

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

1-2017

A two-sided collaborative transparent display supporting workspace awareness

Jiannan LI

Singapore Management University, jiannanli@smu.edu.sg

Saul GREENBERG

Ehud SHARLIN

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Graphics and Human Computer Interfaces Commons](#)

Citation

LI, Jiannan; GREENBERG, Saul; and SHARLIN, Ehud. A two-sided collaborative transparent display supporting workspace awareness. (2017). *International Journal of Human-Computer Studies*. 101, 23-44. Available at: https://ink.library.smu.edu.sg/sis_research/8063

This Journal Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.



A two-sided collaborative transparent display supporting workspace awareness



Jiannan Li, Saul Greenberg*, Ehud Sharlin

Department of Computer Science, University of Calgary, 2500 University Drive NW, Calgary, Alberta Canada T2N 1N4

ARTICLE INFO

Keywords:

Transparent displays
Workspace awareness
Collaborative systems

ABSTRACT

Transparent displays naturally support workspace awareness during face-to-face interactions. Viewers see another person's actions through the display: gestures, gaze, body movements, and what one is manipulating on the display. Yet we can design even better collaborative transparent displays. First, collaborators on either side should be able to directly interact with workspace objects. Second, and more controversially, both sides should be capable of presenting different content. This affords: reversal of images/text in place (so that people on both sides see objects correctly); personal and private territories aligned atop each other; and GUI objects that provide different visuals for feedthrough vs. feedback. Third, the display should visually enhance the gestural actions of the person on the other side to better support workspace awareness. We show how our FacingBoard-2 design supports these collaborative requirements, and confirm via a controlled study that visually enhancing gestures is effective under a range of deteriorating transparency conditions.

1. Introduction

Transparent displays are 'see-through' screens: a person can simultaneously view both the graphics on the screen and the real-world content visible through the screen. Our particular interest is how a transparent display can afford face-to-face collaboration between people situated on *opposite sides* of the screen. For example, consider the simple case of an off-the-shelf transparent display that allows touch interaction on one of its sides. If that display is positioned so that others can view its user through it, collaboration is afforded to some extent. Viewers can see that user's body movements, hand gestures, gaze, as well as what that user is actually manipulating on the display. Similarly, the user can see the viewers, as well as any gestures they make relative to their side of the display. This grounds awareness of mutual action as well as communication.

While an off-the-shelf transparent display affords the limited degree of collaboration as described above, we argue that transparent displays can provide even richer collaboration experiences if they were augmented with four particular features: allowing interactive input on both sides; allowing different content (albeit selectively) on either side; providing public, personal and private supporting the range of individual to group work; and visually augmenting human actions to make them more salient to viewers.

We will explain these ideas shortly. However, because the notion of transparent displays for collaboration is somewhat unusual and spec-

ulative, we begin by justifying why this is a fruitful research area worth pursuing.

1.1. The case for two-sided collaborative transparent displays

Almost all contemporary research on interactive surfaces for collocated collaboration situates people either side-by-side in front of a vertical display, or at various seating positions surrounding a horizontal tabletop display. Within this existing backdrop, it may seem unusual to suggest that collocated people may benefit from working on opposite sides of a single transparent display. Yet there are various reasons why such collaborative transparent displays should be added to our arsenal of techniques.

1.1.1. Reflects real-life practices

Collaborative transparent displays reflect real-life usage practices of people collaborating over glass. Dating back to the mid-20th century, for example, naval operators wrote field information (such as plotting ship direction) on both sides of glass plotting board, as illustrated in Fig. 1. This setup provided various advantages. Both operators had a clear view of the working area, as bodies were not in the way. It reduced interference between operators writing close to each other on the surface (as illustrated in Fig. 1). As operators could write on two sides of the glass, it doubled the space available for input.

* Correspondence to: 709 Larch Place Canmore, Canmore, AB, Canada T1W 1S2.

E-mail addresses: chrisleeseu@hotmail.com (J. Li), saul.greenberg@ucalgary.ca (S. Greenberg), ehud@cpsc.ucalgary.ca (E. Sharlin).




Fig. 1. Operators writing on both sides of a transparent plotting board. Source unknown.

1.1.2. Overcomes environmental separation

Collaborating through the display can overcome particular environmental constraints that require participants to be separated by a divider, i.e., where side by side collaboration is infeasible. For example, Corning Inc (2012) portrays a surgeon in a sterile operating room consulting with a distant colleague through a display wall (Fig. 2). However, we can easily imagine that colleague is standing in an adjacent non-sterile viewing room, where the wall between the rooms comprises display-enabled transparent glass. In this co-located situation, the surgeon can collaborate across this wall with his non-sterile colleague in the other room, where both can study and interact with the displayed medical imagery. Similarly, transparent displays can work as a collaborative yet protective barrier by people separated for security reasons, such as between prisoners/visitors in a jail, between clerks/customers in a bank or jewelry store, and between a taxi driver and her back-seat customers.

1.1.3. Supports opportunistic casual interaction

Transparent displays readily support awareness leading to casual interactions. For example, many contemporary envisionments about near-future work involving a team of collocated people depict various team members working behind transparent displays of various sizes (Shedroff and Noessel, 2012). Co-workers get a sense of what others are doing as they glance around, as they can see the worker's face and hands through the screen as well as what they are working on. In turn, this increases overall situation awareness and creates opportunities for co-workers to interact. An example is one worker noticing another having difficulty with their on-screen work, and coming to their assistance.



Fig. 2. A mock-up scenario showing a surgeon in the sterile operation room asking for advice from his colleague in the other non-sterile room, while studying medical imagery displayed on the transparent wall between them. Source: Corning Incorporated (2012), with permission.

1.1.4. Supports the switch between individual and joint work across desk partitions

If the display can be switched between opaque and transparent modes, it could be used by co-located workers to rapidly switch between individual and joint work across desk partitions. To explain, Danninger et al. (2005) created an LCD glass partition separating the abutting desks of two office workers. To minimize distraction and safeguard privacy, the glass was fully opaque when both were turned away from it. However, if one co-worker knocked on the glass and the other turned to face it, the glass became fully transparent to afford face to face conversation. If this glass was replaced by an interactive display that allowed both opaque and transparent settings (Lindlbauer et al., 2014a, 2014b; Li et al., 2014), that same partition could afford individual work in opaque mode (each working on their own side), and shared work in transparent mode (both working over the common work surface visible to both).

1.1.5. Supports true face to face interaction

A fifth opportunity is suggested by gaming. Console games using vertical displays currently require its players to be in front of the display, where they usually stand or sit side by side. Yet certain console games involve activities normally done through direct face to face play, where the scene and the other person are simultaneously in view (e.g., boxing and tennis games). Games designed for a collaborative transparent display could thus allow players to directly face each other, giving an entirely different feel to game play. This benefit could be applied to any situation where true face to face interaction is desired. In contrast, tabletop and non-transparent vertical displays require participants to either look at the surface or at each other (when face to face) and/or to assume alternate positions (e.g., side to side).

We are not suggesting that collaborative transparent displays should supplant existing digital surface technologies. Indeed, we believe that tabletops and non-transparent wall displays will remain appropriate for a large majority of common situations. Rather, we see collaborative transparent displays as an addition to the repertoire of available surface types, where they are a good match to particular situations such as the samples listed above. We are not the only ones holding this view, as a small community of other researchers are actively researching collaborative transparent displays (e.g., Olwal et al., 2006, 2008; Heo et al., 2013; Kuo et al., 2013; Lee et al., 2014; Li et al., 2014; Lindlbauer et al., 2014a, 2014b).

1.2. Structure of the Paper

In this paper,¹ we contribute to the design of transparent displays supporting collocated collaboration, thus adding to the repertoire of existing collaborative display mediums. Our goal is to elaborate upon 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 (e.g., contrast Fig. 1 with Fig. 2). Our methodology (and the paper structure) roughly follows a multi-step process as detailed below, each offering a particular contribution.

First, we lay the theoretical foundation – drawn from related work – that we use to motivate our design ideas (§2). We know from prior work that seeing the displayed artifacts in the workspace, along with people's bodily actions relative to the artifacts, is critical for efficient collaborative interaction, as it helps communicate and coordinate mutual understanding. This is known as *workspace awareness*, defined as the “up-to-the-moment understanding of another person's interaction with a shared workspace” (Gutwin and Greenberg, 2002). We also

¹ This paper reflects a complete archival report of our multi-year project on collaborative transparent displays. The first part – our theoretical foundation, implementation and related work – expands considerably upon the initial work reported in (Li et al., 2014). The second part – the study – has not been previously published.

know that people tend to tacitly partition a shared workspace into various areas, each with their own utility, e.g., public, personal, and private (Scott et al., 2004; Scott and Carpendale et al., 2010). This is known as *territoriality*. While support for workspace awareness and territoriality is well-studied in tabletop and wall displays, it has not been applied to transparent displays. We thus begin with our intellectual foundation comprising the importance of workspace awareness and territoriality. Later sections elaborate these theories as requirements for collaborative see-through displays.

Second, we briefly survey in §3 related technologies that use a see-through display metaphor. We will see how the see-through display metaphor, along with the theories of workspace awareness and territoriality, has been applied to groupware for distance-separated collaborators. Our work differs in that we focus on collocated rather than remote collaborations. We will also see that several others have built fully interactive collaborative transparent displays along with a few (mostly playful) demonstration applications. Our work builds on those efforts, but with notable differences: our technical infrastructure is novel; we use theory to develop a design rationale and to engineer generalizable interaction techniques; we also identify, study and mitigate problematic situations where transparency is compromised.

Third, we elaborate upon our theoretical foundation to develop requirements for collaborative see-through displays (§4). We will see that such displays have several basic design requirements that go well beyond current transparent display offerings if they are to truly support rich collaboration.

1. *Interactive input on both sides.* Both sides of the display should accept interactive input, preferably by at least touch and / or pen.
2. *Different content.* Both sides of the display should be able to present different content, albeit selectively, while still aligning content across the sides as needed.
3. *Public, personal and private areas.* Although somewhat application-dependent, particular areas of the display should be reserved as territories specifically supporting individual *vs.* group activities.
4. *Augmenting human actions.* If screen contents, lighting and other factors 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.

Within this context, we now define a *two-sided transparent display* as a system that affords interactive input on both sides (point 1), and that is capable of displaying different content (point 2), which in turn makes points 3 and 4 technically feasible.

Fourth, we operationalize these requirements through our implementation of a collaborative transparent display called FACINGBOARD-2. We provide sufficient details of our infrastructure setup (§5) and our test bed application (§6) for the knowledgeable researcher to replicate our system.

Fifth, we revisit what we believe to be a basic design problem with transparent displays, hinted at in point 4 above. Our experiences with both our own and other transparent displays revealed a critical problem: in spite of their name, transparent displays are not always transparent. All trade off the clarity of the graphics displayed on the screen *vs.* the clarity of what people can see through the screen. This compromises what people can see and can severely affect workspace awareness. To mitigate this, we created two methods that track and visually augment human actions. *Touch augmentation* highlights a fingertip with a circular glow that increases in size and intensity during approach, and that changes color upon touch (Fig. 3, top). *Trace augmentation* (Fig. 3, bottom) creates a fading trace that follows the motion of the fingertip (Gutwin, 2002; Gutwin and Penner, 2002). The question is, are these augmentation techniques effective in supporting workspace awareness under degrading transparent display conditions? To answer this question, we conducted a controlled study that investigated how people performed various collaborative tasks while




Fig. 3. Touch *vs.* trace augmentation. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

varying transparency and the augmentation techniques available (§7 and §8). This is followed by several implications that should be considered by both researchers and practitioners (§9).

