UNIVERSITY OF CALGARY

Two-Sided Transparent Display as a Collaborative Medium

by

Jiannan Li

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF SCIENCE

GRADUATE PROGRAM IN COMPUTER SCIENCE

CALGARY, ALBERTA

JANUARY, 2015

THE UNIVERSITY OF CALGARY

FACULTY OF GRADUATE STUDIES

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies for acceptance, a thesis entitled "Two-Sided Transparent Display as a Collaborative Medium" submitted by Jiannan Li in partial fulfillment of the requirements for the degree Master of Science.

____

Supervisor, Ehud Sharlin
Department of Computer Science

____

Co-Supervisor, Saul Greenberg
Department of Computer Science

____

Examiner, Sonny Chan
Department of Computer Science

____

External Examiner, Joshua Taron
Faculty of Environmental Design

____

Date

Abstract

Transparent displays are 'see-through' screens: a person can simultaneously view both the graphics on the screen and real-world content visible through the screen. Interactive transparent displays can serve as an important medium supporting face-to-face collaboration, where people interact with both sides of the display and work together. Such displays enhance workspace awareness, which smooths collaboration: when a person is working on one side of a transparent display, the person on the other side can see the other's hand gestures, gaze, and what s/he is currently manipulating on the shared screen. Even so, we argue that in order to provide effective support for collaboration, designing such transparent displays must go beyond current offerings. We propose using two-sided transparent displays, which can present different content on both sides. The displays should also accept interactive input on both sides and visually augment users' actions when display transparency is reduced. We operationalized these design requirements with our two-sided transparent display prototype, FACINGBOARD-II, and devised a palette of supportive interaction techniques. Through empirical studies, we found that the workspace awareness provided by transparent displays is compromised with degrading display transparency, and that visually enhancing user actions can compensate for this awareness loss.

Publications

Materials, ideas, and figures from this thesis have appeared previously in the following publications:

Li, J., Greenberg, S., Sharlin, E., and Jorge, J. (2014) **Interactive Two-Sided Transparent Displays: Designing for Collaboration**. In *Proceedings of the 2014 conference on Designing interactive systems (ACM DIS '14)*, 395-404.

Li, J., Greenberg, S. and Sharlin, E. (2014) **Enhancing Workspace Awareness on Collaborative Transparent Displays**. Research report 2014-1065-16, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada, October.

Li, J., Sharlin, E., Greenberg, S., and Rounding, M. (2013) **Designing the Car iWindow: Exploring Interaction through Vehicle Side Windows**. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)*, 1665-1670. (In Appendix B)

Acknowledgments

This thesis would not be possible without the help from many individuals.

To my supervisors, Dr. Ehud Sharlin and Dr. Saul Greenberg: thank you for leading me into the exciting realm of human-computer interaction, for your invaluable mentorship, insights, and support in and beyond research. It was the most fulfilling two years in my life so far working with you. Ehud, thank you for always encouraging and supporting me to go forward when I lost confidence or hesitated. Saul, thank you for showing me the true spirit of professionalism and perfectionism.

To all iLabbers: thank you for making Interactions Lab such a friendly, interesting, and creative place. I will remember our lunch table chats, games nights, pizzas on the old tabletop display, and helpful hands from everyone. I am proud that I started to learn the iLab style of presenting research, where the way of telling the story is as important as the story itself; and where aesthetics is as important as algorithms.

To my master buddies, Setareh, David (& Lindsay), and Bon: together we have had much fun and faced many challenges. I wish you all the best in your journey onwards.

To Xing-Dong: thank you for your help and your insights on life as a young researcher. I enjoyed the chats we had in your ‘hacker space’.

To Dr. Anthony Tang: thank you for your encouragement. I still remember the moment that I was in the media room, and you said to me, “You should speak more. Your opinions have their unique value.”

To Dr. Sheelagh Carpendale: thank you for showing me the beauty of information visualization and your generous help when I really needed it.

To my committee members, Prof. Sonny Chan and Prof. Joshua Taron: thank you for being on my committee and your insightful advice.

To my parents, Shuying and He: thank you for ever being my source of energy to wonder, to learn, and to explore.

Dedication

To my parents, Shuying and He.

Table of Contents

List of Figures

List of Tables

Chapter 1. Introduction

This thesis addresses the research question of how to design an interactive transparent display for collaboration. To set the scene, we present the context of our research and the overarching problem we studied. Next we introduce the motivation for choosing this particular research problem, the objectives set out, and the methodology we followed to achieve them. We conclude the chapter with the contributions of our research and an overview of this thesis.

1.1 Background

Small group collaboration activities are common in our daily lives: at work, at school, and at home; for business, for learning and for fun. Often these activities happen around a flat surface, such as a whiteboard, a table or, in this digital age, a computer monitor. People gather around the surface and use it as a convenient container of artifacts employed in work (Scott et al., 2004; Gutwin, 1997) and as a canvas to sketch on. In essence, people use the surface as the space to keep, present, and generate information during collaboration, for example, two designers critiquing a scheme lying on a table, a group of analysts studying data charts drawn on a whiteboard, etc. Recently people have started to collaborate over digitally-augmented surfaces, such as interactive whiteboards and tabletops, in order to take advantage of the potent storage, connectivity, and computation capability of modern computers. Such group activities characterize the context of this thesis, more precisely defined as *collaborative activities carried out by a small-size group of people (2 to 7) around an interactive surface*.

In such group activities, people naturally observe and comprehend others' actions in the workspace so as to coordinate themselves accordingly. In doing so, they make sure their actions serve the collective goal of the group. As an example, a person assembling puzzles with others will observe what pieces other group members are reaching for to avoid conflicts. This up-to-the-moment understanding of others' interactions with the workspace is called *workspace awareness* and acts as important glue for effective collaboration (Gutwin, 1997). Human-computer interaction (HCI) researchers have long been interested in supporting workspace awareness with interactive surfaces to improve the collaborative experience. However, most of the research on this topic concerned existing platforms such as desktop monitors and conventional large interactive surfaces (e.g. Hornecker et al, 2008; Moris et al, 2006; Sugimoto et al, 2004; Tse et al., 2006). This thesis explores an alternative approach that seeks to facilitate face-to-face collaboration with a novel set of interactive surfaces: interactive two-sided transparent displays (Figure 1.1).

Figure 1.1: Two people working face-to-face on an interactive two-sided transparent display.

1.2 Motivation: Why Collaboration on Transparent Displays?

Transparent displays are 'see-through' screens: a person can simultaneously view both the graphics on the screen and real-world content visible through the screen. 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. Researchers are promoting transparent displays in augmented reality applications, where the displayed graphics overlay what is seen through the screen, providing related information, for example by augmenting the real world through a mobile transparent display (Corning Inc., 2011 & 2012; Li et al., 2013; see Appendix B), or by adapting the content of a transparent display to the changing viewing perspectives of people who are moving around it (Olwal et al., 2005). Figure 1.2 present a collection of transparent

Figure 1.2: (1) a transparent display showcase (Humphries, 2010) (2) a transparent display integrated into the car windshield (Lysikatos, 2012) (3) a mobile phone using a transparent display as its screen (phoneArena.com, 2014) (4) a transparent display providing augmented reality information to the passenger in a vehicle (Li et al, 2013).

display products/prototypes: the top-left subfigure presents a large LCD transparent display showcase that displays promotional graphics for the product behind it (Humphries, 2010). The top-right subfigure depicts a display-enabled vehicle windshield that shows dashboard readings right in the driver's line of sight (Lysikatos, 2012). The bottom-left is a mobile phone using a transparent display as its screen (phoneArena.com, 2012). The bottom right demonstrates our work on a transparent display integrated into the side window of a car, superimposing augmented reality information on the surrounding environment for the passenger (Li et al., 2013; Appendix B).

As seen above, most existing transparent display systems are designed for scenarios where users see through and interact with only one side of the screen. We envision a *collaborative transparent display* that acts as a mediator between people working together across both its sides, with each side being fully interactive (e.g. Figure 1.1). We believe that such transparent displays can provide two collaborative benefits 'for free': when a person is working on one side of the transparent screen, people on the other side of it can both see the person, and be aware of what the person is working on. These benefits make it easier for one to perceive cues that help establish workspace awareness and relate them to the contents on the screen. In Figure 1.3 and 1.4 we present a series of mocked-up scenes of cooperative work on a transparent medium, in which two people can write on both sides

Figure 1.3: Two people working on a piece of transparent glass board with marker pens. The image show how the transparent medium helps them make use of **gaze** as an awareness cue.

