

Interactive Two-Sided Transparent Displays: Designing for Collaboration

Jiannan Li¹, Saul Greenberg¹, Ehud Sharlin¹, Joaquim Jorge²

¹Department of Computer Science
University of Calgary
2500 University Dr NW, Calgary, Canada
[jiannali, saul, ehud]@ucalgary.ca

²VIMMI / INESC-ID Instituto Superior Técnico
Universidade de Lisboa
Av. Rov. Pais, Portugal
jorgej@tecnico.ulisboa.pt

ABSTRACT

Transparent displays can serve as an important collaborative medium supporting face-to-face interactions over a shared visual work surface. Such displays enhance workspace awareness: when a person is working on one side of a transparent display, the person on the other side can see the other's body, hand gestures, gaze and what he or she is actually manipulating on the shared screen. Even so, we argue that designing such transparent displays must go beyond current offerings if it is to support collaboration. First, both sides of the display must accept interactive input, preferably by at least touch and / or pen, as that affords the ability for either person to directly interact with the workspace items. Second, and more controversially, both sides of the display must be able to present different content, albeit selectively. Third (and related to the second point), because screen contents and lighting can partially obscure what can be seen through the surface, the display should visually enhance the actions of the person on the other side to better support workspace awareness. We describe our prototype FACINGBOARD-2 system, where we concentrate on how its design supports these three collaborative requirements.

Author Keywords

Two-sided transparent displays; workspace awareness; collaborative systems.

ACM Classification Keywords

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

INTRODUCTION

Transparent displays are 'see-through' screens: a person can simultaneously view both the graphics on the screen and real-world content visible through the screen.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DIS 2014, June 21 - 25 2014, Vancouver, BC, Canada

Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2902-6/14/06...\$15.00.

<http://dx.doi.org/10.1145/2598510.2598518>

Transparent displays are now being explored for a variety of purposes. Commercial vendors, for example, are incorporating large transparent screens into display cases, where customers can read the promotional graphics on the screen while still viewing the showcased physical materials behind the display (e.g., for advertising, for museums, etc.). Researchers are promoting transparent displays in augmented reality applications, where graphics overlay and add information to what is seen through the screen at a particular moment in time. This includes how the real world is augmented when viewed through a mobile device [14, 1] or from the changing view perspectives that arise when people move around a fixed screen [15]. Commercial video visions of the future illustrate various other possibilities. 'A Day Made of Glass' by Corning Inc. [1], for example, illustrate a broad range of applications built upon display-enabled transparent glass in many different form factors, including: handheld phone and pad-sized devices; see-through workstation screens; touch-sensitive display mirrors where one can see one's reflection through the displayed graphics; interior wall-format displays, very large format exterior billboards and walls, interactive automotive photosensitive windows, two-sided collaborative walls (e.g., as in the mock-up of Figure 1), and others.

Our particular interest is in the use of transparent displays in face-to-face collaborative settings, such as in Corning Inc.'s scenario [1] portrayed in Figure 1. Such displays ostensibly provide two benefits 'for free': when a person is working on one side of a transparent screen, people on the other side of it can both see that person and what that



Figure 1. A mocked-up collaborative see-through display. Reproduced from [1]

person is working on. Technically, this is known as *workspace awareness*, defined as the up-to-the-moment understanding of another person's interaction with a shared workspace. As explained in [4], workspace awareness has many known benefits vital to effective collaborations (see §Related work below). While support for workspace awareness is well-studied in tabletop and wall displays, it is barely explored on transparent displays,

In this paper, we contribute to the design of transparent displays for collaborative purposes, thus adding to the repertoire of existing collaborative display mediums. Our goal is to devise a digital (and thus potentially more powerful) version of a conventional glass dry-erase board that currently allows people on either side to draw on the surface while seeing each other through it. As will be explained in a later section, such digital transparent displays have several basic design requirements that go well beyond current offerings if they are to truly support effective collaboration.

1. *Two-sided interactive input*. Both sides of the display must accept interactive input, preferably by at least touch and / or pen.
2. *Different content*. Both sides of the display must be able to present *different* content, albeit selectively.
3. *Augmenting human actions*. Because screen contents and lighting can partially obscure what can be seen through the display, the display should visually augment the actions of the person on the other side to make them more salient.