2. Related work I: theoretical foundations

We see collaborative transparent displays as providing one type of a shared digitally-enabled workspace to the people gathered around it. Because shared workspaces in general are well-researched in computer-supported cooperative work (CSCW), we review two theoretical constructs that we believe are important to the design of collaborative transparent displays: workspace awareness, and territoriality.

2.1. Workspace awareness

In our everyday activities, people naturally stay aware of their surrounding environments and respond accordingly. Human factors research studied how this knowledge of the changing environment – termed *situation awareness* – was availed in highly dynamic and information-rich environments, such as air combat. Situation awareness is described as “knowing what is going on”, where it comprised three key components: the perception of the element within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future (Endsley, 1995).

Researchers in the CSCW community developed a similar concept of awareness involving knowledge of both individual and group activity, information sharing, and coordination in a shared workspace (Dourish and Bellotti, 1992). In particular, 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 (Gutwin and Greenberg, 2002, 1998; Gutwin et al., 1996). 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 not only situation awareness (Endsley, 1995), but distributed cognition (i.e., how cognition and knowledge is distributed across individuals, objects, artefacts and tools in the

environment during the performance of group work, see Hollan et al., 2000).

People achieve workspace awareness through various means (Gutwin and Greenberg, 2002). Using *feedthrough*, they see how the artifacts present within the workspace change as they are manipulated by others. Using *intentional communication*, they hear others talk to them about what they are doing, and they see the communicative gestures others perform over the workspace. Using *consequential communication*, they monitor information produced as a by-product of people's bodies as they go about their activities.

Feedthrough and consequential communication occur naturally in the everyday world. When artifacts and actors are visible, both give off information as a by-product 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* that confirms that an action requested by another person is understood by seeing that action performed. The visibility of gestures also play an important role, where Reetz and Gutwin (2014) found that both large and small gestures form a very observable component of consequential communication.

Similarly, intentional communication involving the workspace is easy to achieve in our everyday world. It includes a broad class of gestures. One example is *deixis*, where a pointing action qualifies a verbal reference (e.g., 'this one here'). Another example is *demonstrations*, where a person demonstrates actions over workspace objects. Intentional communication 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 (Gutwin and Greenberg, 2002).

Gutwin and Greenberg (2002) 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-coupled and tightly-coupled collaboration. Awareness helps people perform these transitions. While a person's focus of attention during loosely-coupled work is primarily on individual work, that person still monitors others' activities to stay aware of opportunities to move into tightly-coupled highly collaborative work.
- *Simplification of communication*. Because people can see the non-verbal actions of others, dialogue length and complexity is reduced (Clark, 1996).
- *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., 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 transparent display design rationale (§4) and our system (§5, §6) build upon Gutwin and Greenberg's (2002) workspace awareness theory. Our hypothesis is that a transparent two-sided display can naturally provide – with a little help – the support necessary for people on either side to maintain workspace awareness. This happens because each can see each other's actions through the workspace relative to the displayed objects (e.g., see Figs. 1 and 2). In §3, we will also review how

these workspace awareness constructs were realized in several types of groupware systems involving a shared workspace, ranging from remote collaboration systems using a see-through display metaphor, to collocated collaboration systems that allowed people to interact on either side of a transparent display.

2.2. Territoriality

Territoriality theory describes how group members partition the shared workspace into zones (areas) of different uses. During collaborative activities, people often use zones located at different positions in the workspace for different purposes. Generally, these zones allow for efficient usage of space (Tang, 1991). For example, at small distances from a workspace area (e.g., meters), zones are equated to social protocols about interpersonal proxemics (Hall, 1966): essentially, the closer one is to a workspace area, the more that area becomes one's own (Vogel and Balakrishnan, 2004). When people surround a workspace, such as in tabletop collaboration, three types of territories can arise (Scott et al., 2004) – personal, public, and storage. Each territory, which may be explicit or tacit, has distinct spatial and functional properties. A *personal territory* is typically one that proximately surrounds the person, and is reserved by that person for his/her individual work. This territory is visible but not accessible to others for the most of the time. A *public territory* is the area where group members share access, usually to collectively pursue the main collaborative task. It often takes up the space that is not occupied by other territories. A *storage territory* serves as the area to store task resources and typically sits atop both personal and public territories. Similar territorial partitions of personal vs. public areas can also be found on vertical workspaces (Azad et al., 2012).

Another type of territory in shared workspaces is the *private territory*, such as the private notebook of a group member. Comparing with personal territories, they ensure a higher level of privacy: neither publicly modifiable nor visible. This distinction between personal and private is important. Early groupware did seek to accommodate and further enforce people's partitioning behavior. One example defines fine-grained access levels on private vs. public objects via what is called user interface coupling (Dewan and Choudhary, 1991), where the coupling level is used to control what particular users see on their display. Another example separates private vs. public territories by device. Private territories are displayed on personal devices (e.g. PDAs and laptops), while public territories are displayed on a shared public workspace (e.g., a table or wall display) (Rekimoto and Saitoh, 1999). The owners of the personal device could see and manipulate objects in the private territory, or transfer objects from their territory to the public territory. However, this binary partition left no room for personal territories, which are only exclusive in terms of access, not of visibility. The visibility of others' personal territories is often critical to group work, as people monitor the activities in these territories to know others' states (Scott et al., 2004) and maintain consequential communication. Later groupware designers paid particular attention to the subtle distinction between private, personal, and public territories. For example, Wu and Balakrishnan's (2003) RoomPlanner had no permanent private territories. However, if a person placed the side of his or her hand on the tabletop to block others from seeing the area behind it, the system recognized that as a gesture that trigger the display of private information. UbiTable by Shen et al. (2003) went even further by providing designated private, personal, and public territories. Like Rekimoto and Saitoh (1999), private territories were workspaces on individuals' laptops. Personal territories covered areas on the tabletop that were close to each group member, visible but not modifiable to others. Public territories were centered within the tabletop, and were shared by all group members.

Territories such as these are important. To quote from Scott et al.'s discussion of territories on tabletops:

“... territories facilitate collaborative interactions on a table by providing commonly understood social protocols that help people to share a tabletop workspace by clarifying which regions are available for individual or joint task work, to delegate task responsibilities, to coordinate access to task resources by providing lightweight mechanisms to reserve and share task resources, and to organize the task resources in the workspace.” (Scott et al., 2004).

The above work suggests that transparent displays can facilitate certain types of collaboration, by including territories with different levels of accessibility and visibility. As we will see, our design rationale recommends such partitioning on collaborative transparent displays. This is also realized in our collaborative transparent display FacingBoard-2, which includes not only public areas for group work, but private storage areas and semi-personal tool palettes, each aligned to appear atop each other in the same location on either side of the display. These will be explained shortly.

3. Related work II: the use of transparency in collaboration

There is a history of work related to the use of transparency, and to the use of transparency in collaboration. We begin with a brief summary of transparent displays in general. We then describe how the see-through display metaphor has been applied to groupware systems supporting remote collaboration. We close by detailing the (few) examples of transparent displays specifically designed to support collaborative work.

3.1. Transparency and transparent displays

Transparency has a good history in graphical user interface design, particularly of layering user interface objects (windows, menus, dialog boxes, etc) over background screen contents. Harrison et. al. (1995) showed that users interacting with semi-transparent user interface objects benefit by staying aware of the screen contents under those objects. Baudisch and Gutwin (2004) improved the readability of text present in either layer through a transparency mechanism called multiblending. Others have considered how transparency in see-through displays (including augmented reality glasses and transparent displays) can be improved, such as by color correction (Sridharan et al., 2013), and transparency level and contrast (Joung et al., 2016).

In spite of the interest in transparency, transparent display hardware is still under active development. Most are either self-contained display panels or projection-based systems. Commercial transparent display panels are typically built upon LCD liquid-crystal or OLED organic light-emitting diode technologies (e.g., Samsung, 2014; Planar Systems, Inc, 2014), with some companies exploring monochromatic transparent displays using liquid crystal or electroluminescent display technology (e.g. Lumineq, 2014; Kent Optonics, 2014). In contrast, projection systems use a projector to project an image onto a material that is both see-through and reflective. Materials are usually special films overlaid onto glass (e.g., Pronova, 2015). However, because projection films may compromise display transparency to achieve image brightness, researchers in material science are actively producing special materials that can achieve a better transparency/image brightness tradeoff (e.g., Sun and Liu, 2006; Downing et al., 1996; Hsu et al., 2014). Artists have also projected images onto translucent fabric (called scrim), so that viewers at an exhibition can see its contents from either side (Wikipedia, 2015). One unusual projection-based system rear-projects images onto a thin plane of water vapor (fog) to create an *immaterial* or *mid-air display* that can be reached through and walked through (Olwal et al., 2006, 2008).

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 (Lee et al., 2013; Li et al., 2013; Corning Inc., 2012) or from the changing view perspectives that arise when people move around a fixed screen (Olwal et al., 2005). Commercial video visions of the future illustrate various other possibilities. ‘A Day Made of Glass’ by Corning Inc. (2012), 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, and others. Others also considered how people working with a transparent vs. conventional display maintain better awareness of what is going on outside the display space (i.e., in the background) (Lindlbauer et al., 2016). Our own interest, however, lies in how transparent displays can be used in collocated collaboration.

3.2. See-through display metaphors in distance-separated collaboration

In the late 1990s, various researchers in 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. Similarly, spoken references by the actor to the action as it was being performed would be much harder to understand.

Some researchers tried to provide this missing information by building special purpose awareness widgets (e.g., Gutwin, Greenberg and Roseman, 1996), 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, (1990, 1991), who developed two video-based systems. VideoDraw (Tang and Minneman, 1990) 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 (Tang and Minneman, 1991) 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. Ishii and Kobayashi (1992) 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 (Ishii and Kobayashi, 1992). As illustrated in Fig. 4, 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 (Tang et al., 2004) and KinectArms (Genest et al., 2013) are both fully digital ‘mixed presence’ groupware system 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.




Fig. 4. Clearboard, with permission from Hiroshi Ishii.

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. We stress that the collocated situation is very different from remote situation. While it is technically possible to use some of the above remote collaboration technologies to support collocated interaction (e.g., to project video into a non-transparent display rather than use a transparent display), using true transparency is a much simpler solution: the real world visible through the display does not have to be digitally replicated. As a result, many of the limitations in the above digital techniques disappear, e.g., calibration issues in maintaining eye contact, true 3D allowing looking around *vs.* tracking head movements to adjust the perspective view (as done in fishbowl VR), potentially better resolution (as one can see the real world rather than a reconstructed world), latency, etc. In addition, the working mode is quite different. Unlike physical transparent displays, systems like Clearboard, VideoDraw and VideoArms require at least two physical displays, with each collaborator working behind their display. This configuration can be unwieldy or impractical in collocated spaces (e.g., two display walls would be required). Alternately, the displays would have to be reworked to provide the illusion that they are a single see-through display, e.g., by placing them back to back.

3.3. Two-sided transparent displays

We have argued that a truly collaborative transparent display requires at least two features beyond conventional transparent displays. First, it must allow for people on either side of the display to interact simultaneously with the displayed graphics while still allowing them to see one another. Second, it ideally allows different content to be selectively projected on either side.

Speaking to the first point, most interactive transparent display systems only recognize the actions of one (but sometimes more) people standing on one side of the display. Still, there are a few instances of two-sided interactivity, typically implemented by using a variety of existing technologies. For example, FacingBoard-1 used two Leap Motion controllers, one per side, to capture the gestures and touches of peoples’ hands relative to the display (Li, 2015). This is illustrated in Fig. 5, where we see two people collaboratively moving a graphical object (a line). The Consigalo FogScreen™ system used IR trackers that track the 3D positions of up to eight IR LEDs placed on objects held by the various participants (Fig. 6) (Olwal et al., 2008). FogScreen™ also provided further control options by augmenting interaction with a wireless joystick held by the user. TransWall used two infrared touch sensor frames mounted on either side to collect multiple touch inputs per side (Fig. 7) (Heo et al., 2013). It also




Fig. 5. FacingBoard-1, our earlier transparent display allowing for two-sided input (here, simultaneous collaborative drawing) (Li, 2015).



