finger. This is a coarse selection, and actually needs to mix natural and inverted translations, as artifacts are laid horizontally on the cabinet. When the options menu is shown, users need to perform a translation along the Z axis (inverted) to select it, which corresponds to a mid-grained interaction (Figure 6.D). Users can then select specific options (fine grained interaction) using natural translations (Figure 6.E).

CASE STUDIES

We used our prototypes to exemplify the different aspects discussed in our design space and to give examples of how our techniques can be used. This section provides technical details related to their implementations but, more importantly, it describes the feedback we obtained while presenting them to different practitioners (musicians, museum staff). This feedback highlights the usefulness of different elements in our prototypes, but also proposes new ways in which our techniques can be applied.

Interactive Museum Cabinet

Our AR cabinet is designed to be included in the archaeological exhibition of the Bristol Museum and Art Gallery museum. The exhibit currently involves an exhibition coordinator, a digital manager and a content curator of archeological content. We report the implementation of our AR cabinet and relevant impressions/conclusions, relative to the use of reflections, gathered from the museum staff.

Current museum cabinets contain the exhibition artifacts, but users need to refer to touch panels by the cabinet to receive information about them. This forces them to divert their attention from the actual artifacts and interaction is often described as boring and not engaging enough. At the same time, when several visitors are discussing about the artifacts, they need to point at them through the glass, which can lead to misunderstandings (e.g. “which artifact are you talking about?”). Also, users can accidentally touch the cabinet, making it dirty or, even, triggering security alarms.

Our prototype allows users to select artifacts by using the reflection of their hands. Several users can mutually see these reflections, which facilitate understanding and discussion. We also emulate a little spotlight attached to the user’s fingers. This allows users to illuminate parts of the artifacts to see them better. Our preliminary impressions seem to indicate that the spotlight facilitates the discoverability of our technique. When users move their hands in front of the cabinet, they soon realize that the spotlight moves in a mirrored way, which facilitates them to realize that it is actually their reflection which allows them to interact. By touching the artifacts, users can get short overviews (name and picture) about the artifacts. They can also get additional content related to the artifact (textual information, explore 3D reconstructions) by touching buttons projected close to the artifact.

The current prototype only contains replicas of actual artifacts. Some of them are modeling clay reproductions, textured using projection mapping. We use a Leap motion to track the users’ fingers and compute the position of the finger reflection inside the cabinet, a projector in the top to project augmentations and a TFT panel in the back to display textual information. The cabinet’s combiner has low reflectivity, which allows reflections to be marginally visible but maximizes the visibility of the objects within. Illumination outside the cabinet only affects the users’ hands, so these are reflected (i.e. but not their faces). Also, the short user distance to the back display (60cm) creates a zone of comfort of only 10 cm in front of it. As a result, a user reading information will not be able to fixate on the hand reflections, avoiding distractions. The dotted connecting line between the artifacts and the dialogues allows progressive adaptation of the focus, when users transition between the dialogues and the actual artifacts.




Figure 7: (A) Our AR museum cabinet allows users to interact with artifacts inside it, keeping track of the actions of other users in front of the cabinet. (B) Design of the cabinet.

We presented the prototype to the museum staff, guiding them through its various features. Besides comments related to adapting cabinets for other sizes and formats (e.g. portrait mode, more non-textual information for kids), they had a very positive reaction to the capability of directly interacting with the artifacts and the fact that several visitors could understand each other’s actions.

They found that the interaction with the menus around the artifacts was initially confusing, but agreed that it was easy to learn after some minutes. They proposed placing a cabinet with reduced functionality in the entrance hall (that displays overviews when users touched the artifacts) to allow users understand the technique, and cabinets with complete functionality inside the exhibition, to get full benefit of their potential.

They found the spotlight especially interesting, as it allows the cabinet to remain almost dark, only illuminating the artifacts when someone is looking at them. This can be useful for delicate artifacts (i.e. sensitive to light radiation) being even possible to dynamically dim the light when users illuminate sensitive parts.