We begin with our intellectual foundation comprising the importance of workspace awareness, and how others have supported it using see-through displays. We then elaborate the above requirements of collaborative see-through displays, with emphasis on how they must support workspace awareness. This is followed by our implementation, where sufficient details are provided for the knowledgeable researcher to replicate our system. Our approach includes particular design features that address (at least partially) the above requirements.

RELATED WORK

Workspace awareness

When people work together over a shared visual workspace (a large sheet of paper, a whiteboard), they see both the contents and immediate changes that occur on that surface, as well as the fine-grained actions of people relative to that surface. This up-to-the-moment understanding of another person's interaction within a shared setting is the *workspace awareness* that feeds effective collaboration [4,6,5]. Workspace awareness provides knowledge about the 'who, what, where, when and why' questions whose answers inform people about the state of the changing environment: Who is working on the shared workspace? What is that person doing? What are they referring to? What objects are being manipulated? Where is that person specifically

working? How are they performing their actions? In turn, this knowledge of workspace artifacts and a person's actions comprise key elements of *situation awareness* (i.e., "knowing what is going on") [2] and *distributed cognition* [10] (i.e., how cognition and knowledge is distributed across individuals, objects, artefacts and tools in the environment during the performance of group work).

People achieve workspace awareness by seeing how the artifacts present within the workspace change as they are manipulated by others (called *feedthrough*), by hearing others talk about what they are doing and by watching the gestures that occur over the workspace (called *intentional communication*), and by monitoring information produced as a byproduct of people's bodies as they go about their activities (called *consequential communication*) [4].

Feedthrough and consequential communication occur naturally in the everyday world. When artifacts and actors are visible, both give off information as a byproduct of action that can be consumed by the watcher. People see others at full fidelity: thus consequential communication includes *gaze awareness*, where one person is aware of where the other is looking, and *visual evidence*, which confirms that an action requested by another person is understood by seeing that action performed.

Similarly, intentional communication involving the workspace is easy to achieve in our everyday world. It includes a broad class of gestures, such as *deixis* where a pointing action qualifies a verbal reference (e.g., 'this one here') and *demonstrations* where a person demonstrates actions over workspace objects. It also includes *outlouds*, where people verbally shadow their own actions, spoken to no one in particular but overheard to inform others as to what they are doing and why [4].

Gutwin and Greenberg [4] stress that workspace awareness plays a major role in various aspects of collaboration.

- *Managing coupling*. As people work, they often shift back and forth between loosely and tightly-coupled collaboration. Awareness helps people perform these transitions.
- *Simplification of communication*. Because people can see the non-verbal actions of others, dialogue length and complexity is reduced.
- *Coordination of action*. Fine-grained coordination is facilitated because one can see exactly what others are doing. This includes who accesses particular objects, handoffs, division of labor, how assistance is provided, and the interplay between peoples' actions as they pursue a simultaneous task.
- *Anticipation* occurs when people take action based on their expectations or predictions of what others will do. Consequential communication and outlouds play a large role in informing such predictions. Anticipation helps people either coordinate their actions, or repair undesired actions of others before they occur.

- *Assistance.* Awareness helps people determine when they can help others and what action is required. This includes assistance based on a momentary observation (e.g., to help someone if one observed the other having problems performing an action), as well as assistance based on a longer-term awareness of what the other person is trying to accomplish.

Our work builds upon Gutwin and Greenberg’s [4] workspace awareness theory. Our hypothesis is that our transparent two-sided display can naturally provide – with a little help – the support necessary for workspace awareness.

See-through displays in remote collaboration

In the late 1990s, various researchers in computer supported cooperative work (CSCW) focused their attention on how distance-separated people could work together over a shared digital workspace. In early systems, each person saw a shared digital canvas on their screen, where any editing actions made by either person would be visible within it. Yet this proved insufficient. Because some systems showed only the result of a series of editing actions, feedthrough was compromised. For example, if a person dragged an object from one place to another, the partner would just see it disappear from its old location and re-appear at its new location. Because the partner could not see the other person’s body, both consequential communication and intentional gestural communication was unavailable.

Some researchers tried to provide this missing information by building special purpose awareness widgets [e.g., 6], such as multiple cursors as a surrogate for gestural actions. Others sought a different strategy: a simulated ‘see-through’ display for remote interaction. The idea began with Tang and Minneman [18,19], who developed two video-based systems. VideoDraw [18] used two small horizontal displays, where video cameras captured and super-imposed peoples’ hands onto the display as they moved over the screen, as well as any drawing they made with marker pens. VideoWhiteBoard [19] used two wall-sized displays, where video cameras captured the silhouette of a person’s body and projected it as a shadow onto the other display wall.