Fig. 6. Consigalo using FogScreen™ (Olwal et al., 2008). With permission, A. Olwal.




Fig. 7. TransWall, a projection-based transparent display. The content on both sides was the same. (Heo et al., 2013). With permission from Woohun Lee.

included acoustic and vibro-tactile feedback, as well as a speaker/microphone that controlled the volume levels of the conversation passing through it.

Second, most transparent displays are currently ‘one-sided’: they display a single image on one side, which the person on the opposite side sees in reverse. Only a very few systems display different content on either side. For example, Hewlett-Packard described a non-interactive see-through display composed of two separate sets of mechanical louvers, which can be adjusted so that observers can see through the

spaces between them (Kuo et al., 2013). At the same time, light can be directed on each set of louvers, thus presenting different visuals on each side. While they envision several uses of their invention, collaboration is not stressed.

Heo et al. (2013) demonstrated TransWall, a high-quality see-through display that allows people on either side of it to interact via direct touch. It is notable here as it uses two projectors on either side (Fig. 7). However, its purpose was to provide an *identical* image on both sides, thereby increasing brightness while minimizing 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 reversed mirror image of the other, as with conventional transparent displays).

FogScreenTM is an immaterial see-through system whose ‘screen’ comprises a thin plane of vaporized water (Fig. 6) (Olwal et al., 2006, 2008) that people can walk through. Its researchers adapted it to implement Consigalo, a multi-user gaming system that can display different content on both sides of the fog, which allows for “individual, yet coordinated imagery” (Olwal et al., 2008). 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 (e.g., showing the back or front of a 3D object on either display side). However, they report that FogScreen’s image quality is relatively poor compared to normal displays.

JANUS is an unusual transparent display that shows different content on its two sides by taking advantage of persistence-of-vision (POV) effects (Lee et al., 2014). It displayed graphics by spinning a blade with an array of tri-color LEDs on each side at a high speed (Fig. 8). The graphics shown on the two sides were independent as the blade was opaque and the two LED arrays responded to separate input signals. As an early research prototype, its limitations include low-resolution, limited display area (the movement range of the blade), and cumbersome hardware.

The Tracs system also deserves mention, for it is the only two-sided collaborative transparent display (albeit with a twist) that includes some notion of territoriality (Lindlbauer et al., 2014a, 2014b). Its display comprises several sandwiched layers: two transparent LCD screens, and a backlit ‘transparency-control layer’ that can be made opaque or transparent. Using this hardware, users can selectively control whether the screen or particular screen regions are non-transparent (each person can only see the contents on their side, i.e., as a private territory), semi-transparent (where people can see through the displayed contents, which are visible to both, i.e. as a public territory), or fully transparent (the contents are hidden but the people

are clearly visible through it). Thus Tracs affords a quite different solution to territories on a two-sided collaborative display, where it dynamically partitions the screen into transparent and non-transparent regions to support both collaborative (group) and individual (private) work.

Our work builds on all the above, with notable differences. We are closest to Consigalo (Olwal et al., 2008) and Janus (Lee et al., 2014): they are the only other transparent display systems that fully allow for different content per side, and where both sides are interactive.² However, those works primarily focused on technical implementation aspects along with proof-of-concept demonstrations involving a few simple (mostly playful) applications. The work we report here—while also contributing technical innovations and benefits (such as improved resolution)—is based on a broader frame of reference. From a collaborative stance, we focus on supporting workspace awareness and territoriality to motivate the design of see-through two-sided interactive displays and interaction techniques. We are especially concerned about situations where the ability for collaborators to see through the display is compromised, where we developed and studied the effectiveness of augmentation techniques to overcome workspace awareness loss.

4. Design rationale for a see-through two-sided interactive display

We previously defined a two-sided transparent display as a system that affords interactive input on both sides, and that is capable of displaying different content. We argue why these capabilities are desired, and how they can be used to develop a myriad of techniques beneficial to collaboration.

4.1. Two-sided interactive input

Collaboration is central to our design thinking. All people – regardless of what side of the transparent display they are on – are considered active participants, where each person can interact simultaneously with the display. From a workspace awareness perspective, we expect people to see each other through the screen, each other’s actions relative to the displayed artefacts, and the effects of those actions on those artefacts. From a territorial perspective, we expect collaborators to have a public area for joint activity, and (depending upon the need) a personal or private area for individual activities.

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. In contrast to small mouse movements, people are able to see body movements through a transparent display. They can thus gather both consequential and intentional communications relative to the workspace, for example, by seeing where others are touching, by observing gestures, by seeing movements of the hands and body, by noticing gaze awareness, and by observing facial reactions.

4.2. Different content on both sides

Excepting a few systems (Olwal et al., 2006; Lee et al., 2014; Lindlbauer et al., 2014a, 2014b), see-through displays universally show the exact same content on either side, although one side would be viewed in reverse. This is called WYSIWIS (what-you-see-is-what-I-see). We argue for a different approach: while both sides of the display




Fig. 8. JANUS, a two-sided emissive transparent display making use of POV effect (Lee et al., 2014). With permission from Woohun Lee.

² In publication order, Consigalo (Olwal et al., 2008) is, to our knowledge, the first two-sided collaborative transparent display system. Second is FacingBoard-2 (Li et al., 2014), followed a few months later after by Janus (Lee et al., 2014) and Tracs (Lindlbauer et al., 2014a; 2014b). These last three systems should be considered contemporaneous research efforts, indicating increased interest in the field.



Fig. 9. The naïve two-projector solution, with unaligned graphics and bleed-through.

will mostly present the same content, different content should be allowed (albeit selectively). This also implies that bleed-through of displayed images from one side to the other is somehow mitigated, as the different content would otherwise create visual noise and interference. Within CSCW, allowing collaborators to mostly see the same thing while still providing for different views is known as *relaxed WYSIWIS* (relaxed what-you-see-is-what-I-see) (Stefik et al., 1987). A variety of reasons supporting different content on both sides are listed below.

4.2.1. 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 for various data abstractions such as text (where reversal affects readability), images such as maps, schematics and blueprints (where orientation matters), and of 3D objects (which will be seen from an incorrect perspective). Unfortunately, the naïve solution of using a projector on each side of the screen to display correctly oriented graphics does not work, as illustrated in Fig. 9. First, the flipped screen images on either side would be severely out of alignment with one another. In Fig. 9, for example, we see that the ‘ABC’ text block on the front left is located horizontally opposite to it on the back. This non-alignment would severely compromise workspace awareness, as a person’s bodily actions as seen through the display will be out of sync with the objects that the other person sees on his or her side (e.g., in Fig. 9 the viewer sees the person’s pointing gesture to an empty area rather than to the ‘ABC’ text block). Another issue is that, in most transparent displays, this non-alignment of graphical objects would create significant visual interference because of bleed-through effects. Bleed-through is also illustrated in Fig. 9 as the greyer image-reversed CBA text block.

We believe that a better – albeit limited – solution applies image reversal selectively to small areas of the screen, while still controlling for bleed-through. For example, consider a screen containing independent blocks of text. If each text block is flipped in place, they would be readable from both sides. If the text block is small (such as a textual label in a bounding box), it can be flipped within its bounding box while keeping that bounding box in exactly the same spot on either side. The same solution can be applied to any other modest-sized visual, such as photos. Similarly, 3D objects can be displayed from their correct perspective, where the true front and back sides of that object are shown aligned on the front and back of the two-sided display (Olwal et al., 2006, 2008). Touch manipulations, gestures and gaze referring to that text or graphical block as a whole are preserved, thus maintaining workspace awareness.

There are limitations. First, this approach does not work for large or full-screen graphics, e.g., a map whose size requires filling the entire display, as gestural references will be grossly unaligned to the graphics shown on both sides. Second, workspace awareness can be compromised if a person is pinpointing a specific sub-area within a block (e.g., a particular word in a text block), as the graphics under the gesture would not be aligned with its counterpart on the other side. This is why we advocate for small blocks, as within-object gestures would be increasingly likely as the block size increases.

4.2.2. Creation of distinct territories

According to territoriality theory, people using a shared visual workspace may require various types of territories, including public, storage, personal and private work areas. These are valuable for a variety of reasons. The public territory should be one held and clearly seen by the group, where it affords joint interactions and clear workspace awareness so all can see what others are doing. Personal territories could collect individual objects and tools that one person is working with or storing, which may differ from another person’s objects and tools. Private territories could hold private information and hide actions that should not be visible to others.

A two-sided display allows for all these work areas. Broadly speaking, we see public territories on such a display as those WYSIWIS regions that include objects that are clearly visible and accessible to all. While objects may be flipped (see previous requirement), they would be visually aligned to appear in the same spot on either side, where people’s actions relative to those objects are easily perceived. In contrast, personal and private work territories are defined areas of the screen that implement relaxed-WYSIWIS. While these territories are aligned to each other on either side, the content on each side may differ substantially (e.g., each may hold tools and objects particular to the individual). Workspace awareness can still be partially supported to varying extents: while one may not know exactly what the other is doing in their personal territory, they will still be able to see that the other is working in that aligned area through their bodily actions.

4.2.3. 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. An example is the brief change of a button’s shading as it is being clicked, or an object immediately disappearing as it is being deleted. This feedback 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 (Gutwin and Greenberg, 1998). 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 rather than the details of that operation. In remote groupware, Gutwin and Greenberg (1998) 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 viewer. In essence, each control or object – likely aligned to the same location on either side of the display – can behave like a mini-personal territory to implement relaxed-WYSIWIS, where it displays differing feedback (to the person doing the action on one side) vs. feedthrough (to the

person viewing the action on the other side).

4.2.4. Personal state

Various interactive objects display their current state. Examples include checkboxes, radio buttons, palette selections, contents of textboxes, etc. In groupware, these objects may be ‘owned’ by individuals, where setting them creates a personal state. An example is a groupware drawing system, where individuals can select their own drawing color by choosing a colored icon from a color palette. Each person should thus be allowed to select these controls and see their states without affecting the other person.

One solution provides each person with a different screen area holding their own controls. Yet this is inefficient in terms of space and clutter, especially if there are many controls. Instead, a two-sided relaxed-WYSIWIS display allows an interactive object 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 can show the color selected by the user on one side as ‘blue’, while simultaneously showing the different color selected by the other user as ‘orange’ on the other side. In such cases, these interactive objects can be considered a mini-public territory (as the objects and actions over them can be done by all) and a mini-personal territory (as the selected visible state of the object is personal and specific per side).

4.2.5. Managing attenuation across the medium

Depending on the technology, image clarity can be compromised by the medium. In our own experiences with a commercial transparent LED display (such as the one shown in Fig. 5), image visibility and contrast through the screen was poor. Projection systems are also problematic. For example, Olwal et al. (2006) describe how their projection-based FogScreen™ transparent display diffuses light primarily in the forward-direction, making rear-projected imagery bright and front-projected imagery faint. Their 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 also adapted by Heo et al. (2013) in their TransWall system. Both systems strove to maintain image brightness, where projected images one either side were precisely aligned to generate the illusion of a single common image per side. Another solution layers two transparent displays together, so that each side is seen at its full brightness. The software used to implement transparency (e.g., alpha-blending techniques, color correction) can also affect what can be seen through the user interface (e.g., Harrison et al., 1995; Baudisch and Gutwin, 2004).