They also proposed an unexpected use for the two sided feature of planar combiners discussed in this paper. A wide angle camera could be placed inside the cabinet, facing towards the users. This would record the user’s actions, but the artifacts’ reflections would also be visible as users interact with them. This image could be displayed in a bigger screen, by the cabinet in the entrance hall, so incoming visitors would see and learn how cabinets work.

The possibility to reveal object’s internals by projecting on the users was found especially attractive for younger audiences and some specific exhibition artifacts (e.g. rich people were buried in nested coffins, which could be revealed with our technique).

AR mirror

This prototype was also the result from our interviews with the museum staff. During one of them, they realized the back panel in our cabinet was placed at approximately the same distance from the combiner as the user’s face. As a




Figure 8: (A) Our AR mirror uses a two way mirror, an LC sandwich, to support several autostereoscopic techniques and allocate real objects. (B) Augmentations overlaid on a user.

result, it would be possible to display augmentations that seemed to overlap their faces (e.g. a pharaoh mask for younger audiences), mimicking the behavior of *Holoflector*.

In order to minimize footprint and allow users to move in front of the mirror, we use a modified LC sandwich arrangement, similar to *MUSTARD* [18]. It allows us to use a variety of 3D techniques for planar displays: perspective corrected *FishTank VR* [36], parallax barriers [23] and multiuser random hole masks [18]. This development also led us to realize the special considerations that need to be considered when mixing planar displays and reflections, and allowed us to solve the limitations in previous approaches to AR mirrors (i.e. simultaneous fixation on reflection and virtual imagery (see Figure 2.C and D)).

The system (Figure 8.A) uses a Kinect to track the user's eyes, required to provide motion parallax and perspective correction, but also to recompute the masks in the LC panels. The Kinect tracks the user's arms and body, which allows us to map augmentations such as hats, glasses or avatars (Fig 8.B).

Using the four steps described earlier and an average adult's arm length $L=70\text{ cm}$, we determined that the display needs to be located 35 cm behind the mirror, and the observer at 67 cm. This gap would still be enough to fit medium size exhibits. Also, the mirror is still in the limits of the users reach, so touch interactions with the surface of the mirror would also be possible.

Digital Musical Instrument

In collaboration with local artists, we used a planar optical combiner to augment a Digital Musical Instrument, with the aim of creating new musical performance experiences and exploring new musical interaction possibilities. The original instrument features a hardware controller with pressure sensors alongside linear and angular potentiometers. Sounds are triggered by hitting the sensors and modulated

by applying pressure to them. These manipulations can be recorded and played as loops afterwards. These complex and subtle interactions of the musician with the instrument sometimes make it hard to understand the connections between his actions and the music [4].

Our prototype uses a two way mirror to combine both spaces, so that they are approximately equally visible. The musician and augmented controller are on one side, and a low fidelity implementation of *DepthCube* [34] is on the other. This volumetric display is built using three modified LG IPS2321P 23", 1980x1080 LCD monitors. We dismantled them from their frames and removed their front and back polarizers. Then we installed them in a transparent mounting frame, and this one inside a box containing a backlighting, front and back polarizers. The box is purposely bigger, as to allow bigger viewing angles for the spectators.

The controller is augmented with light diffusing acrylic and a projector in order to amplify musical gestures. A Kinect records the musician as she creates the loops. These point clouds are shown inside the volumetric display, at the other side of the mirror as if they floated above the controller. This allows spectators to visualize the musicians' actions that produced the sound in the loop. At the same time, it extends the manipulation of loops into the volumetric space above the controller. Loops can be placed over different sensors to control them in different ways.

The mirror combines the augmentations from both sides, showing the musician, the extensions under the controller and the loops above, all inside a common space. Regardless of where the spectators stand, they will always see the combined space (musician and augmentations) through the mirror.