Figure 2. Clearboard, with permission.

Ishii and Kobayashi [11] extended this idea to include digital media. They began with a series of prototypes based on “talking through and drawing on a big transparent glass board”, culminating in the Clearboard II system [11]. As illustrated in Figure 2, Clearboard II’s display incorporated both a pen-operated digital groupware paint system and an analog video feed that displayed the face, upper body and arms of the remote person. The illusion was that one could see the other *through* the screen. Importantly, Clearboard II was calibrated to support gaze awareness. VideoArms [17] and KinectArms [3] are both fully digital ‘mixed presence’ groupware systems that connect two large touch-sensitive surfaces, and include the digitally-captured images of multiple people working on either side. Because arm silhouettes were digitally captured, they could be redrawn on the remote display in various forms, ranging from realistic to abstract portrayals.

Similarly to the above efforts, our work tries to let a person ‘see through’ the display to the other side. It differs in that it is designed to support collocated rather than remote collaborations, as well as to address the nuances and limitations of see-through display technologies.

See-through two-sided transparent displays

Transparent displays are typically constructed by projecting images on translucent panels [15,9], or by using purposefully designed LCD/OLED displays [14,13]. Almost all displays are one-sided. That is, they display a single image on one side, where a person on the opposite side sees it as a reversed image (i.e., they see the ‘back’ of the image). Only a few allow direct interaction (e.g., via touch), but only on one side but not the other. Several notable exceptions are described below.

Hewlett-Packard recently received a patent describing a non-interactive see-through display that can present different visuals on each of its sides [12]. The display is composed of two separate sets of mechanical louvers, which can be adjusted so that observers can see through the spaces between them. At the same time, light can be directed on each set of louvers, thus presenting different visuals on each side. They envision several uses of their invention, but collaboration is not stressed.

Olwal et. al. [16] built FogScreenTM, an unusual see-through system whose screen uses vaporized water as display medium. Two projectors render images on both sides of the fog, which allows for “individual, yet coordinated imagery”. Input is done via 3DOF position tracking of LEDs held by people as tracked by IR cameras. Example uses of different imagery include rendering correctly oriented text and providing different information on either side, and to adapt content to particular viewing directions. However, they do not go into details.

In our own (unpublished) work in spring 2013, we transformed a Samsung transparent display into one that was fully interactive on both sides (Figure 3). We called it



Figure 3. FACINGBOARD-1, our earlier transparent display allowing for two-sided input (here, simultaneous collaborative drawing).

FACINGBOARD-1. Two Leap Motion controllers, one on each side, captured the gestures and touches of peoples' hands relative to the display. Thus people could interact simultaneously through it while at the same time seeing one another. However, both parties saw exactly the same image.

Heo et. al. [8] demonstrated TransWall, a high-quality see-through display that allows people on either side of it to interact via direct touch. It used two projectors to provide an identical bright image on both sides, and to minimize effects of image occlusion that may be caused by one person being in front of a projector. Projectors were calibrated to project precisely aligned images, where people saw exactly the same thing (thus one image would be the mirror image of the other)¹. Two infrared touch sensor frames mounted on either side collected multiple touch inputs per side. The system also included acoustic and vibro-tactile feedback, as well as a speaker/microphone that controlled the volume levels of the conversation passing through it.

Our work builds on the above, with notable differences. From a technical stance, we allow different images to be projected on either side, and both sides are fully interactive. From a collaborative stance, we focus on supporting workspace awareness within such see-through two-sided interactive displays, especially in cases where the ability to see through the display is compromised.

DESIGN RATIONAL FOR SEE-THROUGH TWO-SIDED INTERACTIVE DISPLAYS

Two-Sided Interactive Input.

Collaboration is central to our design. All people – regardless of what side they are on – are active participants.

¹ At the time of this paper's submission, TransWall author Lee told us they working on – but had not yet completed – a system that could project different images. We understand their work is now in submission. While FacingBoard-II predates their work, both should be considered as parallel independent efforts.

As with earlier systems supporting remote collaboration, we expect each person to be able to interact simultaneously with the display. From a workspace awareness perspective, we expect people to see each other through the screen and each other's effects on the displayed artefacts.