While the above solutions work to display the same content, a system that can display different content per side can, as a side-effect, also be able to adjust image brightness and clarity to manage attenuation problems.

4.3. Augmenting human actions to mitigate issues resulting from degrading transparency

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. Depending upon the technology and circumstance, transparency can become degraded. When this happens, it becomes increasingly difficult to see the other person through the screen (including their gestures and actions). Thus workspace awareness can be compromised. Factors that affect transparency include the following, where Fig. 10 selectively illustrates how they are manifested in our own system.

- **Graphics technology.** Different technologies vary greatly in how they draw pixels on a transparent display, e.g., dual-sided projector

systems (Li et al., 2014; Olwal et al., 2008), OLED and LCD screens, and even LEDs moving at high speed (Lee et al., 2014). These interact with the other factors below to affect what people can see through the screen.


- **Screen materials** can afford quite different levels of translucency, where what one sees through the display is attenuated by the material used (e.g., Lee et al., 2014; Li et al., 2014; Olwal et al., 2008). For example, manufactured screens sandwich emissive and conductive layers between glass plates in OLED displays, which affects its transparency. As we will see shortly, our own work uses fabric with large holes in it as the screen material: the trade-off is that larger holes increase transparency, while smaller holes increase the fidelity of the displaying graphics (Fig. 10, with detail shown in Fig. 12).
- **Graphics density.** A screen full of high-density, busy, and highly visible graphics compromises what others can see through those graphics. That is, it is much harder to see through dense cluttered graphics (Fig. 10 right) *vs.* uncluttered graphics (Fig. 10 left)
- **Brightness.** It is harder to see through screens with significant bright and light (*vs.* dark) content, particularly if graphics density is high. Somewhat similarly, if bright projector(s) are used, they can reflect back considerable light, affecting what people see through it (again, compare Fig. 10 right *vs.* left).
- **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 can produce imbalances in what people see. This is akin to a lit room with an exterior window at night time: those outside can see in, while those inside see only their own reflections. For example, the system as shown in Fig. 10 is located in a dark room with blackout curtains to minimize glare and lighting differences.
- **Personal lighting.** If people on the other side of the display are brightly illuminated, they will be much more visible than if they are poorly lit. For example, the configuration in Fig. 10, top includes a light to illuminate the person. That light is off in Fig. 10, bottom.
- **Clothing and skin color** and their reflective properties can affect a person’s visibility through the display. For example, the bare face and hand seen in Fig. 5 top left is reasonably visible. The hand would be far more visible if the person was wearing a white reflective glove, and far less visible if wearing a black glove as in Fig. 13.

Because of these factors, transparency (and thus the visibility of the other person) can alter dramatically throughout a collaborative interactive session. Screen materials and graphics display technology are static factors, but all others are dynamic. Graphics density and brightness of particular display areas can change moment by moment as a function of screen content. Lighting changes as interior lighting is turned on and off, by the exterior light coming into the room (e.g., day *vs.* nighttime lighting), and by shadows. Clothing, of course, will vary by the person.

To mitigate this problem, we suggest augmenting a person’s actions with literal on-screen representations of those actions so they are readily visible by the other person. Examples in our own system (sketched in Fig. 3 and discussed shortly) include highlighting a person’s fingertip with a glow (to accentuate approaching touch selections), and generating graphical traces that outline a finger’s movements (to accentuate simple hand gestures). Yet showing the same visual augmentation on both sides may be less useful, as they may actually interfere with the person performing the action. A two-sided display allows these visual augmentations to be customized not only per action, but also per side. Later sections of this paper will return to this theme, where we will evaluate the effectiveness of particular augmentation schemes when transparency is degraded.



a) sparse graphics, lit person



b) dense graphics, lit person



c) sparse graphics, unlit person



d) dense graphics, unlit person

Fig. 10. The transparency of FACINGBOARD-2 as affected by various graphic density and lighting conditions. The person is located on the other side of the display. **a)** sparse graphics, lit person. **b)** dense graphics, lit person. **c)** sparse graphics, unlit person. **d)** dense graphics, unlit person.





Fig. 11. The FACINGBOARD-2 Setup.



Fig. 12. The FACINGBOARD-2 open-weave projection screen.



Fig. 13. The FACINGBOARD-2 tested application.

5. The FACINGBOARD-2 infrastructure

We implemented our own two-sided collaborative transparent display, which we call 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.³

5.1. Projector and display wall setup

Fig. 11 illustrates our technology setup. We attached fabric (described below) to a 57 cm by 36 cm aluminum frame. Two projectors are mounted back-to-back above the frame along with mirrors. Using two projectors affords a bright image on either side, different graphical projections per side, and 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. Blackout curtains are used, and the

ambient room light is kept somewhat low to minimize glare. However, directional lights (seen in Fig. 11 left at the upper corners of the frame) can illuminate the people on either side.

5.2. 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 the same content, visible from either side. The various glass screens and/or films used in projection systems would not work well for two-sided projection, as those screens or films are designed with the goal of high-clarity bleed-through to their other side to make the projected content visible.

Instead, we explored fabrics 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 each side’s projected contents;
- a person could see through the holes in the open weave to the other side;
- bleed-through 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 (Heo et al., 2013), sound travels easily through openly-woven fabric.

Fig. 12 illustrates how this fabric 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 the facing side of the 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 tinge. 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 inexpensive and easily accessible materials: fabrics for semi-transparent window blinds that are woven out of wide, mostly 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 translucent), 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, compromising image quality and brightness. Low brightness also meant that any bleed-through from the other side would be more visible. Very light materials (e.g., white) reflected the projected content too well, where the overall 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. *Size of holes.* 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 Fig. 12: tobacco thread color, and 10% openness. Openness is a metric used by manufacturers that measure the percentage of light penetration of blinds as determined by its thread width and size of hole.

³ A video illustrating FACINGBOARD-2 is included in Li et al. (2014) and is publicly available at <http://group.lab.cpsc.ualgary.ca/Publications/2014-TransparentDisplay.DIS>

5.3. Input

Raw input is obtained from an off-the-shelf OptiTrack motion capture system. Eight motion capture cameras are positioned around the display (Fig. 11). People 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 could 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 Heo et al., 2013), or vision-based tracking systems (e.g., the Kinect or LeapMotion, as used by Li, 2015), or 6 DOF input devices such as the Polhemus or equivalent (e.g., as used by Olwal et al., 2006). All have their own particular set of advantages and disadvantages (e.g., marker-based or markerless, high or low accuracy, volume of space covered, ability to detect and track in-air gestures in front of but not touching the screen).

5.4. Limitations and practicalities

Our FacingBoard-2 infrastructure works well as a prototyping platform. While it could be the basis for a commercially deployable product, it would be even better if it improved upon several limitations.

First – and common across all transparent displays – the degree of transparency is greatly affected by various factors as already described in Section 4.3. As foreshadowed previously, Fig. 10 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 the figure). The best transparency is in Fig. 10a, where projected graphics are sparse and the person on the other side is well lit. With denser graphics (Fig. 10b) 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 (Fig. 10c), or dense graphics (Fig. 10d).

Second, the fabric used to construct FacingBoard-2 is not ideal. Its threads are not highly reflective, which means that the projected image is not of the brightness and quality one would expect of modern screens. As was seen in Fig. 12, there is also a very small amount of bleed-through of bright image portions to the other side. However, this is barely noticeable if the other side also contains a brightly projected image, and the image resolution is reasonable in spite of the open weave. We believe better fabrics could alleviate these limitations. Display screens (*vs.* projection systems) could also be designed around the same open weave principle. For example, one possibility is to paint a small grid or series of reflective opaque dots onto both sides of an otherwise non-reflective thin transparent surface (or set of sandwiched surfaces).

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. While we minimize occlusion by using downward-angled mirrors (Fig. 11), some occlusion can still happen, for example with taller users over certain screen areas.

6. The FACINGBOARD-2 testbed application

The FACINGBOARD-2 infrastructure is best seen as a medium that allows interaction designers to explore what is possible in a true two-

sided collaborative interactive transparent display. Because our infrastructure offers independent control of both input and output on either side, we could realize various relaxed-WYSIWIS features as motivated by our design rationale in Section 4. To do this, we created a test-bed application: the interactive photo and text label manipulation previously illustrated in Figs. 10–12. Fig. 13 shows a moment in time, illustrating how the system – and the person on the other side – appears to a user on one side. Fig. 14 shows that same moment in time, but this time how it appears to the person on the other side.

6.1. Features

We previously explained how the ability to project different graphics supports relaxed-WYIWIIS, which in turn allows for selective image and text reversal, public to private work territories, semi-personal views of public objects, personal state of controls, different feedback *vs.* feedthrough, and augmenting human actions via visuals. We now illustrate the particular ways FACINGBOARD-2 can be used to achieve these effects. While set within our simple testbed application, we believe these ideas can be generalized to a broad variety of other collaborative transparent display applications.

6.1.1. Public territories

As annotated in Fig. 8, the public territory consumes the majority of the display. Its content is visible to all, and both people can interact with its objects (images and text boxes) simultaneously via direct touch.

6.1.2. Private territories

The system also includes private territories supporting individual storage of photos and text, seen as the white area at the bottom of the display in Fig. 13. Each person's private area is aligned directly atop the other (e.g., compare the location of the private areas between Fig. 13 and Fig. 14). However, its contents are distinct to each viewer, where each person can see and interact with different things. For example, Fig. 13 shows that Person 1 has placed 2 photos in his private area, while Fig. 14 shows how Person 2 has placed a single different photo in his area. Each person can drag objects from the public area to their private area, which causes those objects to disappear from the other person's view. When objects are dragged out of the private area, they reappear in the public area. When a person is manipulating with objects in the private area, the other may see a person's arm movements over that area, but not what is being manipulating. Thus limited workspace awareness is provided (that the person is doing some private work) while still safeguarding privacy (as contents are not visible).

6.1.3. Personal territories showing personal state

The palette of controls, shown on the left side of Fig. 13 and on the right of Fig. 14 are personal territories. Like the private area, the palette is aligned on both sides to appear atop each other. However, like the text and images in a public territory, the actual controls (the buttons) are also aligned on both sides and visible to both people. What makes it a personal territory is that the buttons reflect their state on an individual basis, where selected buttons are shown in white to indicate what that particular person had selected. For example, we see in Fig. 14 that Person 2 has selected the '4px' border thickness and 'Orange' border color, while in Fig. 13 Person 1 has no options selected, as they are in a different drawing mode.

6.1.4. Feedthrough

Within the above personal territories, buttons (all which perform the same function) are aligned. This provides for some workspace awareness. When Person 1 selects a button in their personal palette, Person 2 will see (via transparency) that Person 1 has touched that button. Because this operation can be missed or its details misconstrued, our system adds graphical feedthrough to accentuate a person's




Fig. 14. Image and text reversal in FACINGBOARD-2.

touch action and button selection on the other side. Here, the button as seen on Person 2's side animates for a few seconds (as feedthrough) to reveal Person 1's selection before fading back to its original form. Person 1's feedback differs, where it shows the button briefly highlighted before changing its state. The feedthrough enhances Person 2's awareness of Person 1's actions. Similarly, feedthrough of the other person's interactions with other objects – including those in the public area – can be enhanced in a manner that best reflects the action.

6.1.5. 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, Fig. 13 shows one person's view of the correctly oriented images and text in the public area. However, these images would normally appear mirror-reverse to the person on the other side, as in Fig. 14a. We overcome this problem by selectively flipping images and text in place, as illustrated in Fig. 14b. 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 tool palette and within the private territory is flipped in place to make it readable on either side. While flipped graphics is the system default, users can override this.