Musicians we worked with really appreciated the augmentation possibilities offered by the combined space. They were enthusiastic about the additional interaction capabilities created by the mirror. They also believed that it could dramatically improve spectators' perception and understanding of digital performances. In particular, they pointed out three main aspects of the system that enhanced audience engagement.

They felt that the augmentations could be used to improve the connections between the elements of the instrument such as sensors with the visual representations of the loops. For example, links could be projected between them in both the real and virtual spaces. They were extremely positive about the visibility of the virtual content and wished for dynamic visibility control of both the virtual content displayed on our *DepthCube* and of the musician. A solution would be to increase the brightness of our display by using more transparent LC panels or to switch to greyscale panels to prevent colour filters from interfering. Along with this we can also provide projection mapping on the musician. Finally, in addition to the *Depth Cube* they

were really enthusiastic about the possibility of using their own reflection or the audience's reflection as a way to reveal and play with augmentations that could be placed anywhere in the combined space. Instead of revealing the internals of an object as shown in Figure 4B, spectators could reveal the components of the instrument through more active interaction with the optical combiners; for example, by moving parts of their body/ reflection through the combined space.

THROUGH THE DISTRIBUTED GLASS

One of the implicit limitations of the design space explored in this paper is the fact that the mirror can only overlap the spaces immediately in front and behind it. While this lends itself to a variety of applications (surgery rooms, security checks, museums, etc.); it does not allow collaboration with remote participants.

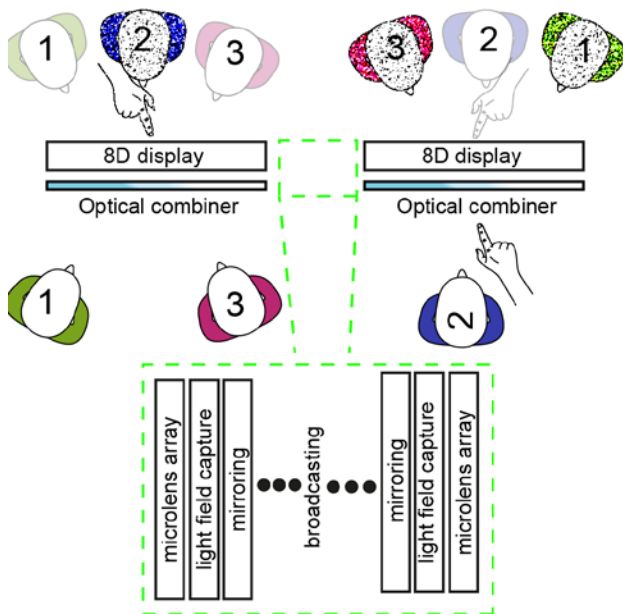


Figure 9: A schematic of how a distributed combining glass could be implemented.

Advances in display technology seem to indicate that such scenarios would be possible. Lightfield cameras that capture a low angular resolution lightfield are commercially available⁸ and Hirsch et al. [15] proposed a display that captured an approximately 7x7 lightfield in front of it, to illuminate a virtual scene inside the display.

Instead of projecting the captured lightfield forward, one could reflect it, mimicking the behavior of a mirror. This would create a digitized version of the reflection, which could then be then broadcasted to enable collaborations like the ones proposed in this paper in a distributed environment. Users would stand in front of the optical

combiner, seeing their own reflections and the digitized reflections of other remote users (see Figure 9).

Unlike computational holography displays, the lightfield is not computed, but could directly be captured and retransmitted. Advances in camera sensor integration and appropriate lenses to increase angular resolution of the lightfield would allow the complete lightfield to be digitized. Improvements in broadband connections would also be necessary as to retransmit such lightfields. Such a scenario would provide exciting possibilities, such as telemedical assistance, interaction with people in restricted environments like nuclear power-plants, hard-to-reach areas and military combat environments. Challenges such as bandwidth and angular range of the captured lightfield need to be overcome before such future scenarios can be realised.