While such systems could be operated with a mouse or other indirect pointing device, our stance is that workspace awareness is best supported by direct interaction, e.g., by touch and gestures that people perform relative to the workspace as they are acting over it. Thus if people are able to see through the display, they can gather both consequential and intentional communications relative to the workspace, e.g., by seeing where others are touching, by observing gestures, by seeing movements of the hands and body, by noticing gaze awareness, by observing facial reactions.

Different Content on Both Sides

Excepting FogScreen™ vapour display [16], see-through displays universally show the exact same content on either side (albeit one side would be viewed in reverse). We argue for a different approach: while both sides of the display will mostly present the same content, different content should be allowed (albeit selectively) for a variety of reasons as listed below. Within CSCW, this is known as *relaxed WYSIWIS* (relaxed what-you-see-is-what-I-see).

Managing attenuation across the medium. Depending on the technology, image clarity can be compromised by the medium. For example, Olwal et al. [16] describe how their FogScreen™ diffuses light primarily in the forward-direction, making rear-projected imagery bright and front-projected imagery faint, thus requiring two projectors on either side. In our own experiences with a commercial transparent LED display (such as the one in Figure 3), image contrast was poor. One solution is to display content on both sides, rather than relying on the medium to transmit one-sided content through its semi-transparent material. This solution was adopted by Heo et. al. [8] in their TransWall system to maintain image brightness, where both projected images were precisely aligned to generate the illusion of a single common one-sided image.

Selective image reversal. Graphics displayed on a 'one-sided' traditional transparent display will appear mirror-reversed on the other side. While this is likely inconsequential for some applications, it can matter in others. This is especially true of reversed text (which affects readability), photos where orientation matters (maps, layouts, etc.), and of 3D objects (which will be seen from an incorrect perspective). The naïve approach, using two projectors, is to simply reverse one of the projected images, thus making them both identical from both viewers' perspectives. The problem is that the image components are no longer aligned with one another. This would severely compromise workspace awareness: a person's bodily actions as seen through the display will not be 'in sync' with the objects that the other person sees on his or her side.

A better solution applies image reversal selectively to small areas of the screen. For example, consider flipping blocks of text so that they are readable from both sides. If the text block is small (such as a textual label in a bounding box), it can be flipped within the bounding box while keeping that bounding box in exactly the same spot on either side. The same is true for any other small visuals, such as 3D objects. Thus touch manipulations, gestures and gaze made over that text or graphic block as a whole are preserved. However, it has limits: reversal may fail if a person is pinpointing a specific sub-area within the block, which becomes increasingly likely at larger reversed area sizes.

Personal work areas. Shared workspaces can include personal work areas. These are valuable for a variety of reasons. For one, they could collect individual tools that one person is using. During loosely coupled work, they could hold information that a person is gathering and working on, but that is not yet ready to show to others. They could even hold private information that one does not wish to share. A two-sided display allows for both shared and personal work areas. For example, an area of the screen (aligned to each other on either side) can be set aside as a personal work area, where the content on each side may differ. Workspace awareness is still partially supported: while one may not know exactly what the other is doing in their personal area, they will still be able to see that the other is working in that area.

Feedback vs. feedthrough. In many digital systems, people perform actions quite quickly (e.g., selecting a button). Feedback is tuned to be meaningful for the actor. For example, the brief change of a button's shading as it is being clicked or an object disappearing as it is being deleted suffices as the actor sees it as he or she performs the action. Alternately, pop-up menus, dialog boxes and other interaction widgets allow a person to perform extended interactions, where detailed feedback shows exactly where one is in that interaction sequence. Yet the same feedback may be problematic if used as feedthrough in workspace awareness settings [5]. The brief change of a button color or the object disappearing may be easily missed by the observer. Alternately, the extended graphics showing menus and dialog box interactions may be a distraction to the observer, who perhaps only needs to know what operation the other person is selecting. In remote groupware, Gutwin and Greenberg [5] advocated a variety of methods to portray different feedthrough vs feedback effects. Examples include making small actions more visible (e.g., by animations that exaggerate actions) and by making large distracting actions smaller (e.g., by showing a small representation indicating a menu item being selected, rather than the displaying the whole menu). The two sided display means that different feedback and feedthrough mechanisms can be tuned to their respective audience.