6.1.6. 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 (as mentioned above), add a red border to it, change the border thickness, as well as the background color of the text. These changes appear only on that person's side. For example, in Fig. 14a, Person 2 has kept the image and text reversed, as he wishes to point out their fine details. This makes its contents identically aligned to what the other person sees in Fig. 13, where fine-grained gestures will point to the correct internal parts of the object. Later, as seen in Fig. 14b, he has reversed the text and images so they are now correctly oriented for personal viewing. Fig. 14 also shows how Person 2 has added a red border to an image and has colored a text object in orange, which differs from what Person 1 sees in Fig. 13.

6.1.7. Augmenting human actions

As previously described (and elaborated shortly), the transparency and thus 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. In particular, our work considers how mid-air finger touches and movements could be augmented. While just a subset of all actions possible, tracking fingers is important. It supports awareness of another's basic mid-air gestures made over the 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).

Our first solution (Fig. 15), called *augmented touch*, enhances touch actions. We enhance awareness by displaying a small visualization (a modest-sized dot) on the spot where the fingertip orthogonally projects onto the display. The dot only appears on the other side of the display, as it could otherwise mask the person's fine touch selections. For example, in Fig. 13 Person 1 is touching a photo and no dot is visible. However, Person 2's view of the workspace from the other side (Fig. 14a,b) reveals a gold dot marking Person 1's touch. Fig. 15a-c shows how the actual size of the dot varies as a function of the distance between the fingertip and the display. The dot is small when the finger is far from the surface (Fig. 15a), gets increasingly larger as the finger moves towards the surface (Fig. 15b) and is at its largest when touching the surface (Fig. 15c). When a touch occurs, the dot's color also changes.

Our second solution, called *augmented traces*, enhances gestural acts. As seen in Fig. 16, an ephemeral trail follows a person's in-air finger 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. We derived augmented traces from telepointer traces as used in remote groupware (Gutwin and Penner, 2002).

6.2. Testbed experiences: the problem of varying transparency


We created the FACINGBOARD-2 application as a testbed. We did this to experience what collaboration was like through a two-sided transparent display, and whether the particular features above worked to support those collaborations. Our experiences were generally positive, with one major exception. When working with our earliest version, which did not include the touch or trace augmentation, we became increasingly concerned about the changes in transparency that occurred. As already discussed, many factors affect the moment-by-



a) Small dot to reflect distant finger



b) Dot's size increases with approaching finger




c) Dot at full size, color change indicating touch

Fig. 15. Enhancing touch actions. The person is on the other side of the screen.

moment transparency of the display as a whole, as well as the transparency of particular areas of the display (e.g., as affected by graphics density and image brightness). As transparency became increasingly compromised, we found it increasingly effortful to see and track the other's actions through the screen, which led to a perceived loss of workspace awareness. As a consequence, we added the touch and trace techniques mentioned above as part of our iterative development.

Our personal experiences with these augmentation techniques suggest that they do mitigate the transparency issue, at least to some extent. Still, there were several questions that deserved answering at a more precise level, questions that have not been addressed in the workspace awareness literature. First, what is the severity of the problem, i.e., the extent of workspace awareness loss as a function of degraded transparency? Second, what is the efficacy of our touch and

**Fig. 16.** Enhancing gestural events through traces. The person is on the other side of the screen.

trace augmentation methods over different transparency conditions? While we felt they helped in low transparency conditions, we had no clear evidence that this was actually the case. There was also the chance that our visual augmentations could interfere with the viewer's interpretation of the scene when transparency was either uncompromised or somewhat compromised: the viewer would then have to track both the other person as seen through the screen and the augmented visual on the screen, which could increase cognitive load.

Consequently, we investigated the relationship between workspace awareness, degrading transparency, and augmentation methods over a variety of tasks, as discussed next.


7. Study methodology

Our study concerns itself with the interplay between transparency and workspace awareness, and the efficacy of particular augmentation techniques. For terminology convenience, the *viewer* is the person (the participant) who observes the *actions* of the *actor* (the experimenter) on the other side of the display. Our first hypothesis is that viewer's workspace awareness degrades as transparency is compromised. Our second hypothesis is that this degradation can be mitigated by enhancing the actor's actions via touch and trace augmentation methods.


We decided upon a controlled laboratory study designed to probe the relationship between transparency, display density, and trace augmentation across a variety of workspace awareness tasks. Using this methodology, we could control and empirically measure the effects of display transparency and augmentation on workspace awareness, something which could not be easily probed or quantified in the more casual 'real world' study. We could also control for the way people performed tasks, which again would be difficult to do in a real world setting where participants may develop workarounds to overcome workspace awareness deficits (e.g., by relying heavily on speech).

As we will detail below, we used artificial patterns instead of photographs and text (Fig. 17) to control for transparency across the entire screen. These patterns allowed us to examine a range of transparencies, from quite transparent to barely see-through. Our controlled study also relied on three simple experimental tasks, whose interaction mechanics are common to many real world situations (Gutwin and Penner, 2002). Because each task relies on the viewer's ability to maintain workspace awareness, the viewer's accuracy and success rate at correctly completing a task provides a measure of workspace awareness.⁴


⁴ A video illustrating the study and its conditions is viewable at <http://grouplab.cpsc.ualgary.ca/grouplab/uploads/Publications/Publications/2014-TransparentStudy.Report2014-1065-16.mp4>.




Level 1 transparency / front lit actor:
actor clearly visible



Level 2 transparency / front lit actor:
body somewhat visible, hand visible



Level 3 transparency / front lit actor:
body barely visible, hand somewhat visible



Level 4 transparency / no front lighting
body / hand barely visible

Fig. 17. The 4 transparency conditions with trace augmentation (on blue trail). All show the actor as seen through the screen. The actor is tracing a route within the route task. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

7.1. Independent variables

7.1.1. Transparency

We vary transparency as an independent variable. We use four transparency levels. Each comprises a particular mix of graphical density patterns projected onto the viewer's side of the display, and lighting on the actor. To explain, Fig. 17 illustrates the 4 transparency conditions.⁵ All sub-figures show the actor in the same pose indicating a route through several circles, with trace enhancement turned on (the route task will be described shortly). The actor in all but the bottom right is front-lit. At the top left of Fig. 17 is level 1, the most transparent condition, where the actor's hand, arm, body and eye gaze are clearly visible through the display. The top right is level 2, where we increase the graphical density by projecting a pseudo-random pattern

comprising a ratio of 25% white to black pixels. The actor's arm and hand are still clearly visible, but details of his body and eye gaze are harder to make out. The bottom left is level 3: the ratio is 67% and the actor's details become even more difficult to see (although the hand remains reasonably visible). The bottom right is level 4: the ratio remains at 67% but the actor is no longer front-lit. Here, the actor – while still discernable – is barely visible.

7.1.2. Augmentation: enhancing touch and gestures

As previously explained, we developed two feedthrough augmentation techniques that try to enhance the viewer's visibility of the actor's touch and gestural actions. The *augmented touch* technique draws a circular glow on the screen location corresponding to the actor's finger. The glow becomes larger and visually more intense as the actor's finger approaches the display, where the glow changes color when the display is actually touched (Fig. 15). The *augmented trace* technique draws a fading line on the display, where the line follows the path of the actor's finger (Fig. 16). We treat augmentation as an independent variable,

⁵ To make the figure photographs legible in this manuscript, we altered the lighting somewhat from the actual experimental conditions, and portray the actor without gloves. However, the images are reasonable approximations of what study participants saw.

where it is either present or absent. The particular augmentation technique used (touch *vs.* trace) depends upon the particular task associated with each study.

7.2. Tasks and dependent variables

We developed three tasks that exemplify common real-world activities that people would be expected to perform on a two-sided display, where our tasks are variations of those described and developed in Gutwin and Penner (2002). The experimenter is the actor, while the participant is the viewer. The viewer's performance over these tasks in our 8 conditions are our dependent variables, where they serve as a measure of their ability to maintain workspace awareness.

7.2.1. The shape task/error rate

Shape gestures refer to finger movements that trace geometric shapes that convey symbolic meanings, e.g., a character, a rightwards gesture indicating direction. Shape gestures can appear anywhere, and are not necessarily associated with workspace artifacts.

The *shape task* involves shape gesture actions. The actor uses his finger to 'write', as a shape gesture, a horizontally-reversed English letter over a randomly selected quadrant just above the display surface (reversal correctly orients the letter to the viewer). The viewer's task is to say out loud the letter s/he saw. We note that this task also requires the viewer to disambiguate those parts of the gesture that are not part of the letter (e.g., when the person's finger approaches and leaves the display surface). For augmentation conditions, we use the trace augmentation technique, with 8 trials per block.

Error rate is the dependent variable: the number of incorrectly recognized or missed shapes over the total number of shapes presented per condition.

7.2.2. Route task/accuracy rate

Route gestures are paths going through some objects in the workspace. Routes can suggest actual paths in the space, transitions between object states, or groupings of objects. Unlike shape gestures, they are made relative to the workspace and its artifacts.

The *route task* involves route gesture actions. A 16×10 grid of circles are aligned to appear on the same locations on both the actor's and viewer's sides of the screen. The actor then gestures a path through a particular sequence of circles (illustrated in Fig. 17). While routes differ between trials, all paths go through five circles with one turn in the middle. The viewer's task is to reproduce that path by touching the circles the path went through. We use the trace augmentation for the augmentation conditions, with 8 trials per block.

Accuracy rate is the dependent variable: the number of correct responses over the total number of responses per conditions. Correct responses are those where the viewer has correctly indicated the circles the route gesture went through.

7.2.3. The point task/response time, response error, miss rate

The previous tasks are examples of tightly-coupled collaboration: both actor and viewer focus their attention on the gesture as it is being performed. We wanted to see what would happen in *mixed-focus collaboration*, where participants pursue individual work while still monitoring group activities (Gutwin and Greenberg, 1998, 2002). As previously mentioned, workspace awareness is particularly important for mediating the shift from loosely to tightly coupled group work, for it helps create opportunities to coordinate mutual actions.

The *point task* measures, in part, a viewer's ability to stay aware of the actor's touch action during mixed-focus collaboration. We describe it from the actor's and then the viewer's perspective.

– *Actor's viewpoint and task.* A randomly-positioned circle appears only on the actor's side of the display. The actor taps that circle, which then disappears. After a pseudo-random time interval, a new

circle appears elsewhere on the display, the process repeats.

– *Viewer's viewpoint and task.* To emulate mixed focus collaboration, the viewer, while performing individual work, has to simultaneously follow the actor by monitoring and repeating the actor's touch actions. Thus the viewer has to concurrently perform an individual task and a group task (called the follower task).

- a) *Task 1: Individual task.* Solid squares pseudo-randomly appear on only the viewer's side of the display. The viewer was asked to tap those squares as they appear.
- b) *Task 2: Follower task.* The viewer was asked follow the actor's actions. To do so, the viewer has to monitor the touch actions made by the actor, and then tap the spots that the actor touched. The viewer was told that the follower task took precedence over the individual task. That is, as the viewer perform task 1, they have to simultaneously monitor the workspace for the actor's touch actions, where the viewer has to react as quickly and as accurately as possible to indicate where the actor had touched.

On average the ratio of individual to follower task episodes were ~3:1, but were interleaved irregularly to make their timing unpredictable to the viewer. We use the touch augmentation for the augmentation conditions, with 80 trials (60 individual task; 20 follower task) per bloc.

Three metrics measured awareness as a dependent variable. *Response time* is the elapsed time between the touch from the actor and the following responding touch from the viewer. *Response error* is the distance between the location touched by the actor and the location touched by the viewer. *Miss rate* is the rate where participants failed to react to a touch by the actor, e.g., because the viewer didn't notice the touch or failed to see where the touch occurred.