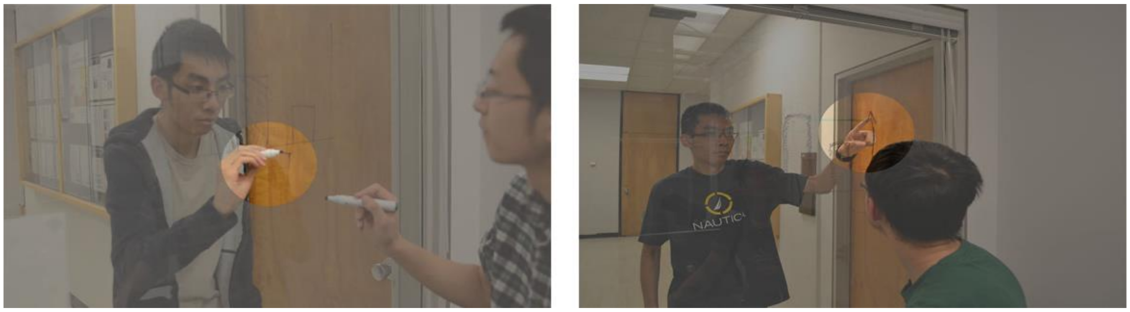

Figure 1.4: Two people working on a piece of transparent glass board with marker pens. The images show how the transparent medium helps them make use of **direct interaction** (*left*) and **deixis** (*right*).

of a piece of glass board with marker pens. These scenes take three important awareness cues—gaze, direct interaction, and gestures—as examples, and show how a collaborative transparent display helps people capture them.

Because of their unique properties, collaborative transparent displays can introduce new capabilities to collaborative environments. For example, they can be seamlessly integrated into existing windows and transparent walls now prevalent in workplaces. Video visions of the future ‘A Day Made of Glass’ released by Corning Inc. (2011 & 2012), for instance, illustrate a broad range of applications built upon display-enabled transparent glass in many different form factors, including a collaborative wall installed in a hospital (Figure 1.5, next page). Across this wall a surgeon in the sterile operation room can collaborate with his colleagues in the other non-sterile room, talking face-to-face while studying the medical imagery of the patient displayed on the transparent interactive wall. Collaborative transparent displays can also serve in data visualization analysis, gaming, tutoring, and many other applications. In this thesis, we focus our attention to the fundamental interface features and how it can support collaboration, and leave the exploration of the range of possible applications of collaborative transparent displays for future work.

1.3 Thesis Objective

The overall objective of this thesis is to **explore the design of collaborative transparent displays**. This overreaching goal is divided into the following three sub-objectives:

Figure 1.5: 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 (Corning Inc., 2011&2012).

First, **determine a number of design requirements for transparent displays to effectively support collaborative work.** Even though researchers have demonstrated transparent displays for collaborative activities (e.g. Heo, 2013; Olwal, 2008), as far as we know no prior work has gone beyond hardware specification and configuration. Therefore our first objective is to determine the design requirements articulating desired interface features of collaborative transparent displays.

Second, **develop the hardware and software that fulfills these requirements.** Because there is no transparent display specifically designed for around the guidelines from goal #1, we need to design and implement a test bed from ground up to help us uncover and address the challenges that arise in operationalizing the requirements we suggested.

Third, **investigate and evaluate the design via empirical studies**. Finally, we would like to understand particular challenges encountered in the design process and to validate our approaches to addressing these challenges through empirical studies. In particular, we focus on the degradation of workspace awareness in collaborative transparent displays that arises when display transparency is compromised by low lighting and displaying dense graphics, and how the techniques of visually augmenting a person's actions can overcome this degradation.

Achieving the first goal will deepen our understanding of the basic requirements behind a transparent display serving as a collaboration mediator. The second goal will operationalize the requirements, creating a platform worthy of critique. Satisfying the third goal will help identify significant usability issues and investigate the effectiveness of a possible solution.

1.4 Methodology

We used the following methodology to meet the research goals stated above:

For goal #1, drawing upon Computer-Supported Collaborative Work (CSCW) theories on workspace awareness and workspace territoriality, and our exploratory experiment with the first prototype, FACINGBOARD-I, we set out three design requirements for collaborative transparent displays: *interactive two-sided input*, *showing different content on both sides*, *and augmenting human actions*. We name transparent displays that can show different content on both sides **two-sided transparent displays**.

For goal #2, we created the second prototype FACINGBOARD-II, which aimed to realize the requirements determined in our first goal. It comprised a fabric-based two-sided transparent display, a finger tracking system, and demonstrative graphical user interfaces. We also explored various rich interaction features enabled by this setup—including visualization techniques that augment human actions—to expand the design space of collaborative transparent displays.

For goal #3, we conducted a controlled study to investigate how the capability of a two-sided transparent display to provide workspace awareness is affected by display transparency, and what is the efficacy of action augmentation techniques in compensating potential awareness loss. We analyzed the study results and present our interpretations in this thesis.

1.5 Contributions

This research provides five novel contributions to the state of the art of human-computer interaction and interface design of collaborative transparent displays:

First, this was the first research work providing general design guidelines for collaborative transparent displays.

Second, we contributed what was arguably the first interactive two-sided transparent display systems that can present different content on both its sides. As of today, only two other systems currently have this capability. Our system, published in June this year (Li et al., 2014), predates the later work of Lee et al. (2014). While there was one earlier work (Olwal et al., 2008) that has similar capabilities, it is an entertainment system that projects images onto fog rather than a screen. Our system construction is also novel: to our knowledge, the use of openly-woven fabric has not been explored before in implementing transparent displays.

Third, this was the first research work that explored interaction techniques supporting collocated collaboration on two-sided transparent displays. In particular, we devised novel techniques that leverage the unique collaborative benefits enabled by showing different content on both sides of the display. As we will see, this includes selective text and image reversal for legibility, private areas, semi-personal view of public objects, and two techniques that augment user actions when display transparency is compromised.

Forth, this research work was the first that investigated how transparency of transparent displays can be compromised, and how this in turn can severely affect workspace awareness.

Fifth, our work was the first to evaluate the efficacy of using visualization techniques that augment actions to compensate for awareness loss resulting from reduced display transparency.

The scope of this research focuses primarily on workspace awareness in collaborative transparent displays. While limited, we hope it will provide a foundation for future explorations exposing other prospects and issues concerning the utility and usability of such systems.

1.6 Frame of Reference

The two other collaborative transparent displays (Lee et al., 2014; Olwal et al., 2008) primarily focused on their technical implementation aspects along with proof-of-concept demonstrations involving a few simple (mostly playful) applications. Our own work—while also covering technical aspects and proof-of-concept applications— uses a broader frame of reference. It begins with low-level theories of collaboration, in particular, workspace awareness theories, which are used to motivate particular designs. The thesis subsequently evolves via the process of combining theories, design, and experiments to gradually develop our interaction design for collaborative transparent displays, as well as to critique some of the solutions found.

1.7 Thesis Overview

This document unfolds as follows. Chapter 2 reviews the related work and intellectual foundation of this research, in particular, transpernt display technologies, the importance of workspace awareness in cooperative environments, groups' partitioning behavior of workspaces, and how others have supported collaboration using transparent displays or transparent display metaphors.

Chapter 3 reports our prototyping efforts. We first introduce the exploratory prototype FACINGBOARD-I, the lessons learned from it, and the design requirements derived from combining these lessons with CSCW theories on workspace awareness and workspace territoriality (Contribution #1). Then we describes how we operationalized these design requirements with the second prototype, FACINGBOARD-II, including its hardware and software implementation details and the interaction features it enabled (Contribution #2 and #3).

Chapter 4 documents a thorough user study investigating how display transparency affects the capability of a two-sided transparent display to provide workspace awareness. The study also looked into the efficacy of the augmentation techniques we proposed to neutralize such influence (Contributions #4 and #5).

Chapter 5 discusses the results of the study and their implications (Contributions #4 and #5).

Chapter 6 concludes by reflecting on the contributions and limitations of this work, and laying out possible paths for future research.

Chapter 2. Related Work

This chapter seeks to familiarize our readers with the intellectual basis from which our later discussion on collaborative transparent displays can emerge. More specifically, it briefly introduces existing technologies, theories, and systems that have enlightened our research.

People has explored various technological approaches for building transparent displays. We review the mainstream of these approaches and discuss how they have informed the technological choices of our prototypes.

Next, we describe workspace awareness theories, which explains the mechanisms through which people stay aware of others' states in collaboration. We will use these theories throughout this thesis as a lens to examine cooperative activities and to guide our design of a collaborative transparent display.

We also discuss theories of workspace territoriality, which describes how collaborators partition a workspace. These theories have informed the other important aspect of our interaction design—supporting natural workspace partitioning.

Finally, we review prior research on supporting collaboration with transparent displays or transparent display metaphors.

2.1 Transparent Display Technology

In this section, we will review a broad range of technological approaches people have taken to make transparent displays (not including input technologies). Because of our focus, we will particularly make a distinction between technologies that enable two-sided transparent displays, i.e. displays that can present different content on both sides, and those that cannot. We begin with emissive display technologies, which generate images directly on the screen, followed by projection-based systems.