Personal state. Various widgets display their current state. Examples include checkboxes, radio buttons, palette

selections, contents of textboxes, etc. In groupware, each individual should be allowed to select these controls and see these states without affecting the other person, e.g., to select a drawing color from a palette. A two-sided relaxed WYSIWIS display allows a widget drawn at identical locations to show different states that depend upon which side it is on and how the person on that side interacted with it. For example, a color palette may show the currently selected color as 'blue' on one side, and 'orange' on the other.

Augmenting Human Actions.

Despite their names, transparent displays are not always transparent. They all require a critical tradeoff between the clarity of the graphics displayed on the screen vs. the clarity of what people can see through the screen. Factors that affect transparency include the following.

- **Graphics density and brightness.** A screen full of high-density and highly visible graphics compromises what others can see *through* those graphics. It is harder to see through cluttered (vs. sparse) graphics on a screen.
- **Screen materials.** Different screens comprise materials with quite different levels of transparency (or translucency).
- **Projector brightness.** If bright projector(s) are used, they can reflect back considerable light, affecting what people see through it. It is harder to see through screens with significant white (vs. dark) content.
- **Environmental lighting.** Glare on the screen as well as lighting on the other side of the screen can greatly affect what is visible through the screen. Similarly, differences in lighting on either side of the screen produces imbalances in what people see (e.g., consider a lit room with an exterior window at night time: those outside can see in, while those inside only see their own reflections).
- **Personal lighting.** If people on the other side of the display are brightly illuminated, they will be much more visible than if they were poorly lit.

To mitigate these problems, we suggest augmenting a person's actions with literal on-screen representations of those actions. Examples to be discussed in our own system include highlighting a person's fingertips (to support touch selections), and generating graphical traces that follow their movements (to support simple hand gestures).

THE DESIGN OF THE FACINGBOARD-2 SETUP

To our knowledge, no other transparent screen-based system offer a full range of two-sided interactive capabilities, including the ability to display different graphics on either side (but see [16]). Consequently we implemented our own display wall, called FACINGBOARD-2. Because it uses mostly off-the-shelf materials and technology, we believe that others can re-implement or vary its design with only modest effort as a DIY project.

Projector and Display Wall Setup

Figure 4 illustrates our setup. We attached fabric (described below) to a 57 cm by 36 cm aluminum frame. Two




Figure 4. The FACINGBOARD-2 Setup

projectors are mounted back-to-back above the frame along with mirrors, which affords different graphics per side, and which minimizes occlusion and glare through the screen. Projections are reflected through the mirrors at a downwards angle onto both sides of the fabric. A separate computer controls each projector, and both run our distributed FACINGBOARD-2 software that coordinates what is being displayed. Lighting is also controlled. Room light is kept low to minimize glare, while directional lights illuminate the people on either side.



Figure 5. Our open-weave projection screen

Projection Fabric

The most fundamental component of our system is a transparent display that could show independent content on either side. Most existing displays do not allow this. Current LED / OLED screens inherently display on one side. The various glass surfaces and/or films used in projection systems would not work well for two-sided projection, as the projected contents are designed with the goal of high-clarity bleed-through to the other side.

Instead, we explored materials comprising openly-woven but otherwise opaque materials (i.e., a grid of thread and holes) as a two-sided projection film. The idea is that these fabrics provide ‘mixed transparency’:

- images can be projected on both sides of the film, where the threads would reflect back and thus display the projected contents;
- a person could see through the holes in the open weave to the other side;
- bleedthrough would be mitigated if the thread material were truly opaque;
- while large solid displays can attenuate acoustics to the point that either side requires microphones / speakers [8], sound travels easily through openly-woven fabric.

Figure 5 illustrates how this works in FACINGBOARD-2. First, it shows the open weave of the fabric (the inset shows a close-up of it). Second, it shows the graphics (the ‘WallST’ photo) projected onto this facing side opaque weave. Third, it shows the person on the other side as seen through the fabric’s holes. Finally, it shows only minor bleed-through from the projection on the other side, visible as a slight greenish tint. This is caused by projected light from the other side bouncing off the horizontal thread surfaces, and because the fabric threads are not entirely opaque.

We used cheap and easily accessible materials: fabrics for semi-transparent window blinds that are woven out of wide, opaque threads forming relatively large holes. Choosing the correct blind material was an empirical exercise, as they vary considerably in the actual material used (some are not fully opaque), the thread color, the thread width, and the

hole size. Our investigation exposed the following factors as affecting our final choice of materials.