7.3. Study design

We ran three studies. Each study is similar in form, except that participants performed a different task (shape, route and point), each with their own dependent variables. All are based upon a within-subject (repeated measures) ANOVA factorial design: *transparency* (4 levels)×*augmentation* (2 levels), or 8 different conditions per task. All used the same participants as viewers, where each participant did all three tasks over all 8 conditions in a single 90 min session. For each condition, subjects underwent many repeated trials. Transparency levels are as described above. Augmentation type varies per task, and is either present (augmentation on) or absent (augmentation off).

7.4. Hypotheses

Our null hypothesis is suggested by our study design.

Across the four transparency levels and the presence or absence of augmentation, there is no difference in a participant's ability to

- a) recognize the shape as measured by the error rate,
- b) trace a route as measured by the accuracy, and
- c) observe touches as measured by the response time, the response error, and the miss rate.

Before running the study, we made several prediction.

1. Participants' performance with no augmentation would generally deteriorate as transparency was compromised, although we could not predict the degree of deterioration.
2. In the level 4 low transparency condition, augmentation would improve performance when compared to no augmentation, as it would supply otherwise hard-to-see workspace awareness information. However, the performance would be less than in the level 1 high transparency condition that offers richer workspace awareness information.

3. In the level 1 high transparency condition, augmentation would decrease performance when compared to no augmentation, as the visualization would compete with people's perceptions of the other person through the screen.

In all cases, we could not predict the actual amount of performance differences. We were also uncertain about the performance outcomes in the level 2 and level 3 transparency conditions.

7.5. Materials

The study was conducted on our two-sided transparent display prototype. As detailed in Section 5, it is a 57×36 cm two-sided transparent display, where projectors on each side can project different visuals without significant bleed-through. An OptiTrack Flex 13 motion capture system tracked a marker placed on the index finger of gloves worn by participants. We developed dedicated software modules to display screen contents for each task, and collect data about user actions.

7.6. Participants

Twenty-four participants (10 female and 14 male) between the ages of 19 and 41 were recruited from a local university for this study. All were experienced in some form of touch screen interactions (e.g., phones, surfaces). All were right-handed. Each participant received a \$15 payment.

7.7. Procedure

After being briefed about the study purpose, the participant completed a demographics questionnaire. Participant then performed the shape, route and point task in that order. For each task, the participants were instructed on what they had to do, and then did 9 blocks: a practice block and then eight counter-balanced blocks corresponding to the eight previously described conditions. After completing each task, the experimenter led the participant through a semi-structured interview, where the participant was asked to comment about his or her experiences with the various conditions, as well as the strategies used to perform tasks.

8. Results

8.1. Statistical analysis method

We ran a two-way repeated measures ANOVA for each of the measures obtained from the three tasks, with sphericity assumed. For sphericity-violated cases, we used Greenhouse-Geisser corrections. For post-hoc tests, we used the test of simple main effects with Bonferroni corrections. The level of significance was set at $p < 0.05$.

8.2. The shape task

In the shape task, the actor wrote, as a gesture, a horizontally reversed capital letter; the viewer's task was to say what letter he or she saw. The error rate of the shape task was then calculated as the ratio of misrecognized letters in each condition for each participant.

8.2.1. Results

Our analysis reported a significant main effect for transparency ($F_3, 69=12.458, p < 0.05$), augmentation ($F_1, 23=42.037, p < 0.05$), and the interaction between them ($F_3, 69=14.73, p < 0.05$).

Fig. 18 graphically illustrates the means of the error rate and our post-hoc test results. The green and blue lines represent the augmentation on vs. off conditions respectively, while the four points on those lines are the values measured at each of the four transparency levels,




Fig. 18. Shape task results. Error rate plotted by condition. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

with level 1 on the left and level 4 on the right. The vertical red lines indicate where the post-hoc test reported a significant difference between the augmentation off vs. on values at a particular transparency level. For example, we see that the red lines indicate a significant difference in the error rate between the augmentation on/off conditions at levels 2, 3 and 4. The numbers in the colored box next to particular points indicate which transparency levels differed significantly on a given augmentation condition. For example, with augmentation off, we see from the numbers in the blue box that: level 1 differs significantly from levels 3 and 4; and levels 2 and 3 differ from level 4. However, with augmentation on, there are no significant differences in the error rate at any transparency level.

8.2.2. Discussion

The null hypothesis for the shape task is rejected. First, without augmentation, there is a notable increase in the error rate as display transparency decreases, where most pairwise differences between these means are statistically significant (Fig. 18, blue line). Differences are practically significant as well, where the error rate of ~10% in the most transparent condition increases to ~44% in the least transparent condition (see the blue line data points in Fig. 18).

Second, with augmentation, the error rate is constant regardless of the transparency level, with no significant difference seen across any of the transparency levels when augmentation is used (Fig. 18, green line). Notably, the error rate is low at ~6%. This sharply contrasts with augmentation off conditions, where the error rate increases as transparency decreases.

Third, the presence or absence of augmentation does not affect error rate in highly transparent conditions, i.e., using augmentation when it is not needed does not incur a negative effect (compare the first points in Fig. 18's green vs. blue lines, where differences are not significant).

In summary, the results indicated that people have much more difficulty correctly recognizing shape gestures as transparency is compromised (without augmentation). They also indicate that the trace augmentation mitigates this problem, where people are able to maintain a largely stable and fairly low error rate ($M=6.0\%$, $SD=0.013$) that is equivalent to highly transparent conditions. That is, the trace augmentation supports people's ability to perceive the other's gestural shapes as transparency deteriorates.




Fig. 19. Route task results. Accuracy rate plotted by condition. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

8.3. The route task

In the route task, the actor gestured a path through a particular sequence of circles shown on the display. The viewer's task was to reproduce the path by touching particular circles that the path went through. The accuracy of the route task was then calculated as the ratio of correctly reproduced paths to the total paths in each condition.

8.3.1. Results

Our analysis discovered a significant main effect for transparency ($F(3, 69)=7.240$, $p < 0.05$), augmentation ($F(1, 23)=42.037$, $p < 0.05$), and the interaction between them ($F(3, 69)=4.515$, $p < 0.05$). Fig. 19 graphically illustrates the means of the accuracy rate and our post-hoc test results, where their portrayal is similar to Fig. 18.

8.3.2. Discussion

The null hypothesis for the route task is rejected. First, without augmentation the accuracy decreases noticeably as display transparency deteriorates (Fig. 19, blue line), where we see statistically significant differences between the accuracy at transparency level 1 and all other levels. The differences are also practically significant: the ~91% accuracy in the most transparent condition degrades to ~62% in the least transparent condition.

Second, accuracy across transparency levels in augmentation-on conditions is constant at a high level (~85–90%): the slight downward trend is not significant (Fig. 19, green line). For transparency level 4, accuracy is significantly higher with augmentation than without.

Third, the presence or absence of augmentation does not affect accuracy in highly transparent conditions, i.e., it does not incur a negative effect (compare 1st points in Fig. 19's green vs. blue lines, where differences are not significant).

To sum up, the results indicate that people have much more difficulty accurately perceiving the route gesture when display transparency is compromised (without augmentation). The results also indicate that trace augmentation alleviates these difficulties at low levels of transparency. That is, the trace augmentation supports people's ability to perceive the other's path drawing gestures relative to objects as transparency deteriorates.

8.4. The point task

In the point task, the viewer was asked to: (a) carry out a separate




Fig. 20. Point task results: response time by condition. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

independent task, and (b) simultaneously monitor and respond to the actors' touch actions on the display by touching the location where the actor had just touched. There were three dependent variables. Response time is the average elapsed time between the actor's touch and the responding viewer's touch. Response error is the distance between the location touched by the actor and the corresponding location touched by the viewer. Miss rate is the rate where viewers failed to react to the actor's touch. Each is discussed in turn.

8.4.1. Results: response time

Our analysis revealed a significant main effect for response time for transparency ($F(3, 69)=20.731$, $p < 0.05$), augmentation ($F(1, 23)=4.517$, $p < 0.05$), and the interaction between them ($F(3, 69)=4.620$, $p < 0.05$). Fig. 20 graphically illustrates the means of the response time and our post-hoc test results.

8.4.2. Discussion: response time

The null hypothesis is rejected. First, without augmentation, response time tends to increase as display transparency decreases (significant differences are visible between these means in Fig. 20, blue line). The differences are also practically significant, with response times of ~700 ms increasing to ~1000 ms between the most to least transparent conditions.

Second, with augmentation the response time exhibits a statistically significant but somewhat modest increase from transparency level 1 (~700 ms) to level 2 (~800 ms), with no further increase afterwards (Fig. 20, green line).

Third, for levels 1 and 2 transparency, adding augmentation neither increases nor reduces the response time with respect to similar conditions without augmentation i.e., it does not incur a negative effect. Yet augmentation is beneficial in low transparency conditions (compare Fig. 20 data points between the green and blue lines).

In summary, the results indicate that people pursuing their own individual tasks while simultaneously monitoring another person's touches are somewhat slower to respond when transparency is compromised (without augmentation). The results also indicate that the touch augmentation method mitigates this somewhat: their response time increases only slightly in low transparency conditions.




Fig. 21. Point task results: *response error* by condition. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

8.4.3. Results: response error

Our analysis revealed a significant main effect on response error for transparency ($F(3, 69)=11.676$, $p < 0.05$), augmentation ($F(1, 23)=48.508$, $p < 0.05$), and the interaction between them ($F(3, 69)=13.270$, $p < 0.05$). Fig. 21 graphically illustrates the means of the response error and our post-hoc test results.

8.4.4. Discussion: response error

The null hypothesis is rejected. First, without augmentation the response error increases as display transparency deteriorates (significant differences are visible between these means in Fig. 21, blue line). The differences are also practically significant, where the response error of ~28 mm in the most transparent condition increases threefold to ~99 mm in the least transparent condition.

Second, with augmentation the response error is constant regardless of the transparency levels, with no significant differences between them (Fig. 21, green line). Furthermore, the response error stays low (at ~33 mm) when augmentation is present; this contrasts dramatically to the statistically significant increase in response error without augmentation when display transparency is compromised (compare green and blue lines in Fig. 21).

Third, the presence or absence of augmentation does not affect error rate in highly transparent conditions, i.e., it does not incur a negative effect. Yet it is beneficial in all other conditions when transparency is compromised (compare Fig. 21 data points between the green and blue lines).

In summary, the results indicate that people are less precise when display transparency is compromised (without augmentation). The results also indicate that the touch augmentation method mitigates this considerably.

8.4.5. Results: miss rate

Our analysis found a significant main effect on the miss rate for transparency ($F(3, 69)=23.249$, $p < 0.05$), augmentation ($F(1, 23)=21.300$, $p < 0.05$), and the interaction between them ($F(3, 69)=15.434$, $p < 0.05$). Fig. 22 graphically illustrates the means of the response time and our post-hoc test results.




Fig. 22. Point task results: *miss rate* by condition. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

8.4.6. Discussion: miss rate

The null hypothesis is rejected. First, without augmentation the miss rate increases sharply as transparency is reduced where a significant difference is seen between the first 3 levels *vs.* the 4th level (Fig. 22, blue line). This difference is practically significant, where the miss rate jumps from ~6% in the most transparent condition to ~43% in the least transparent condition.

Second, with augmentation the miss rate remained invariably low at ~8% (Fig. 22, green line).

Third, the presence or absence of augmentation does not affect error rate in highly transparent conditions, i.e., it does not incur a negative effect. Yet it is beneficial in all other conditions when transparency is compromised (compare Fig. 22 data points between the green and blue lines).