2.1.1 Emissive Transparent Displays

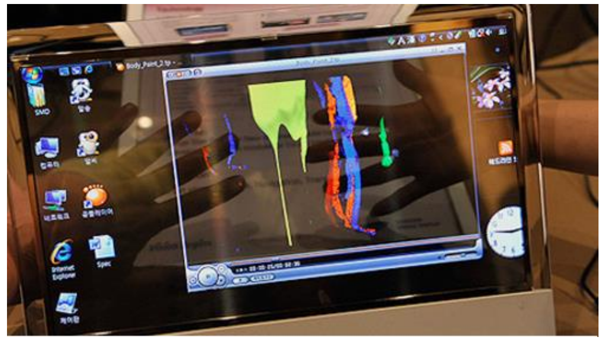

Figure 2.1: An OLED-based transparent display (EarlyTechNews, 2014)

LCD (liquid-crystal display) and OLED (organic light-emitting diode) are the two most common backbones for conventional displays such as TVs and mobile phone screens, and are unsurprisingly the most mature options to make transparent displays (see Figure 2.1 for an example). Several companies have already been marketing their showcase products incorporating transparent displays based on LCD or OLED (e.g. Samsung, 2014; Planar Systems, Inc., 2014). As the core of modern displays, they have both been improved for years. To date, LCD/OLED-based transparent displays offer the best overall image quality, with the same level of resolution and maximum display colors as conventional displays and better contrast ratio than, for example, projection-based alternatives. Discussing their technical details is beyond the scope of this thesis, but we will point out a few facts relating to making transparent displays. First, both LCD and OLED panels consist of display units integrated into a piece of highly transparent material. The light from the display units transmit to both sides of the panel, allowing its content to be visible to viewers on either side. The tiny display units, though not visible individually, make the panel hazy and less transparent. Second, OLED units generate light on their own while LCD units do not. Therefore LCD-based transparent displays must rely on strong ambient light or backlight units to illuminate the image.

Figure 2.2: A monochromatic transparent display based on liquid crystal (Kent Optronics, 2014)

Manufactures have also made monochromatic transparent displays with liquid crystal or electroluminescent display technology (e.g. Lumineq, 2014; Kent Optronics, 2014; see Figure 2.2). Though unable to render vivid colorful image as LCD/OLED displays do, they feature lower cost, higher transparency, and stronger resilience to environmental factors, lending themselves to industrial and outdoor conditions. In terms of overall structure, they resemble that of LCD/OLED displays, integrating display units into fully transparent material. Likewise they show the same content on both sides.

Figure 2.3: JANUS, a two-sided emissive transparent display making use of POV effect (Lee et al., 2014).

Lee et al. built an emissive transparent display, JANUS (2014), which differed from above as it could show different content on its two sides. Making use of the persistence-of-vision (POV) effect, JANUS displayed graphics by spinning a blade with an array of tri-color LEDs on each side at a high speed (Figure 2.3). 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 far as we know, JANUS was the first emissive transparent display that supported two-sided display capability[1]. As an early research prototype, its limitations include low-resolution, limited display area (the movement range of the blade), and cumbersome hardware.

2.1.2 Projection-based Transparent Displays

Many other transparent display systems were implemented through projection on a see-through panel. An immediate problem concerning this type of setup is much of the light from the projector penetrates through the transparent panel instead of being scattered, which leads to low image brightness. To alleviate this problem

Figure 2.4: TransWall, a projection-based transparent display. The content on both sides was the same. (Heo et al. 2013)

[1] JANUS is the most similar system to our work. As mentioned, it appeared after our own work, which, as we will see, is a projection-based display.

builders of these systems attached special projection films onto the panel, which hit a good balance between translucency and reflectiveness. The film was also required to diffuse the projected light rather than reflect it specularly, as most transparent panels do. Using diffuse refection, which sends light in all directions, helped to achieve a wider viewing angle and avoid annoying shiny glare. Commercial vendors of this type of film usually call it 'holographic projection film'. To further enhance image brightness, some people used two projectors to project precisely aligned images on both sides of the film, such as in TransWall (Heo et al. 2013; see Figure 2.4).

Figure 2.5: A transparent display that selectively scattered light at a particular wavelength (Hsu et al. 2014)

Because all currently available projection films works in a way that sacrifices display transparency to some extent for image brightness, researchers of material science suggested making display panels with special materials for better transparency. This includes frequency-conversion materials that convert projected ultraviolet light to visible light (Sun and Liu, 2006; Liu and Sun, 2006), or infrared light to visible light (Downing et al., 1996), and material that selectively scatter light of a particular wavelength while being almost transparent to other wavelengths (Hsu et al., 2014; see Figure 2.5). These unique materials were created by adding particular nanoparticles (e.g. Sun and Liu, 2006; Hsu et al., 2014) or metallic chemical elements (e.g. Downing et al., 1996) to glass. Featuring high transparency, the new materials may be the foundation of next-generation transparent displays, but currently they are still at early experimental stage and offer limited display capability. For example, the display prototypes of Downing et al. (1996) and Hsu et al. (2014) only supported a limited number of colors.

There are other approaches that use immaterial screens as the projection medium. Fog display systems form a 'wall' for projection by trapping suspended particles, such as water droplets, in between two thin sheet of air. The fog of particles moves in a non-turbulent flow (laminar flow) so that people can see stable projected image (see Figure 2.6).

These immaterial displays have no clearly visible boundaries with surrounding environment, creating the illusion of image floating in air. Because very little of the projected light is reflects by the fog layer back towards the projection source, the image is primarily visible to viewers on the opposite side of the projector, i.e. those viewing rear-projected image (DiVerdia et al., 2006;). Therefore, fog displays are able to show different contents on both sides, if two projectors are used, one per side. Because of their vapor state, the image quality of fog displays is limited in terms of resolution, contrast, and stability.

Figure 2.6: Consigalo, a two-sided transparent display based on FogScreen^{TM} (Olwal et al. 2008)

Another transparent display system design that can show independent content on both sides is described in a patent Hewlett-Packard recently received (Kuo et al., 2013). The display is composed of two separate sets of mechanical louvers, which can be adjusted so that observers could see through the spaces between them. At the same time, light can be directed on each set of louvers, thus presenting different visuals on each side. Their approach differs significantly from others with respects to material choice. The panel that is projected onto is not a piece of uniformly transparent material, but consists of interwoven hollow and opaque areas. Two opposite sides of the opaque areas can show independent projected images. Such panels cannot support optimal display resolution but they afford the capability of showing different content on both sides.

After reviewing the approaches above and experimenting with some of them, we have devised our own solution for building two-sided transparent displays. The design and implementation details are illustrated in the next chapter.

2.2 Workspace Awareness

Workspace awareness theories try to explain how people gather and process information from the surrounding environment to coordinate themselves in collaboration. We have used them to systematically analyze cooperative activities and make informed design decisions throughout this project. In this section, we will explain the role of workspace awareness in collaboration and the mechanism by which people gather awareness information, and how collaborative transparent displays can support these mechanisms.

2.2.1 Workspace Awareness in Collaboration

In our everyday activities, people naturally stay aware of their surrounding environments and respond accordingly. For example, before crossing a street, a pedestrian would check the traffic conditions and stop if there are vehicles approaching. These behaviors are usually too common and natural to be consciously noticed but are critical for people to perform all kinds of tasks. Human factors research studied how this knowledge of the changing environment was availed in highly dynamic and information-rich environments, such as air combat. They called it "situation awareness" (Endsely, 1995). Situation awareness comprises 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. In the above pedestrian example the person first looks at the road and listens to the sound of engines. S/he then estimates the distance and speed of the coming vehicle and decides to halt or proceed. Like in many other activities, such a process does not occur just once; because the traffic conditions are constantly changing, it keeps running until the person safely arrives at the other side.

Researchers in the computer-supported cooperative work (CSCW) community developed a similar concept of awareness involving information sharing, knowledge of group and individual activity, and coordination in a *shared workspace* (Dourish and Bellotti, 1992). A *shared workspace* refers to the shared space where group work (e.g. design sessions and business planning) is carried out, such as a table or a whiteboard. The notion of awareness in CSCW captures the similar idea of knowing what is going on in the workspace, where people receive, generate and modify information.

When someone is working alone, his or her awareness only involves the workspace and the domain task. If others join to form a group, the person has to take on a third type of awareness into account, that of their co-workers, if s/he is to benefit from effective collaboration. It includes not only who is in the group but also much richer real-time information as to the fine-grained actions of people relative to the workspace. This up-to-the-moment understanding of another person's interaction within a shared setting is the *workspace awareness* that feeds effective collaboration (Gutwin et al, 1996; Gutwin and Greenberg, 1998, 2002). It provides knowledge about the 'who, what, where, when and why' questions whose answers inform people about the state of other active group members: 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? Knowing these states allows people to coordinate with the group accordingly. In essence, workspace awareness serves as the glue that integrates individual contributions into collective productivity. Gutwin and Greenberg (2002) stress that workspace awareness plays a major role in various aspects of collaboration, which are listed as follows.