- 1) *Thread color.* Very dark (e.g., black) materials did not reflect the projected content well. This meant that any bleed-through would be more visible. Very light materials (e.g., white) reflected the projected content too well, where the brightness of the display limited how people could see through it.
- 2) *Thread width.* Wider threads reflect back more projected pixels and thus enhance display resolution. However, threads that are too wide also bounce light through to the other side (e.g., when the projection hits the top horizontal surface of the thread), which increases bleed-through.
- 3) *Hole size.* The holes must be large enough to let light pass through (thus ensuring transparency). However, holes that are too large compromise image fidelity.

After testing various materials, we chose the blind fabric seen in Figure 4: tobacco thread color, and 10% openness (a factor measuring the percentage of light penetration of blinds as determined by its thread width and hole size).

Input

Raw input is obtained from an off-the-shelf OptiTrack motion capture system. Eight motion capture cameras are positioned around the display (Figure 4). Participants on either side wear distinctive markers on their fingertip,

whose positions are tracked by the cameras and captured as 3D coordinates. The FACINGBOARD-2 software receives these coordinates and converts them into semantically meaningful units, e.g., as gestural mid-air finger movements relative to the display, and as touch actions directly on the display. Our current implementation is able to track separate finger motions on either side within a volume of at least 50 cm by 36 cm by 35 cm, and supports single touch point on each side. The software does not yet recognize one person's multi-touch, nor does it track other body parts (such as head orientation for approximating gaze awareness direction). This would be straightforward to do, and will be implemented in future versions.

We note that our choice of the OptiTracks motion capture system was driven by convenience: we had one, they are highly accurate, and they are reasonably easy to program. Other input technologies could be substituted instead. These include touch sensor frames (e.g., as used by [8]), or vision-based tracking systems (e.g., the Kinect), or 6 DOF input devices (e.g., Polhemus). All have their own particular set of advantages and disadvantages (e.g., marker-based or markerless, high or low accuracy, ability to detect and track in-air gestures in front of but not touching the screen).

Limitations and Practicalities

Our FACINGBOARD-2 setup works well as a prototyping platform, but still has a ways to go before it could be



Figure 6. The transparency of FACINGBOARD-2 as affected by various graphic density and lighting conditions.

considered a commercially deployable product.

First – and common across all transparent displays – the degree of transparency is greatly affected by various factors as already described in prior sections. Figure 6 illustrates how the transparency effect of FACINGBOARD-2 is affected by several of these factors (although due to limitations of photographing our setup, the transparency is actually better than what is shown in Figure 6). The best transparency is in Figure 6a, where projected graphics are sparse and the person on the other side is well lit. With denser graphics (6b) it is somewhat harder to see the person through it. If the other person is not lit, he can be even harder to see through either sparse (6c), or dense graphics (6d).

Second, the fabric used to construct FACINGBOARD-2 is not ideal. The threads are not particularly reflective, which means that the projected image is not of the brightness and quality one would expect of modern screens. As was seen in Figure 5, there is a very small amount of bleed-through of bright image portions to the other side. However, this is not noticeable if the other side also contains a brightly projected image. We believe better fabrics or screens could alleviate these limitations. One possibility is to paint a small grid or series of reflective opaque dots onto both sides of a thin transparent surface.

Third, as typical with all projection systems, image occlusion can occur when a person interposes part of their body between the projector and the fabric. We minimize occlusion by using downward-angled mirrors (Figure 4).

DESIGNING FACINGBOARD-2 RELAXED WYSIWIS


Our test-bed application is illustrated in Figure 7a: an interactive photo and text label manipulation. It includes a public area (top central), a private area (bottom), and a personal palette (left), all which will be discussed below. Because we had independent control of both input and output on either side, we were able to realize the various relaxed-WYSIWIS features as described in our Design Rational section.

Selective image and text reversal. As mentioned, graphics displayed on a ‘one-sided’ traditional transparent display will appear mirror-reversed on the other side. For example, Figure 7a shows one person’s view of the correctly oriented images and text in the public area, while in Figure 7b it appears in reverse to the person on the other side. We overcome this problem by selectively flipping images and text in place (Figure 7c). Each image and text block is precisely aligned to display at the exact same location on both sides, but its contents on one side are flipped to maintain the correct view orientation. Similarly, the text shown in the personal palette and private is flipped in place to make it readable on either side.