In summary, the results indicate that people, when pursuing their own individual tasks while simultaneously monitoring another person's touches, are much more likely to miss the other person's touch actions when transparency is compromised (without augmentation). The results also indicate that the touch augmentation method mitigates this: the miss rate remains low under all transparency conditions.

8.5. Overall discussion of results

The above results, when considered collectively, consistently show that decreasing display transparency reduces a viewer's awareness of the actor's actions on the other side of a transparent display. This is as predicted. Across all three tasks and as reflected by all five measures, participants' performance with no augmentation generally deteriorated as transparency was compromised. Differences were both statistically and practically significant. While we could not predict the actual amount of performance differences ahead of time, we now see that the degradation of transparency imposes a severe performance penalty in all measures.

The same results also show that augmentation techniques mitigate awareness loss when display transparency is compromised. Again, this was true across all tasks and all measures, where differences were both statistically and practically significant. We had predicted improvement at only very low transparency, and were thus pleased to see it work at moderate levels of transparency as well.

We also saw – to our surprise – that the augmentation techniques did not have a negative effect in situations where they were not strictly necessary, i.e., high transparency conditions when the actor's actions are clearly visible. This was contrary to our prediction. Across all tasks and for 4 of the 5 measures, the presence or absence of augmentation had little effect on participants' performance at the highly transparent level. On the other hand, we also saw that augmentation almost always had a beneficial effect when transparency was degraded when compared to the no-augmentation condition.

However, the results also reveal subtleties. While all measures in all tasks show that augmentation helps overcome the degradation in people's performance as transparency declines, it is not always continuous. For example, consider the response time measure in the point task, as illustrated in Fig. 20, where there is a difference between the response time in the augmentation-on condition between levels 1 and 2. Thus we see an (isolated) case where workspace awareness has degraded, but augmentation does not appear to help. Our post-study interviews of participants suggest why this is so. Most reported that their strategy was to watch for movements of other body parts of the actor before the finger was close to the screen (e.g., raising the arm and moving the hand towards the screen). This consequential communication signaled that a touch was soon to occur. Participants said they found it increasingly difficult to see those body movements as transparency decreased, and consequently they reacted more slowly. For example, at transparency level 2 (Fig. 17, upper right), people found it more difficult to see initial arm movements, but they could still see the hand as it approached the display. While touch augmentation provided information about where the fingertip was and its distance to the screen, it did not signal the earlier actions of other body parts and thus had no net benefit. When transparency was compromised even further at levels 3 and 4, participants had more difficulty seeing the un-augmented approaching finger (Fig. 20, blue line). In those cases, augmentation helped signal the approach at closer ranges, thus enabling people to react faster as compared to no augmentation (Fig. 20, green line).

Overall, we conclude that augmentation can supply the information necessary for people to maintain workspace awareness as transparency degrades. In those cases where augmentation may not provide any benefit (such as highly transparent situations where the actor is clearly visible), augmentation can still stay on as it has no negative effects. Keeping augmentation on at all times is useful, as our results also show that the degradation of workspace awareness varies (more or less) as a function of transparency degradation: there is no clear threshold that defines when augmentation should be turned on.

9. Implications on the design of two-sided transparent displays

Providing necessary workspace awareness is crucial for the utility and usability of collaborative transparent displays. Therefore, their hardware and software interface design should guarantee reasonable support for the cues that comprise workspace awareness. We offer two implications for addressing this awareness requirement.

9.1. Implication 1: controlling transparency

Transparent displays are often portrayed as fully transparent in commercial advertisements, many research figures, and even futuristic visions of technology. We suspect that their graphics density and lighting are tuned to show such displays at their best. Yet transparent displays are not invariantly transparent. The consequence (as our results clearly show) is that degrading transparency can greatly affect how collaborators maintain mutual awareness.

One partial solution is to control display transparency as much as possible. Our experimental setup and study confirmed that high graphics density and dim lighting on the actor can reduce what one

can see through the display. This can be partially remedied by design. For lighting, the system could incorporate illumination sources (perhaps integrated into the display frame) that brightly illuminates the collaborators. For graphics density, applications for transparent displays should distribute graphics sparsely on the screen, with enough clear space between its elements to permit one to see through those spaces. Colors, brightness and textures can be chosen to find a balance between seeing the displayed graphics and seeing through those graphics.

Another partial solution controls for external factors. This includes the ambient light that may reflect off the display, and even the color of surrounding walls and furniture. For example, we surrounded our own display with blackout curtains both to block out light and to provide a dark background. Another controllable factor is the color of the collaborators' clothes (bright colors are more reflective than dark colors) and how that color contrasts with the surrounding background. For example, participants can wear white reflective gloves to better illuminate their hand movements to others.

Another partial solution relies on the display technology itself. For example, our display is based on a mesh fabric that only allows a certain amount of light to pass through it. Other technologies, such as JANUS (Lee et al., 2014), may afford more light transmission. However, we should not expect technical miracles, as we believe that all technologies will be affected to some extent by the factors mentioned in Section 4.3.

Another issue may arise, where the degree of transparency required for the moment may be context-dependent. That is, while high transparency may be desired during tightly coupled interactions, the fine-grain fidelity that results may prove distracting in either loosely-coupled interaction, or where there is little need to know what is occurring in the surrounding environment (Lindlbauer et al., 2016). This suggests a form of dynamic transparency that adjusts itself to the degree of awareness desired.

In practice, we expect that the ability to control for the above factors is highly dependent on context. Designers may be able to devise (or recommend) specific transparency modulation mechanisms if they know where the display is used what tasks people are carrying out on it, and the degree of collaboration desired. However, we expect most installations will limit what designers can control. Nonetheless, we can still enhance workspace awareness by augmenting user actions, as discussed next.

9.2. Implication 2: augmenting user actions

We argued previously that display transparency may be compromised, thus limiting the fidelity of what people see through the display. One solution would be to mitigate the various root causes behind transparency degradation, such as to improve the underlying technology to afford better transparency, presenting only sparse screen contents, and controlling lighting and shading. Unfortunately, these approaches may be unfeasible (e.g., better technology), or unavoidable (the user needs dense graphics, the environment constrains how much lighting can be adjusted). This is why we suggested augmentation techniques as another solution.

Our study revealed that augmentation techniques can mitigate awareness loss when display transparency is compromised. In spite of the simplicity of our techniques (revealing the motion of a single finger), they proved effective. This clearly suggests that – at the very least – designers should visually augment a person's dominant finger movements. This is somewhat generalizable, as that finger often signals pointing gestures, is the focal point of input interaction for touch-based displays, and hints at where the actor is directing one's gaze.

However, we can do even better. While seeing finger movement is helpful, body language is far richer. In daily face-to-face activities, we maintain workspace awareness by observing movements of multiple body parts (including gaze awareness) and interpret those sequences in

relation to the workspace. We need to develop augmentation techniques that capture that richness, where we expect it will be helpful across a broader variety of tasks and situations. Examples include systems that: represent the entire hand, that change the representation as a function of distance; that show where a person is looking; that show the entire arm (Tang et al., 2004), or that even show the entire body (Tang and Minneman, 1991).

Of course, there are challenges to this. Technical challenges include tracking. Graphical challenges include designing an easily understood representation that does not occlude, distract, or otherwise interfere with a person's view of the workspace: recall that workspace awareness involves a view of the participant, the workspace artifacts, and the participant's actions relative to those artifacts.

In summary, simple augmentation techniques will likely work well for mitigating awareness loss in many scenarios. However, new techniques and representations should be developed to better match the situation, display and task.

9.3. Discussion and limitations

Our controlled study was, to our knowledge, the first of its kind and, as typical with such studies, has limitations.

First, we used only four transparency levels. While these were chosen to capture a range from highly to barely transparent, it does not cover the full transparency spectrum nor expose other factors that could affect transparency.

Second, our manipulation of graphical density was artificial, where we used a random pixel pattern containing a well-defined ratio of bright vs. dark pixels as a wash. While very useful for understanding transparency effects in particular conditions, real world graphical displays have other characteristics that could prove important. Future work could test how people maintain awareness through (say) a document editor, a photo-viewing application, and/or a running video, each which may change transparent levels across sub-areas of the screen on a moment by moment basis.

Third, the three study tasks were artificial. We do consider these tasks reasonable representatives of what people do during collaboration, as they include typical tracing gestures and touch actions that people perform during cooperative work (Gutwin and Penner, 2002). However, they are not inclusive of all gestures, nor would they cover all interaction nuances. Related to this is that our augmentation methods only matched what we thought would be critical actions within these tasks, i.e., finger touch and movement gestures. We did not attempt, for example, to augment gaze awareness. Future work should, of course, test people doing real tasks, where people may exhibit more complex interaction and gestural patterns of behaviours.

Even so, our own everyday qualitative use of our display running the testbed application illustrated in Fig. 13 aligns with the quantitative results produced by our study. That is, in spite of the artificiality of our study tasks, we do not have any reason to believe that they would not apply to real world tasks.

Our study (along with our design rationale) has laid a strong foundation for understanding the strengths and limits of two-sided collaborative transparent displays. It exposed how compromised transparency can severely affect workspace awareness and thus performance of even simple tasks. It also revealed how this performance loss can be largely overcome by simple augmentation methods.

10. Conclusion

In this paper, we provided reasons behind why we should add two-sided interactive transparent displays to our repertoire of interactive surfaces. We first laid a theoretical foundation, where we summarized the relevant workspace awareness and territoriality theory essential to collaborative surfaces. We then applied these theories to create a design rationale for see-through two-sided interactive displays, where

we argued for two-sided interactive input, different content on both sides, and the ability to augment human actions to overcome display technology limitations.

We then described the design of our FACINGBOARD-2 system, whose characteristics emerged from our design rationale. 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. We showed how the FACINGBOARD-2 infrastructure has the ability to project different graphics without significant bleed-through a mesh-like fabric, which in turn supports relaxed-WYTWIS and transparency. This in turn allows for many software effects supporting collaboration: selective image and text reversal; various territories including public, personal and private areas; semi-personal views of public objects, personal state of controls, different feedback vs. feedthrough, workspace awareness in general, and several ways to augment human actions via visuals. We also highlighted some of the design trade-offs 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.

Yet we also unearthed a significant problem in two-sided transparent displays: they are not always transparent. This is the reason why we created several augmentation techniques that visualize people's actions as touch dots and traces on the screen. To investigate this problem and possible solution in more detail, we performed a controlled study that examined the effect of display transparency on people's awareness of others' actions, and the effectiveness of augmentation techniques that visually enhance those actions. Our analysis confirms that people's awareness is severely reduced when display transparency is compromised, and that augmentation techniques can mitigate this awareness loss. Based on our findings, we suggested a few implications for collaborative transparent display designers.

Our design iterations and study of two-sided collaborative displays have unearthed exciting possibilities. Yet we recognize that the present work is just the beginning of an exploration of what is possible on this new medium.

Acknowledgments

Joaquim Jorge, of the Universidade de Lisboa, contributed to the early stages of this work, where he was a co-author in (Li et al., 2014). We greatly appreciate his input. Special thanks to Sutapa Dey who helped in our pilot studies. This work is supported by the NSERC-AITF-SMART Industrial Chair in Interactive Technologies, and NSERC's Discovery Grant and Surfnetwork.