- *Managing coupling*. People are often engaged in mixed-focus collaboration, where they shift back and forth between loosely and tightly-coupled group work. In loosely-coupled work, their focus of attention are primarily on individual work, while monitoring others' activities for opportunities of collaboration (tightly-coupled work). Awareness helps people perform these transitions. For example, two mechanics installing a machine together are assembling individual parts on their own, while watching for each other's progress so that they can connect components properly.
- *Simplification of communication*. Because people can see the non-verbal actions of others, dialogue length and complexity is reduced.
- *Coordination of action*. Fine-grained coordination is facilitated because one can see exactly what others are doing. This includes who accesses particular objects, handoffs, division of labor, how assistance is provided, and the interplay between peoples' actions as they pursue a simultaneous task.
- *Anticipation* occurs when people take action based on their expectations or predictions of what others will do. Such predictions are largely informed by observing others'

actions and hearing their verbal utterances. Anticipation helps people either coordinate their actions, or repair undesired actions of others before they occur.

- *Assistance*. Awareness helps people determine when they can help others and what action is required. This includes assistance based on a momentary observation (e.g., to help someone if one has 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.

2.2.2 Gathering Workspace Awareness Information

Researchers have identified three main sources of workspace awareness information and three corresponding mechanisms of information gathering: consequential communication for bodies, feedthrough for artifacts, and intentional communication for conversation and gestures (Gutwin and Greenberg, 2002).

Consequential Communication

Much of one's interactions with the workspace are carried out through bodily actions—holding a pen and writing on a piece of paper, reaching out the arm to fetch a tool, turning the head to examine a certain spot, etc. Others who observe and comprehend these actions can naturally derive a wealth of information as to answering the "who, what, where, when and why" questions about the actor. This mechanism of information transferring as a consequence of one's activities is called *consequential communication*. While observed actions are not undertaken for communication purposes, the observers understand these actions, interpret them, and become informed.

Feedthrough

Artifacts are objects that people make use of in the workspace. Artifacts can inform viewers of the actions and intentions of their users: for example, pencils signal drawing or writing, rulers signal measuring, and piles of files signal reading or sorting. The characteristic sound that artifacts make during their usage can also reveal what is being manipulated and how. For example, the scraping sound of a pencil tip against paper and the snipping sound of scissors indicate that these artifacts are being used. In sum, the movement and sound of an artifact not only sends direct feedback to its user, but also give off information about that user's actions on workspace artifacts to observers. This mechanism of providing feedback to the observer is called *feedthrough* (Dix, 1994). In a computer-supported workspace, physical artifacts are commonly replaced by virtual graphical user interfaces (GUIs). Similarly, GUIs with characteristic appearance and sound can provide feedthrough to people in the workspace.

Intentional Communication

Awareness information is also transferred through intentional communication when people explicitly express their thoughts via verbal utterances and gestures. In explicit interpersonal conversation, they state their opinions, comment on others, request help etc. People also use *outlouds,* where they 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). Gestures are another form of intentional communication. Previous research has pointed out their significant role in group communication (Tang, 1991). People employ a broad class of gestures, such as *deixis*, where a pointing action qualifies a verbal reference (e.g., 'this one here'), and *demonstrations* that embody abstract concept through hand movements.

2.2.3 Workspace Awareness and Collaborative Transparent Displays

Collaborative transparent displays can provide support for all three mechanisms to maintain workspace awareness in group work.

Across collaborative transparent displays, one can readily see the position, posture, and movement of the person on the opposite side, having easy access to elements in *consequential communication*. These elements include *gaze*, namely where one is looking, which gives off information about the person's current focus of visual attention, and *visual evidence*, which confirms that an action requested by another person is understood by seeing that action performed. *Feedthrough* largely merges into consequential communication as people can see the person on the other side and the GUIs that person is manipulating at the same time.

Collaborative transparent displays promote *intentional communication* by helping establishing the connection between the verbal or gestural expression and the context within which it is interpreted, which is usually the content on the screen. On collaborative transparent displays, the person who speaks or gestures and the content related are in sight at the same time for the viewer.

In sum, collaborative transparent displays can support workspace awareness by arranging both the person carrying out actions and the visual workspace (within which the actions are interpreted, in the same field-of-view). Our interaction design, which will be presented in later chapters, has further explored techniques to ease all three mechanisms for awareness information gathering, especially in conditions where transparency is compromised.

2.3 Territoriality in Shared Workspaces

Territoriality theory describes how group members partition the shared workspace into zones of different uses. As we will see, our interaction design tries to support such partitioning behavior on collaborative transparent displays.

During collaborative activities, people use tacit zones located at different relative positions to them in the workspace for different purposes. Generally, these zones allow for efficient usage of space (Tang et al., 1991). For example, at small distances from a workspace area (e.g., meters), zones are often defined by social protocols about interpersonal proxemics (Hall, 1966): the closer one is to the workspace area, the more that space becomes one's own. When people surround a workspace, such as in tabletop collaboration, three types of territories can be identified (Scott et al., 2004)—personal, public, and storage, each with 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 *group territory* is the area where group members share access to it, usually to collectively pursue the main collaborative task. It usually takes up the space that is not occupied by personal territories. A *storage territory* serves as the area to store task resources and typically sits atop both personal and group territories. Similar partitions of personal versus group can be found on vertical workspaces as well (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 seeking to accommodate and further enforce people's partitioning behavior, such as those described by Rekimoto et al. (2002), supported private territories with devices separated from the shared workspaces (e.g. PDAs and laptops) so that only their owners could see and manipulate that 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 (see Section 2.2.2). Later groupware designers paid particular attention to the subtle distinction between private, personal, and public territories. For example, Wu et al.'s RoomPlanner (2003) had no permanent private territories. However, it supported a gesture that could temporarily make a personal territory private. To perform the gestures, a user placed the side of his or her hand on the tabletop to form a horizontal line, blocking others from seeing the area behind it. The gesture could trigger displaying private information, or allow for private voting. UbiTable by Shen et al. (2004) went even further by providing designated private, personal, and public territories. Private territories were workspaces on individuals' laptops. Personal territories covered areas on the tabletop that ware close to each group member, visible but not modifiable to others. Public territories sat around the center of the tabletop and were shared by all group members.

The above work suggests that collaborative transparent displays should have areas with different levels of accessibility and visibility. Our solution provides private and public areas, and semi-personal states of public content. It will be detailed in the next chapter.

2.4 Supporting Remote Collaboration using Transparent Display Metaphors

In the field of CSCW, researchers explored using transparent display metaphors for remote collaboration, where network-connected remote collaborators were presented with the illusion of working on the two sides of a virtual transparent display. Although not involving physical transparent displays, their works suggested the potential of facilitating collaboration with transparent displays.

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 (see Section 2.2) was compromised. For example, if a person dragged an object from one place to another, the partner would just see it disappear from its old location and re-appear at its new location. Because the partner could not see the other person's body, both consequential communication and intentional gestural communication was unavailable.

Some researchers tried to provide this missing information by building special purpose awareness widgets (e.g., Gutwin et al., 1996), such as multiple cursors as a surrogate for gestural actions. Others sought a different strategy: a simulated 'see-though' display for remote interaction. The idea began with Tang and Minneman (1990; 1991), who developed two video-based systems, VideoDraw and VideoWhiteboard. 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 drawings 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 (see Figure 2.7). 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

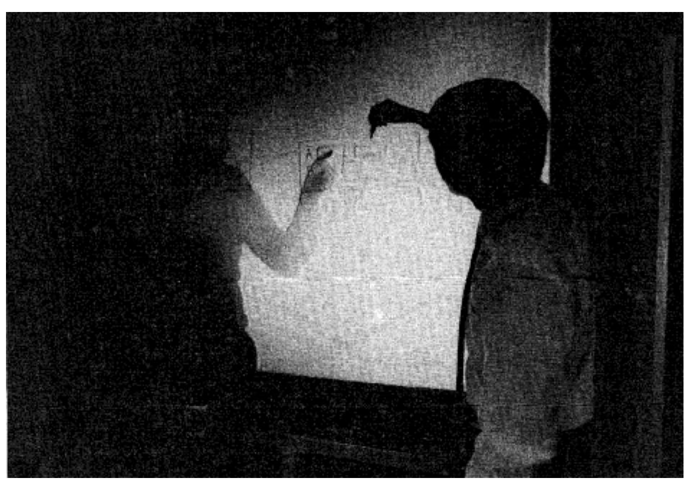

Figure 2.7: A person works with a remote collaborator on VideoWhiteboard (Tang and Minneman, 1991)

a big transparent glass board", culminating in the ClearBoard II system (Ishii and Kobayashi, 1992). As illustrated in Figure 2.8, 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) were both fully digital 'mixed presence' groupware system that connected two large touch-sensitive surfaces, and included 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.

Note that on ClearBoard, the pair of remote users could share a common orientation of drawings on the screen because they were presented with a mirror-reversed video feed of the collaborator. The same solution does not apply to actual transparent displays, where the problem of reversed content orientation will arise. With two-sided transparent displays, which can present different graphics on both sides, the problem can be solved by selectively reversing shared content on the screen. We will discuss this advantage of two-sided transparent displays, along with their other collaborative benefits, in the next chapter.