Personal work areas. While the public work area is visible to both people (albeit with flipped content), the contents of the private area are distinct to the viewer. For example, Figure 7a shows how Person 1 has 2 photos in his private



a) Person 1's view: photos / text correctly oriented



b) Person 2's view on the other side, showing how photos and text would normally appear as reversed



c) His relaxed-WYSIWIS view; text/photos unreversed

Figure 7. Relaxed WYSIWIS in FACINGBOARD-2.

area, while 7b,c shows how Person 2 has only 1 (different) photo. Each person can drag objects to / from their personal area, which causes them to disappear / reappear from the other person's view.

Semi-personal view of public objects. Each person is selectively able to modify the appearance of the text and images seen in the public view. Using the palette controls, they can reverse a selected object, add a red border to it, change the border thickness, as well as the background color of the text. These changes appear only on one side. For example, in Figure 7b, Person 2 has reversed his image as he wishes to point to fine details of it: this makes its contents identically aligned to what the other person sees. In Figure 7b,c, he has added a red border to an image and has colored a text object in orange, which differs from what Person 1 sees in Figure 7a.

Personal state. The palette controls, which are otherwise aligned on both sides, reflect their state on a personal basis, where selected radio buttons are shown in white. For example, we see in Figure 7b,c that Person 2 has selected the '4px' border thickness and 'Orange' border color, while in Figure 7a Person 1 has no options selected.

Feedthrough. When Person 1 selects a button in their personal palette, the button on Person 2's side animates for a few seconds longer than on Person 1's side. This enhances Person 2's awareness of Person 1's actions.

Augmenting human actions. As described above, the visibility of what a person sees through the medium can vary considerably. To mitigate this, we augment a person's actions with literal on-screen representations of those actions. Our initial work considers how mid-air finger movements and touches could be augmented. While simple, tracking fingers supports awareness of another's basic mid-air gestures made over a work surface (e.g., deixis and demonstrations), of intents to execute an action (e.g. a mid-air finger moving towards a screen object) and of actual actions performed on the display (e.g., touching to select and directly manipulate an object).

We enhance awareness by displaying a small visualization (a modest-sized dot) on the spot where the fingertip orthogonally projects onto the display. This dot only appears on the other side of the display, as it could otherwise mask the person's fine touch selections. For example, in Figure 7a Person 1 is touching a photo and no dot is visible. However, Person 2 sees the dot on their side (Figure 7b,c) Figure 8a-c shows how the actual size of the dot varies as a function of the distance between the fingertip and the display, i.e., the dot is small when the finger is far from the surface (8a), gets increasingly larger as the finger moves towards the surface (8b) and is at its largest when touching the surface (8c). When a touch occurs, the dot's color also changes.

We also use traces [7] to enhance gestural acts. As seen in Figure 8d, an ephemeral trail follows a person's finger



a) Tracking dot small to reflect distant finger



b) Tracking dot's size increases with approaching finger



c) Tracking dot at full size, color change indicating touch



d) Traces enhance gestural paths

Figure 8. Enhancing touch and gestural events. The person is on the other side of the screen.

motion, with its tail narrowing and fading over time. This enhances people's ability to follow gestures in cases where transparency is compromised (e.g., over dense graphics), as well as how people can interpret demonstration gestures.

DISCUSSION AND CONCLUSIONS

We are currently running a controlled study to investigate the effects on participants' performance when human actions are enhanced under different transparency conditions (such as those in Figure 6). We have several tentative findings. In poor transparency conditions without augmentation, participants said they could follow other's actions as long as they deliberately and consciously tried to do so. However, if participants were focused on other areas of the display (e.g., as in loosely coupled work), they had difficulty retaining their peripheral awareness of other's actions (which was not the case in high transparency situations). Thus our initial observations enforce our hypothesis that augmenting human actions is valuable, especially in low-transparency situations.

FACINGBOARD-2 is best seen as a design medium that allows designers to explore what is possible in a true two-sided interactive transparent display. Our particular motivation was to explore how it could best serve as a collaborative medium. We showed how the ability to project different graphics supports relaxed-WYIWIS, which in turn allows for selective image and text reversal, personal work areas, semi-personal views of public objects, personal state of controls, different feedback vs. feedthrough, and augmenting human actions via visuals. We also highlight some of the design tradeoffs entailed by face-to-face collaboration through an interactive semi-transparent medium, as well as limitations in our chosen materials. Even so, we expect advances in materials, technology and sensing will extend our ability to design interesting features and products in future two-sided mediums.