References

- Azad, A., Ruiz, J., Vogel, D., Hancock, M., Lank, E., 2012. Territoriality and behaviour on and around large vertical publicly-shared displays. In: Proceedings ACM Conference on Designing Interactive Systems (DIS '12). ACM, New York, NY, USA, pp. 468–477.
- Baudisch, P., Gutwin, C., 2004. Multiblending: displaying overlapping windows simultaneously without the drawbacks of alpha blending. In: Proceedings ACM Conference on Human Factors in Computing Systems (CHI '04). ACM, New York, NY, USA, pp. 367–374.
- Clark, H., 1996. *Using Language*. Cambridge University Press, United Kingdom.
- Corning, Incorporated, 2012. A Day Made of Glass 2. (www.youtube.com/watch?v=jZkHpNnXLB0) (Retrieved February 10, 2015).
- Danninger, M., Vertegaal, R., Siewiorek, D.P., Mamuji, A., 2005. Using social geometry to manage interruptions and co-worker attention in office environments. In: Proceedings Graphics Interface (GI '05). pp. 211–218.
- Dewan, P., Choudhary, R., 1991. Flexible user interface coupling in a collaborative system. In: Proceedings ACM Conference on Human Factors in Computing Systems (CHI '91). ACM, New York, NY, USA, pp. 41–48.
- Dourish, P., Bellotti, V., 1992. Awareness and coordination in shared workspaces. In: Proceedings ACM Conference on Computer Supported Cooperative Work (CSCW'92). ACM, New York, NY, USA, pp. 107–114.
- Downing, E., Hesselink, L., Ralston, J., Macfarlane, R., 1996. A three-color, solid-state, three-dimensional display. *Science* 273, 1185–1189.
- Endsley, M., 1995. Toward a theory of situation awareness in dynamic systems. *Hum. Factors* 37 (1).

- Genest, A., Gutwin, C., Tang, A., Kalyn, M., Ivkovic, Z., 2013. KinectArms: a toolkit for capturing and displaying arm embodiments in distributed tabletop groupware. In: *Proceedings ACM Conference on Computer Supported Cooperative Work (CSCW'13)*. ACM, New York, NY, USA, pp. 157–166.
- Gutwin, C., 2002. *Traces: visualizing the immediate past to improve group interaction*. *Proc. Graph. Interface (GI'02)*, 43–50.
- Gutwin, C., Greenberg, S., 1998. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In: *Proceedings ACM Conference on Computer Supported Cooperative Work (CSCW'98)*. ACM, New York, NY, USA, pp. 207–216.
- Gutwin, C., Penner, R., 2002. Improving interpretation of remote gestures with telepointer traces. In: *Proceedings ACM Conference on Computer Supported Cooperative Work (CSCW'02)*. ACM, New York, NY, USA, pp. 49–57.
- Gutwin, C., Greenberg, S., 2002. A descriptive framework of workspace awareness for real-time groupware. *J. CSCW* 11 (3–4), 411–446.
- Gutwin, C., Greenberg, S., Roseman, M., 1996. *Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation*. *Proceedings of HCI on People and Computers XI (HCI '96)*. Springer-Verlag, London, UK, 281–298.
- Hall, E.T., 1966. *The Hidden Dimension*. Doubleday, New York, USA.
- Harrison, B., Ishii, H., Vicente, K., Buxton, W., 1995. Transparent layered user interfaces: an evaluation of a display design to enhance focused and divided attention. In: *Proceedings ACM Conference on Human Factors in Computing Systems (CHI '95)*. ACM, New York, NY, USA, pp. 317–324.
- Heo, H., Kim, S., Park, H., Chung, J., Lee, G., Lee, W., 2013. TransWall. *ACM SIGGRAPH '13 Emerging Technologies*. Article No. 14.
- Hollan, J., Hutchins, E., Kirsh, D., 2000. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM TOCHI* 7, 2.
- Hsu, C.W., Zhen, B., Qiu, W., Shapira, O., DeLacy, B.G., Joannopoulos, J.D., Soljacic, M., 2014. Transparent displays enabled by resonant nanoparticle scattering. *Nat. Commun.* 5, 3152.
- Ishii, H., Kobayashi, M., 1992. ClearBoard: a seamless medium for shared drawing and conversation with eye contact. In: *Proceedings ACM Conference on Human Factors in Computing Systems (CHI '92)*. ACM, New York, NY, USA, p. 525.
- Joung, H.Y., Kim, S.Y., Im, S.H., Huh, B.K., Kim, H., Kwon, G.H. and Park, J., 2016. Optimal user interface parameters for dual-sided transparent screens in layered window conditions. In: *Human-Computer Interaction: Interaction Platforms and Techniques*, Proc. 18th International Conference on Human Computer Interaction (HCI International 2016), LNCS 9732, Springer, pp. 162–169.
- Kent Optronics, 2014. Transparent Display. Retrieved August 5, 2014 from (<http://www.kentoptronics.com/transparent.html>).
- Kuo, H., Hubby, L.M., Naberhuis, S., Birecki, H., 2013. See-through Display. US Patent 8,462,081 B2. Filed Mar 9. 2011. Issued Jun 11. 2013.
- Lee, H., Cho, S., Hong, J., Lee, G., Lee, W., 2014. JANUS. In: *Proceedings of ACM SIGGRAPH 2014 Emerging Technologies*, Article No. 15.
- Lee, J., Olwal, A., Ishii, H., Boulanger, C., 2013. SpaceTop: Integrating 2D and spatial 3D interactions in a see-through desktop. In: *Proceedings ACM Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, pp. 189–192.
- Li, J., 2015. *Two-Sided Transparent Display as a Collaborative Medium* (Master's thesis). Department of Computer Science, University of Calgary, Calgary, Alberta, Canada.
- Li, J., Greenberg, S., Sharlin E., Jorge, J., 2014. Interactive two-sided transparent displays: designing for collaboration. In: *Proceedings ACM Conference on Designing Interactive Systems (DIS '14)*. ACM, New York, NY, USA, pp. 395–404.
- Li, J., Sharlin, E., Greenberg, S., Rounding, M., 2013. Designing the car iWindow: exploring interaction through vehicle side windows. In: *Proc. ACM CHI Ext. Abstracts*, pp. 1665–1670.
- Lindlbauer, D., Liliya, K., Walter, R., Müller, J., 2016. Influence of display transparency on background awareness and task performance. In: *Proceedings ACM Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, pp. 1705–1716.
- Lindlbauer, D., Aoki, T., Walter, R., Uema, Y., Höchtl, A., Haller, M., Inami, M., Müller, J., 2014a. Tracs: transparency-control for see-through displays. In: *Proceedings ACM Symposium on User Interface Software and Technology (UIST'14)*. ACM, New York, NY, USA, pp. 657–661.
- Lindlbauer, D., Aoki, Uema, Y., Höchtl, A., Haller, M., Inami, M. and Müller, J., 2014b. A collaborative see-through display supporting on-demand privacy. *SIGGRAPH Emerging Technology '14*, Vancouver, Canada.
- Lumineq, 2014. Transparent electroluminescent displays, TASEL. Retrieved August 5, 2014 from (<http://lumineq.com/en/products/tasel>).
- Olwal, A., DiVerdi S., Rakkolainen, I., Höllerer, T., 2008. Consigalo: multi-user face-to-face interaction on immaterial displays. In: *Proceedings ACM Conference on Intelligent Technologies for Interactive Entertainment (INTETAIN '08)*. Article No. 8.
- Olwal, A., Lindfors, C., Gustafsson, Kjellberg, T., Mattsson, L., 2005. ASTOR: an autostereoscopic optical see-through augmented reality system. In: *Proceedings IEEE Mixed and Augmented Reality*, pp. 24–27.
- Olwal, A., DiVerdi S., Candussi, N., Rakkolainen, I., Höllerer, T., 2006. An immaterial, dual-sided display system with 3D interaction. In: *Proceedings IEEE Virtual Reality Conference (VR'06)*, March 25–29 (Virginia, USA), IEEE, pp. 279–280.
- Planar Systems, Inc, 2014. Planar LookThru Series. Retrieved August 5, 2014 from (<http://www.planar.com/products/transparent-displays/>).
- Pronova, 2015. HoloPro, the Patented Solution for Living Projections on Glass. Retrieved January 31, 2015 from (<http://www.holopro.com/en/products/holopro.html>).
- Reetz, A., Gutwin, C., 2014. Making big gestures: effects of gesture size on observability and identification for co-located group awareness. In: *Proceedings ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, pp. 4087–4096.
- Rekimoto, J., Saitoh, M., 1999. Augmented surfaces: a spatially continuous work space for hybrid computing environments. In: *Proceedings ACM Conference on Human Factors in Computing Systems (CHI '99)*. ACM, New York, NY, USA, pp. 378–385.
- Samsung, 2014. 22" Transparent Display Showcase NL22B-T. Retrieved August 5, 2014 from (<http://www.samsung.com/ca/business/business-products/large-format-display/specialized-display/LH22NLBTVTC/EN>).
- Scott, S.D., Carpendale, S., 2010. Theory of tabletop territoriality. In: Müller-Tomfelde, C. (Ed.), *Tabletops - Horizontal Interactive Displays*. Springer (HCI Series), London, 375–406, (ISBN:978-1-84996-112-7).
- Scott, S.D., Carpendale, M.S.T., Inkpen, K.M., 2004. Territoriality in collaborative tabletop workspaces. In: *Proceedings ACM Conference on Computer Supported Cooperative Work (CSCW'04)*. ACM, New York, NY, USA, pp. 294–303.
- Shedroff, N., Noessel, C., 2012. Make it so: interaction design lessons from Science fiction. *Rosenfeld Media*.
- Shen, C., Everitt, K.M., Ryall, K., 2003. UbiTable: impromptu face-to-face collaboration on horizontal interactive surfaces. In: *Proceedings ACM Conference on Ubiquitous Computing (UbiComp'03)*. ACM, New York, NY, USA, pp. 281–288.
- Sridharan, S.K., Hincapié-Ramos, J.D., Flatla, D.R., Irani, P., 2013. Color correction for optical see-through displays using display color profiles. In: *Proceedings ACM Symposium on Virtual Reality Software and Technology (VRST '13)*. ACM, New York, NY, USA, pp. 231–240.
- Stefik, M., Bobrow, D., Foster, G., Lanning, S., Tatar, D., 1987. WYSIWIS revised: early experiences with multiuser interfaces. *ACM Trans. Inf. Syst.* 5 (2), 147–167.
- Sun, X., Liu, J., 2006. Light emitting material integrated into a substantially transparent substrate. US Patent 6986581 filed 3 Nov 2004, and issued 17 Jan 2006.
- Tang, J., Minneman, S., 1990. Videodraw: a video interface for collaborative drawing. In: *Proceedings ACM Conference on Human Factors in Computing Systems (CHI '90)*. ACM, New York, NY, USA, pp. 313–320.
- Tang, J., Minneman, S., 1991. VideoWhiteboard: video shadows to support remote collaboration. In: *Proceedings ACM Conference on Human Factors in Computing Systems (CHI '91)*. ACM, New York, NY, USA, pp. 315–322.
- Tang, J.C., 1991. Findings from observational studies of collaborative work. *Int. J. Man Mach. Stud.* 34, 143–160.
- Tang, A., Boyle, M., Greenberg, S., 2004. Display and presence disparity in mixed presence groupware. In: *Proc. Australasian User Interface Conf. Vol. 28*. Australian Computer Society, Inc., pp. 73–82.
- Vogel, D., Balakrishnan, R., 2004. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. In: *Proceedings ACM Symposium on User Interface Software and Technology (UIST'04)*. ACM, New York, NY, USA, pp. 137–146.
- Wikipedia, 2015. Scrim (material). Retrieved February 1, 2015 from (<http://en.wikipedia.org/wiki/Scrim>) (material).
- Wu, M., Balakrishnan, R., 2003. Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. In: *Proceedings ACM Symposium on User Interface Software and Technology (UIST'01)*. ACM, New York, NY, USA, pp. 193–202.