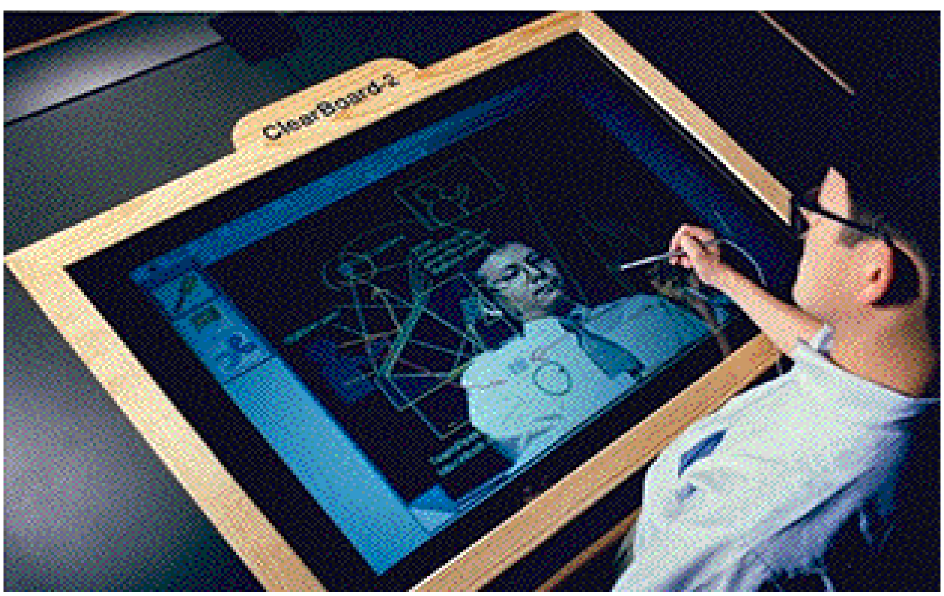

Figure 2.8: ClearBoard II (Ishii et al., 1992)

2.5 Supporting Collocated Collaboration Using Transparent Displays

Supporting collocated collaboration on a physical interactive transparent display is still an emerging and less explored realm.

Dating back to World War II, air traffic controllers used to write field information on both sides of glass plotting boards (Figure 2.9). They did so to reduce interference between the controllers that wrote closely to each other on the surface, demonstrating another benefit (though not the focus of this research) of using transparent displays for collaboration—expanded input space.

Ishii and Kobayashi (1992) started their exploration of the ClearBoard project with a preliminary prototype, ClearBoard-0, which was largely similar to the plotting board that air controllers historically used. They were interested in problems and prospects of such displays as a metaphor for remote collaboration, and moved to video-based systems for connecting spatially distributed collaborators.

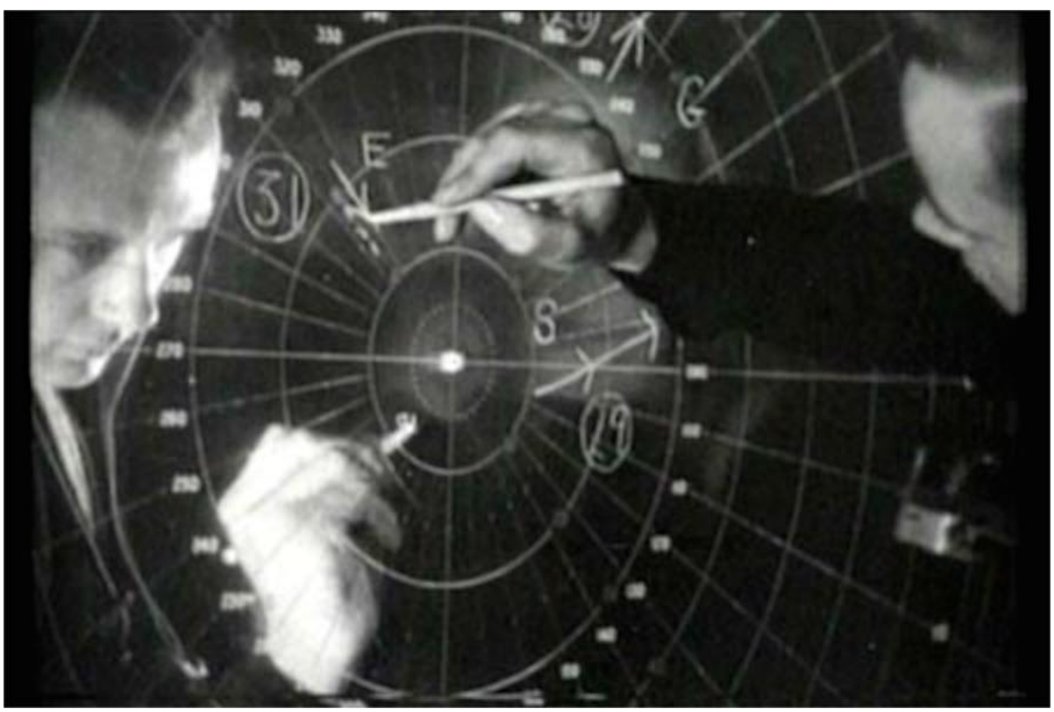

Figure 2.9: Air controllers writing on both sides of a transparent plotting board. (Gessler, 2014)

Recently researchers started to build interactive transparent displays for collaboration which allowed for direct input on both sides. Olwal et al. (2008) built Consigalo (Figure 2.6), a collaborative gaming system based on an immaterial see-through display, FogScreen^{TM}. It could show different content on both sides. Input was done via three degree-of-freedom position tracking of LEDs held by people and tracked by infrared cameras. People on either side of the display could play a collaborative game against the other group on the opposite side, in which they competed by picking up falling shapes projected on the screen. Though Olwal et al. stressed playfulness added by face-to-face communication, they mainly focused on technical implementations and did not go into details of interaction design.

Heo et al. (2013) demonstrated TransWall, a see-through display whose capabilities were illustrated by various entertainment applications (Figure 2.4). It used two projectors to provide an identical bright image on both sides. Two infrared touch sensor frames mounted on either side collected multiple touch inputs per side and allowed people on either side of the display to interact via direct touch. The system also included acoustic and vibro-tactile feedback, as well as a speaker/microphone that controlled the volume levels of the conversation passing through it. With its high-quality image and rich output, TransWall enabled an array of collaborative drawing and musical games, some of which highlighted the unique benefits of collaboration on transparent displays. However, they did not consider the demands and challenges in general cooperative tasks. TransWall also showed identical content on both sides, entailing problems such as that text on one side would appear reversed on the other side.

Our work builds upon the works above, with notable differences. From a technical stance, it allows different images to be projected on either side. From a collaborative stance, it focuses on supporting workspace awareness within such see-through two-sided interactive displays, especially in cases where the ability to see through the display is compromised.

2.6 Summary

In this chapter, we reviewed related research and commercial products that had informed our technological and interaction design for collaborative transparent displays.

We first reviewed technological approaches for making transparent displays, including technologies based on emissive displays and on projection. We particularly examined whether and how these approaches enable two-sided transparent displays, which can present different content on both sides. OLED/LCD-based technologies overall offer the best graphics quality but they are not applicable for two-sided transparent displays. The existing approaches that support the capability of showing different content on both sides include emissive displays using the persistence-of-vision effect and projection-based systems that project graphics on fog displays or mechanical louvers. Our own approach for making two-sided transparent displays will be introduced in the next chapter.

We then explained workspace awareness theories. Workspace awareness is people's up-to-the-moment understanding of others' interactions within the shared workspace. People naturally make use of this knowledge to coordinate themselves with other group members in collaboration. There are three primary sources of workspace awareness information and three corresponding mechanisms for information gathering: consequential communication for bodily actions, feedthrough for artifacts, and intentional communication for conversations and gestures. We discussed how collaborative transparent displays can facilitate these mechanisms. We will use these theories to design appropriate interactions for collaborative transparent displays in the following chapter.

We introduced theories of workspace territoriality, which illustrates people's natural behavior of partitioning workspaces into zones and using them for different purposes. We described four typical territories found in shared workspaces—public, personal, private, and storage territories and their properties. We also reviewed interactive systems that permitted various levels of partitioning. As we will see, based on prior research, our interface will be designed to support workspace partitioning on collaborative transparent displays.

We described earlier research works that sought to support remote collaboration with transparent display metaphors using network-transmitted video and more recent works that designed for collocated collaboration with physical transparent displays. While these systems inspired us, we hinted at how our own work differs from them from both technical and collaborative stances.

Chapter 3. Designing Collaborative Transparent Display Interfaces

With related theories, systems, and technologies in mind, we set out to explore the design of collaborative transparent displays. This chapter documents our thinking and prototyping efforts. We first describe our preliminary prototype, FACINGBOARD-I, and the results from an informal evaluation of this prototype. Then we discuss three requirements we suggest for transparent displays to provide sufficient support for cooperative work. Finally, we present our second prototype, FACINGBOARD-II: its hardware/software configuration and interaction techniques, all of which were devised to address the proposed requirements.

3.1 Early Exploration: FACINGBOARD-I

As the first step to approach the goal of designing collaborative transparent displays, we built the first prototype, FACINGBOARD-I, to obtain first-hand experience and probe into the design space.