Our design iterations on two-sided collaborative displays has unearthed exciting possibilities. Yet we recognize that the present work is just the beginning of our explorations of what is possible on this medium. We are continuing our controlled study to understand both opportunities and limits in human performance. We are creating a suite of applications suitable for this medium. We are also elaborating on the various effects described in this paper.

Acknowledgements. Funds provided by NSERC-AITF-SMART Industrial Chair in Interactive Technologies, NSERC's Discovery Grant and Surfnet Network, and FCT grant CEDAR PTDC/EIA-EIA/116070/2009. Special thanks to Sutapa Dey who helped in our pilot studies.

REFERENCES

1. Corning, Inc. (2011, 2012) A Day Made of Glass I and II. Youtube: v=6Cf7IL_eZ38 and v=jZkHpNnXLB0. Retrieved December 31, 2013.
2. Endsley, M. (1995) Toward a Theory of Situation Awareness in Dynamic Systems. *Human Factors*, 37(1).
3. Genest, A., Gutwin, C., Tang, A., Kalyn, M. and Ivkovic, Z. (2013) KinectArms: a Toolkit for Capturing and Displaying Arm Embodiments in Distributed Tabletop Groupware. *Proc. ACM CSCW*, 157-166.
4. Gutwin, C. and Greenberg, S. (2002) A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *J. CSCW*, 11(3-4):411-446.
5. Gutwin, C. and Greenberg, S. (1998) Design for Individuals, Design for Groups: Tradeoffs between Power and Workspace Awareness. *Proc. ACM CSCW*.
6. Gutwin, C., Greenberg, S. and Roseman M. (1996) Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation. *Proc. HCI*, Springer, 281-298.
7. Gutwin, C. and Penner, R. (2002) Improving Interpretation of Remote Gestures with Telepointer Traces. *Proc. ACM CSCW*, 49-57.
8. Heo, H., Kim, S., Park, H., Chung, J., Lee, G. and Lee, W. (2013) TransWall. *ACM SIGGRAPH '13 Emerging Technologies*. Article No. 14.
9. Hirakawa, M. and Koike, S. (2004) A Collaborative Augmented Reality System using Transparent Display. *Proc. IEEE Multimedia Software Engineering*, 410-416.
10. Hollan, J., Hutchins, E. and Kirsh, D. (2000) Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research. *ACM TOCHI*, 7(2).
11. Ishii, H., and Kobayashi, M. (1992) ClearBoard: a Seamless Medium for Shared Drawing and Conversation with Eye Contact. *Proc. ACM CHI*, 525.
12. Kuo, H., Hubby, L. M., Naberhuis, S. and Birecki, H. (2013). See-through Display. US Patent 8,462,081 B2. Filed Mar 9. 2011. Issued Jun 11. 2013.
13. Lee, J., Olwal, A., Ishii, H. and Boulanger, C. (2013) SpaceTop: Integrating 2D and Spatial 3D Interactions in a See-through Desktop. *Proc. ACM CHI*, 189-192.
14. Li, J. Sharlin, E. Greenberg, S. and Rounding, M. (2013) Designing the Car iWindow: Exploring Interaction through Vehicle Side Windows. *Proc. ACM CHI Ext. Abstracts*, 1665-1670.
15. Olwal, A. Lindfors, C., Gustafsson, Kjellberg, T. and Mattsson, L. (2005) ASTOR: an Autostereoscopic Optical See-through Augmented Reality System. *Proc. IEEE Mixed and Augmented Reality*, 24-27.
16. Olwal, A., DiVerdi, S., Rakkolainen, I. and Hollerer, T. (2008) Consigalo: Multi-user Face-to-face Interaction on Immaterial Displays. *Proc. INTETAIN*, #8, ICST.
17. Tang, A., Boyle, M. and Greenberg, S. (2004). Display and Presence Disparity in Mixed Presence Groupware. *Proc. Australasian User Interface Conf* (Vol. 28. Australian Computer Society, Inc., 73-82.
18. Tang, J. and Minneman, S. (1990) Videodraw: A Video Interface for Collaborative Drawing. *Proc ACM CHI*.
19. Tang, J. and Minneman, S. (1991) VideoWhiteboard: Video Shadows to Support Remote Collaboration. *Proc. ACM CHI*, 315-322.