CONCLUSIONS

In this paper, we have explored the design potential of planar optical combiners as a means to merge physical spaces in front and behind the combiners with interactive virtual display spaces. The merged space offers a wide range of interactive capabilities. We discuss the design considerations pertinent for a designer using such optical combiner setups. In addition to this, we implemented three independent prototypes that leverage different aspects of the interactive capabilities. We demonstrated two of the prototypes to the museum staff and gathered initial impressions from the curators of the museum. The third prototype (AMI) was deployed for a musical performance for musicians. Finally we discuss how the optical combiner approach can be further extended to operate in a distributed setup allowing remote collaboration and natural telepresence. We believe that our exploration of the design space adds to the existing work and extends it into new and exciting directions.

ACKNOWLEDGEMENTS

We thank Alex Jones, Nick Janaway and the personal of the Bristol Museum and Art Gallery for their advice in designing and evaluating our prototypes. We thank Matt Sutton for taking the images and the video. This work is supported by the European Research Council (Starting Grant Agreement 278576) and the EU Marie Curie Action (IEF 2012 Number 330770) both under the Seventh Framework Programme.

REFERENCES

1. Anderson, F., Grossman, T., Matejka, J., and Fitzmaurice, G. YouMove: enhancing movement training with an augmented reality mirror. In Proc. of UIST'13. 2013. ACM.
2. Benko, H. and Wilson, A.D. Pinch-the-sky dome: freehand multi-point interactions with immersive omnidirectional data. In Proc. of CHI'10 Extended Abstracts. 2010. ACM.