FACINGBOARD-I was a straightforward and preliminary realization of the concept that two people work together on both sides of a transparent display and interact with it using direct manipulation. Based on a transparent LCD display, it showed the same content on both sides. We developed four collaborative sketch tools and a two-player game to experiment with the setup and conducted an informal evaluation to solicit feedback from people.

3.1.1 Implementation of FACINGBOARD-I

The main body of FACINGBOARD-I consisted of a 22 inch Samsung LCD transparent display and two Leap Motion sensors on its two sides (Figure 3.1, next page). As described in Chapter 2, the LCD transparent display provided high-quality image, but the brightness

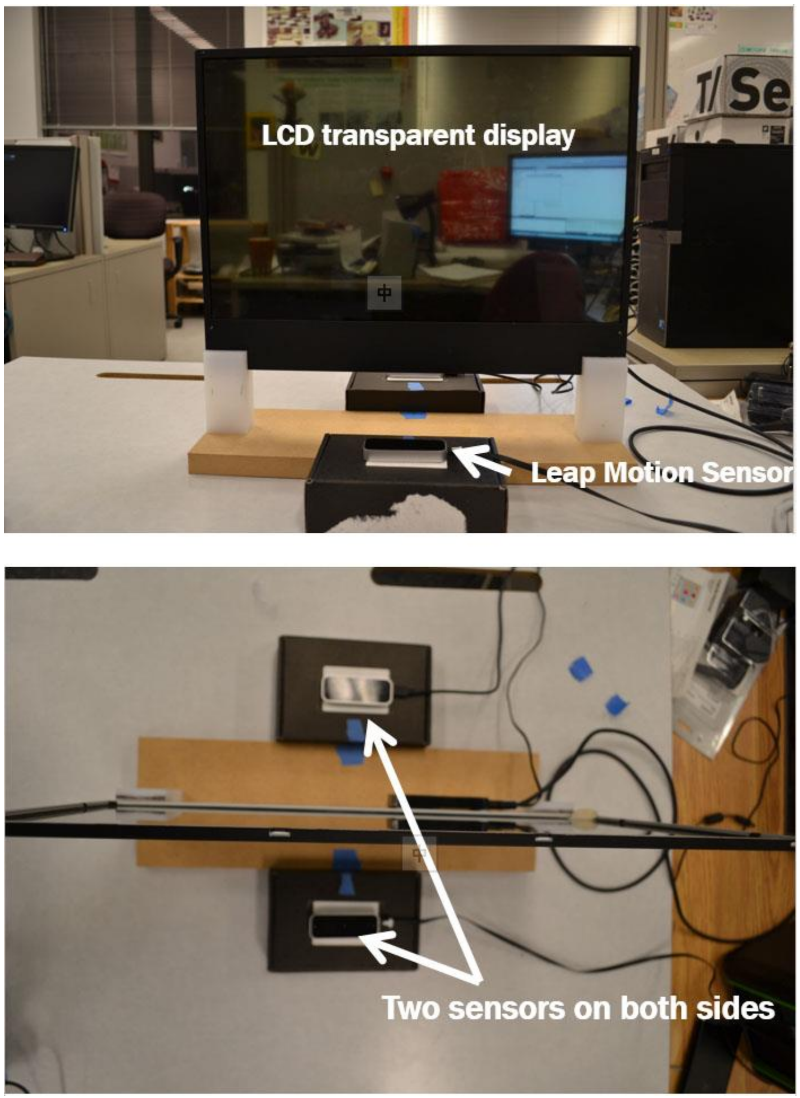

Figure 3.1: (*top*) The setup of FACINGBOARD-I prototype. The Leap Motion sensor tracks finger position. (*bottom*) One sensor is positioned on each side and tracks the interaction on this side.

of its image depended on ambient lighting. The images rendered on both sides were exactly the same. A Leap Motion sensor was used on each side of the display in order to track 3-dimentional finger positions. Because of its limited capture volume one sensor only tracked one person's finger movement on one side. The two sensors were positioned distantly to prevent cross talk.

The transparent display and one of the sensors were connected to a desktop computer running a C#-based display controller program. The display acted as the secondary monitor of the desktop. Because one computer could only control one Leap Motion sensor directly, the other sensor was connected to another desktop computer, which communicated with the main program via a network. Although Leap Motion sensors could track all fingers of two hands, to simplify prototyping we assumed the user only touched the display with his/her pointing finger, and that two fingers were used, at most. The sensors were positioned at a fixed location in relation to the display. A touch point was registered when the fingertip was close enough to the screen, after which the program translated the 3-dimensional fingertip position to a 2-dimensional screen coordinate. Several filtering algorithms were incorporated to filter out noise and reduce misrecognition of other long and thin objects, whose 3-dimensional shapes were similar to that of a finger.

3.1.2 Creating and Playing with FACINGBOARD-I

In order to learn people's reaction to this rather uncommon form of interaction and collaboration, we implemented several tentative applications for FACINGBOARD-I. This included four collaborative sketching tools: *pencil*, *note*, *two-finger line*, and *four-finger quadrilateral*. We also implemented a two-player game, *Pac-chaser*. The purpose of these applications was not to support practical tasks but rather to elicit users' reactions and to inform interface design possibilities.

The *pencil* tool allowed the person on either side to draw on the canvas freely with the finger (Figure 3.2a, next page).

The *note* tool supported note sharing between people on the two sides. The user of note first defined a virtual note visually similar to a paper Post-it™ with a pinch gesture. Then s/he could create text and graphics on this area. Since the content appeared as reversed for the person on the opposite, a "reverse" button was added aside the virtual note and tapping it flipped the content horizontally (Figure 3.2b, next page). The *note* tool was created to mitigate the reverse orientation problem hindering information sharing, especially text sharing between two parties across FACINGBOARD-I, which only showed the same image on both sides. However, it was a partial solution as whether flipped or not the content was inevitably reverse for one of the two sides.

The *two-finger line* tool, borrowed from early groupware GroupSketch (Greenberg et al., 1992), drew a line segment on the canvas connecting two ends specified by one touch

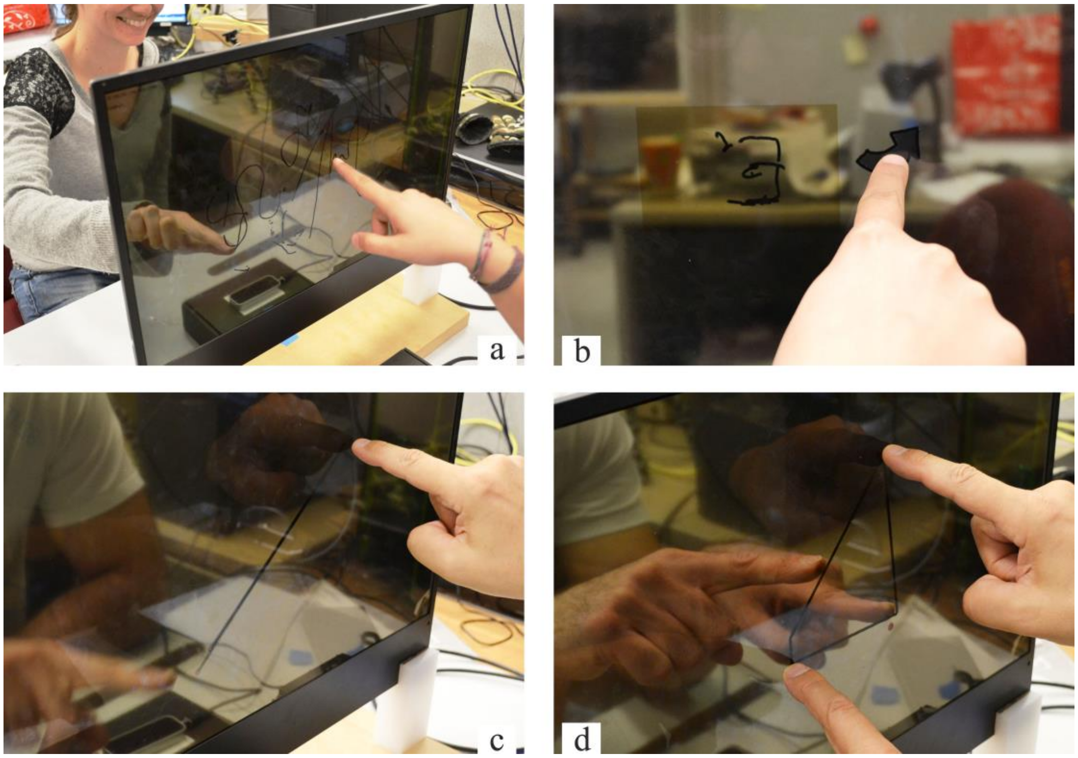

Figure 3.2: (*a*) Two people are drawing with the *pencil* tool. (*b*) A person is touching the "reverse" button of the *note* tool. (*c*) Two people are interacting with the *two-finger line* tool. (*d*) Two people are interacting with the *four-finger*

point on each side of the display (Figure 3.2c). Two people could move both end points simultaneously and the line changed accordingly. As a tool operated by cooperative gesture, *two-finger line* attempted to encourage participation, collaboration and awareness (Morris et al., 2006). Its variation, *four-finger quadrilateral*, drew a quadrilateral whose four vertices were specified by two touch points on each side of the screen (Figure 3.2d).

Besides the collaborative sketching tools above, we also designed a simple two-player game, *Pac-chaser*, to investigate the prospect of using FACINGBOARD-I for multi-player video games. *Pac-chaser* was inspired by the classic video game, Pac-man. In the game, a player moved, rotated and resized his/her Pac-man-shaped avatar on the display, trying to grow larger than the opponent, chasing it, and eventually swallowing it (Figure 3.3, next page). The player needed to be careful to avoid obstacles as collision led to losing the game.

Figure 3.3: Two people are playing the *Pac-chaser* game.

3.1.3 Informal Evaluation and Discussion

Seven graduate students from a computer science laboratory were invited to try and comment on FACINGBOARD-I. They were not given any specific tasks, as the purpose was to see people's reaction to this interaction form in general. The feedback they provided was mixed but overall enlightening:

"Quick and direct feedback (from the other person)"

"I like that you can see your opponent's face" (in Pac-chaser)

As anticipated, most participants liked the fact they could see their co-workers through the display and they considered such awareness as quick and direct. Some of them particularly stressed that it was fun to see the other's face in Pac-chaser. In contrast with conventional video games where players' attention is directed at the screen, FACINGBOARD-I enabled players to see each other's facial expressions readily. These expressions of concentration, or delight, or anxiety provided emotional feedback to players, enriching their game experience and connecting them more tightly.

"It's not easy to do real tasks on it"

The participants found the sketch tools on FACINGBOARD-I did not provide full support for practical cooperative tasks, mostly because the text written on one side looked reversed on the other side. While the *note* permitted content flip, people did not want to be constrained in the *note* area. Furthermore, information on *note* was only legible for one person, which prevented it from being an information container both parties could refer to simultaneously. However, in actual cooperative activities, people regularly rely on such containers for discussion and demonstration of ideas. Lack of such a shared information container rendered FACINGBOARD-I's support of collaboration incomplete.

"Sometimes I cannot see you very well, especially when there's something on the screen"

Some participants pointed out that the transparent display did not always look 'transparent'. As introduced in Chapter 2, the LCD transparent display was slightly hazy because of the display units embedded in the panel. Once overlaid with displayed graphics, hands and faces seen through the hazy panel became less clear. Our later observation (see Chapter 4) confirmed the visual awareness of other people through a transparent display is not always guaranteed. It is subject to the transparency of the display material, the density of graphics shown on the display, and environmental factors such as ambient lighting, and others.

The evaluation confirmed the unique benefits of collaboration on transparent displays for providing workspace awareness. It also revealed two limitations of FACINGBOARD-I that could impede collaboration: 1) only showing the same content on both sides and 2) compromised awareness. The findings suggest critical problems to be addressed in following design efforts.

3.2 Design Rationale for Collaborative Transparent Displays

Reflecting on the lessons learned from FACIGNBOARD-I we set out to determine a list of requirements for a collaborative transparent display. Such a display should support people's routine actions in practical group work, where they generate, modify, and refer to information in textual and graphical form, and provide sufficient workspace awareness regardless of content on the screen and environmental conditions. For this purpose, we advocate using a fully two-sided transparent display that enables independent input and output on both sides, and that augments less visible human actions. We formulate this as three design requirements*: two-sided interactive input, different content on both sides, and augmenting human actions*. We discuss each of these themes in the sub-sections, below.

3.2.1 Two-Sided Interactive Input

Collaboration is central to this design. All people—regardless of what side of the display they are on— are active participants. As with earlier systems supporting remote collaboration, we expect each person to be able to interact simultaneously with the display (Figure 3.4). From a workspace awareness perspective, we expect people to see each other through the screen and each other's effects on the displayed artefacts.

While such systems could be operated with a mouse or other indirect pointing device, our stance is that workspace awareness is best supported by direct interaction, e.g., by touch

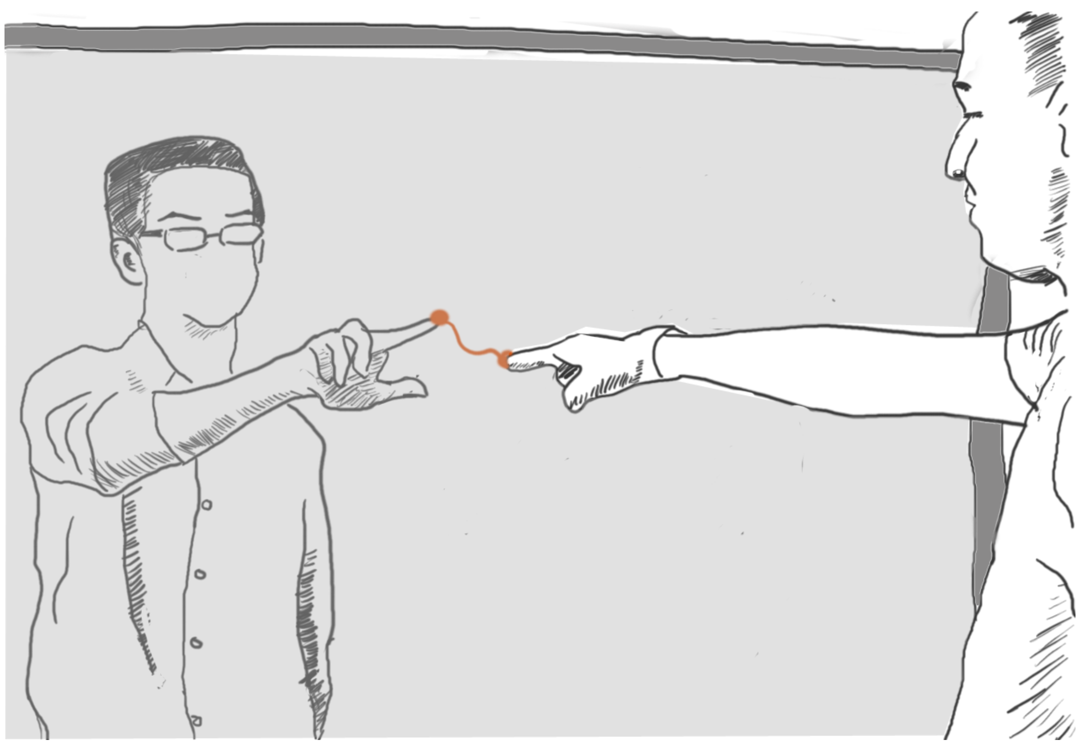

Figure 3.4: Two people interacting with the display simultaneously with direct touch

and gestures that people perform relative to the workspace as they are acting over it. Thus if people are able to see through the display, they can gather both consequential and intentional communications relative to the workspace, e.g., by seeing where others are touching, by observing gestures, by seeing movements of the hands and body, by noticing gaze awareness, by observing facial reactions.

3.2.2 Different Content on Both Sides

Except the FogScreen™ vapor display (Olwal et al., 2008) and JANUS (Lee et al., 2014), see-through displays universally show the exact same content on either side (albeit one side would be viewed in reverse). We argue for a different approach: while both sides of the display will mostly present the same content, different content should be allowed (albeit selectively) for a variety of reasons as listed below. Within CSCW, this is known as relaxed WYSIWIS (relaxed what-you-see-is-what-I-see). Figure 3.5 illustrates this concept: when a person is moving a triangle-shaped marker on a map, people on her opposite side of the display may see she is moving a circle-shaped marker.

Managing attenuation across the medium. Depending on the technology, image clarity can be compromised by the medium. For example, Olwal et al. (2008) describe how their FogScreen™ diffuses light primarily in the forward-direction, making rear-projected imagery bright and front-projected imagery faint, thus requiring two projectors on either side. In our own experiences with FACINGBOARD-I (which was LCD-based), image contrast was poor. One solution is to display content on both sides, rather than relying on

Figure 3.5: (*left*) Lisa is touching an orange triangle on her side of the display. (*right*) People on the other side of the display see a blue circle at the location where Lisa is touching.

the medium to transmit one-sided content through its semi-transparent material. This solution was adopted by Heo et. al. (2013) in their TransWall system to maintain image brightness, where both projected images were precisely aligned to generate the illusion of a single common image.

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

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

Private work areas. As introduced in Section 2.3, shared workspaces can include private work areas. These are valuable for a variety of reasons. For one, they could collect individual tools that one person is using. During loosely-coupled work, they could hold information that a person is gathering and working on, but that is not yet ready to show to others. They could even hold private information that one does not wish to share. A two-sided display allows for both shared and private work areas. For example, an area of the screen (aligned to each other on either side) can be set aside as a private work area, where the content on each side may differ. Workspace awareness is still partially supported: while one may not know exactly what the other is doing in their private area, they will still be able to see that the other is working in that area.