⁸ <https://www.lytro.com/camera/>

3. Bertalini, M., Mirrors and the mind. *The Psychologist*, 2010. 23(2): p. 112.
4. Berthaut, F., Marshall, M., Subramanian, S., and Hachet, M. Rouages: Revealing the Mechanisms of Digital Musical Instruments to the Audience. In *Proc. of NIME*. 2013.
5. Bimber, O. and Raskar, R., Spatial Augmented Reality: Merging Real and Virtual Worlds. 2005.
6. Butler, A., Hilliges, O., Izadi, S., Hodges, S., Molyneaux, D., Kim, D., and Kong, D. Vermeer: direct interaction with a 360 degree viewable 3D display. In *Proc. of UIST'11*. 2011. ACM.
7. Campbell, F.W., The depth of field of the human eye. *Journal of Modern Optics*, 1957. 4: p. 157.
8. De la Riviere, J.B., Dittlo, N., Emmanuel, O., Kervégant, C., and Courtois, M. Holocubtile: 3D multitouch brings the virtual world into the user's hands. In *Proc. of ITS'10*. 2010. ACM.
9. Eisert, P., Rurainsky, J., and Fechteler, P. Virtual Mirror: Real-Time Tracking of Shoes in Augmented Reality Environments. In *Proc. of ICIP'07*. 2007. IEEE.
10. Foley, J.D., van Dam, A., Feiner, S.K., and Hughes, J.F., *Computer graphics: principles and practice* (2nd ed.). 1990: Addison-Wesley Publishing Co., Inc. 1174.
11. Fujinami, K., Kawsar, F., and Nakajima, T. AwareMirror: a personalized display using a mirror. In *Proc. of 3rd international conference on Pervasive Computing*. 2005. Springer-Verlag.
12. Fujinami, K., Ushida, K., Tanaka, Y., Naemura, T., and Harahima, H. i-mirror: An Interaction/Information Environment Based on a Mirror Metaphor Aiming to Install into Our Life Space. In *Proc. of ICAT*. 2002.
13. Hachet, M., Bossavit, B., Cohe, A., and Riviere, J.P. Toucheo: multitouch and stereo combined in a seamless workspace. In *Proc. of UIST'11*. 2011. ACM.
14. Hilliges, O., Kim, D., Izadi, S., Weiss, M., and Wilson, A.D. HoloDesk: direct 3d interactions with a situated see-through display. In *Proc. of CHI'12*. 2012. ACM.
15. Hirsch, M., Izadi, S., Holtzman, H., and Raskar, R. 8D display: a relightable glasses-free 3D display. In *Proc. of SIGGRAPH'12 Posters*. 2012. ACM.
16. Hoffman, D.M., Girshick, A.R., Akeley, K., and Banks, M.S., Vergence-accommodation conflicts hinder visual performance and cause visual fatigue. *Journal of Vision*, 2008. 8(3): p. 33.
17. Holliman, N.S., Dodgson, N.A., Favalora, G.E., and Pockett, L., Three-Dimensional Displays: A Review and Applications Analysis. *IEEE Trans. on Broadcasting*, 2011. 57(2): p. 362-371.
18. Karnik, A., Mayol-Cuevas, W., and Subramanian, S. MUSTARD: a multi user see through AR display. In *Proc. of CHI'12*. 2012. ACM.
19. Korsch, D., *Reflective Optics*. 1991: Academic Press.
20. Krueger, M., *Artificial Reality II*. 1991: Addison Wesley
21. Li, W.H.A. and Fu, H. Augmented reflection of reality. In *Proc. of SIGGRAPH'12 ETech*. 2012. ACM.
22. Percival, A.S., *The prescribing of spectacles*. 1920: J. Wright.
23. Perlin, K., Paxia, S., and Kollin, J.S. An autostereoscopic display. In *Proc. of SIGGRAPH'00*. 2000. ACM.
24. Poston, T. and Serra, L. The virtual workbench: dextrous VR. In *Proc. of VRST'64*. 1994. World Scientific Publishing Co., Inc.
25. Raskar, R., Welch, G., Low, K.-L., and Bandyopadhyay, D. Shader Lamps: Animating Real Objects With Image-Based Illumination. In *Proc. of 12th Eurographics Workshop on Rendering Techniques*. 2001.
26. Regenbrecht, H.T., Franz, E.A., McGregor, G., Dixon, B.G., and Hoermann, S., Beyond the looking glass: Fooling the brain with the augmented mirror box. *Presence: Teleoper. Virtual Environ.*, 2011. 20(6).
27. Sato, H., Kitahara, I., and Ohta, Y., MR-Mirror: A Complex of Real and Virtual Mirrors, in *Virtual and Mixed Reality*, Shumaker, R., Editor. 2009, Springer Berlin Heidelberg. p. 482-491.
28. Sheard, C., The prescription of prisms. *American Journal of Optometry*, 1934. 11: p. 364.
29. Shibata, T., Kim, J., Hoffman, D.M., and Banks, M.S., The zone of comfort: Predicting visual discomfort with stereo displays. *Journal of Vision*, 2011. 8(11): p. 11.
30. Starch, D.A., Demonstration of the trial and error method of learning. *The Psychological Bulletin*, 1910. 5.
31. Stratton, G.M., A mirror pseudoscope and the limit of visible depth. *Psychological Review*, 1898. 5(6): p. 632.
32. Stratton, G.M., Upright Vision and the Retinal Image. *Psychological Review*, 1897. 4(2): p. 182.
33. Stratton, G.M., Vision without inversion of the retinal image. *Psychological Review*, 1897. 4(5): p. 463.
34. Sullivan, A. DepthCube solid-state 3D volumetric display. In *Proc. of SPIE*. 2004.
35. Valkov, D., Steinicke, F., Bruder, G., and Hinrichs, K. 2d touching of 3d stereoscopic objects. In *Proc. of CHI'11*. 2011. ACM.
36. Ware, C., Arthur, K., and Booth, K.S. Fish tank virtual reality. In *Proc. of CHI '93* 1993. ACM.
37. Yoshida, T., Shimizu, K., Kurogi, T., Kamuro, S., Minamizawa, K., Nii, H., and Tachi, S. RePro3D: full-parallax 3D display with haptic feedback using retro-reflective projection technology. In *Proc. of ISVRI'11*. 2011. IEEE