Feedback vs. feedthrough. In many digital systems, people perform actions quite quickly (e.g., selecting a button). Feedback is tuned to be meaningful for the actor. For example, the brief change of a button's shading as it is being clicked or an object disappearing as it is being deleted suffices as the actor sees it as he or she performs the action. Alternately, pop-up menus, dialog boxes and other interaction widgets allow a person to perform extended interactions, where detailed feedback shows exactly where one is in that interaction sequence. Yet the same feedback may be problematic if used as feedthrough (see Section 2.2.2) 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. 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 enables that different feedback and feedthrough mechanisms can be tuned to their respective audience.

Personal state. Various widgets display their current state. Examples include checkboxes, radio buttons, palette selections, content of textboxes, etc. In groupware, each individual should be allowed to select these controls and see these states without affecting the other person, e.g., to select a drawing color from a palette. A two-sided relaxed WYSIWIS display allows a widget drawn at identical locations to show different states that depend upon which side it is on and how the person on that side interacted with it. For example, a color palette may show the currently selected color as 'blue' on one side, and 'orange' on the other. Personal states do not align exactly with the conventional notion of personal territories in shared workspace as they are not visible to others; however, they provide a subtle middle ground between fully private and public, as private auxiliaries associated with public widgets.

3.2.3 Augmenting Human Actions

As we saw on FACINGBOARD-I, despite their names transparent displays are not always transparent. They all require a critical trade-off between the clarity of the graphics displayed on the screen vs. the clarity of what people can see through the screen. Factors that affect transparency include the following.

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

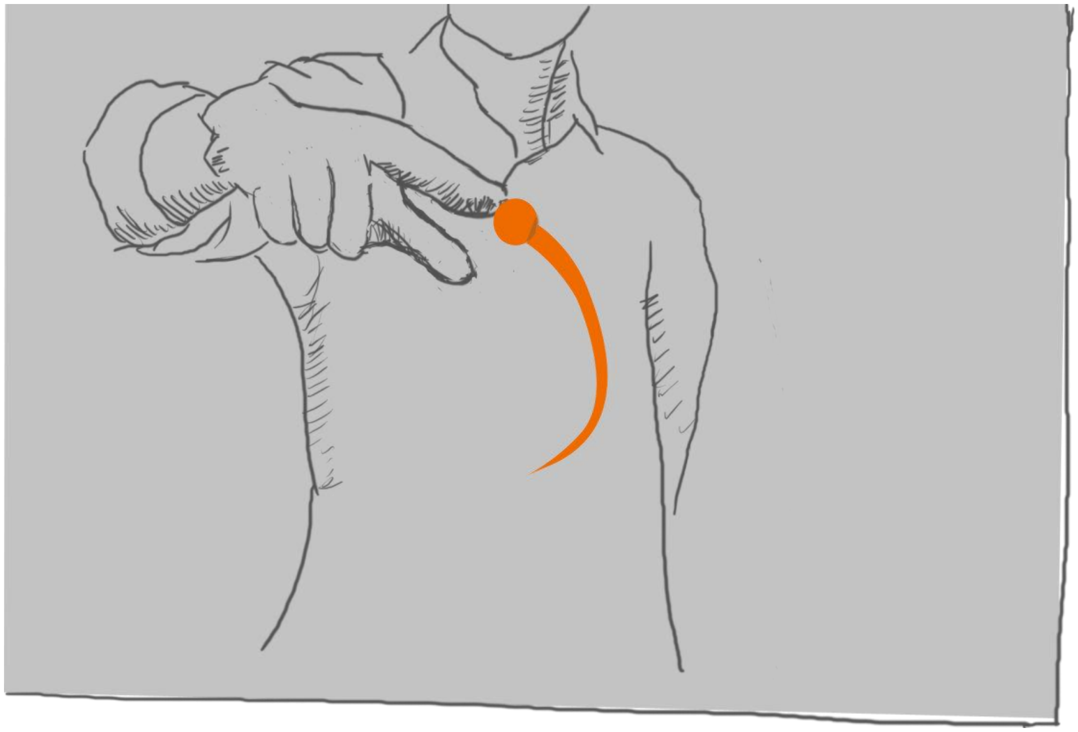

Figure 3.6: The person on the opposite side of the actor sees the finger movement augmented with an on-screen trace.

To mitigate these problems, we suggest augmenting a person's actions with literal on-screen representations of those actions. Examples to be discussed in our own system include highlighting a person's fingertips (to support touch selections), and generating graphical traces that follow their movements (to support simple hand gestures). The latter is illustrated in Figure 3.6.

3.3 The Implementation of FACINGBOARD-II: a Two-Sided Transparent Display

To our knowledge, no other transparent screen-based systems offer a full range of two-sided interactive capabilities, i.e. the ability to display different graphics on either side (but see Olwal et al., 2008; Lee et al., 2014). Consequently we implemented our own display wall, called FACINGBOARD-II. Because it used 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.

3.3.1 Projector and Display Wall Setup

Figure 3.7 and 3.8 (next page) illustrate our 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, which affords different graphics per side, and which minimizes occlusion and glare through the screen. Projections are reflected through the mirrors at a downwards angle onto both sides of the fabric. A separate computer-controlled each projector, and both run our distributed FACINGBOARD-2 software that coordinates

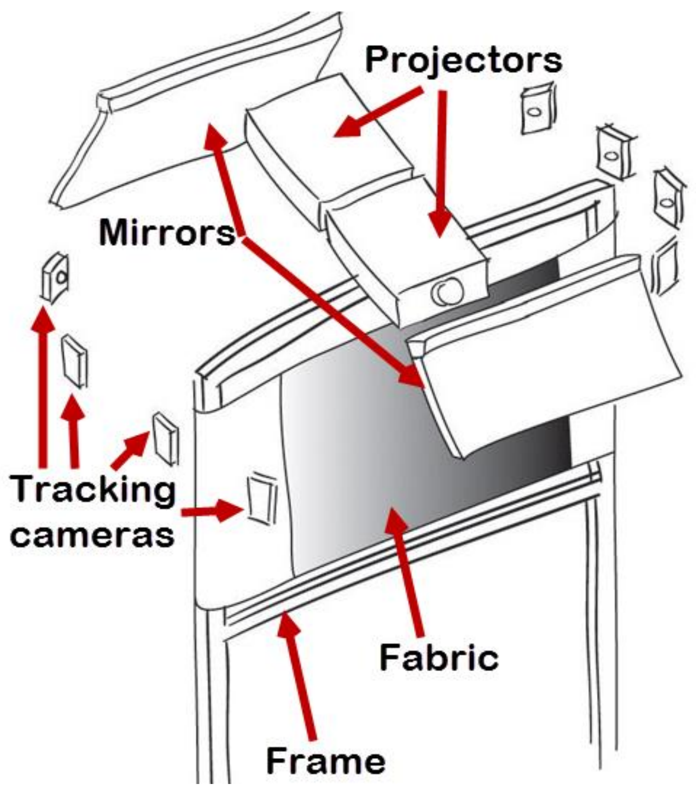

Figure 3.7: The setup of FACINGBOARD-II

Figure 3.8: The actual FACINGBOARD-II setup

what is being displayed. Lighting is also controlled. Room light is kept low to minimize glare, while directional lights illuminate the people on either side.

3.3.2 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 LCD / OLED screens inherently display on one side. The various glass surfaces and/or films used in projection systems would not work well for two-sided projection, as the projected content are designed with the goal of high-clarity bleed-through to the other side.

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

- images can be projected on both sides of the film, where the threads would reflect back and thus display the projected content;
- a person could see through the holes in the open weave to the other side;
- bleedthrough would be mitigated if the thread material were truly opaque;
- while large solid displays can attenuate acoustics to the point that either side requires microphones / speakers (Heo et al., 2013), sound travels easily through openly-woven fabric.

Figure 3.9 illustrates how this works in FACINGBOARD-II. First, it shows the open weave of the fabric (the inset shows a close-up of it). Second, it shows the graphics (the 'WallST' photo) projected onto this facing side opaque weave. Third, it shows the person on the other side as seen through the fabric's holes. Finally, it shows only minor bleed-through from the projection on the other side, visible as a slight greenish tint. This is caused

Figure 3.9: A close look of the projection fabric

by projected light from the other side bouncing off the horizontal thread surfaces, and because the fabric threads are not entirely opaque.

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

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

After testing various materials, we chose the blind fabric seen in Figure 3.8: tobacco thread color, and 10% openness (a factor provided by the manufacturer that purportedly represents the percentage of light penetration of blinds as determined by its thread width and hole size).

3.3.3 Input

Raw input is obtained from an off-the-shelf OptiTrack motion capture system. Eight motion capture cameras are positioned around the display (Figure 3.7). Participants on either side wear distinctive markers on their fingertip, whose positions are tracked by the cameras and captured as 3D coordinates. The FACINGBOARD-II 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 a 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 (Microsoft, 2014)), or 6 DOF input devices (e.g., Polhemus (Polhemus, 2014)). All have their own particular set of advantages and disadvantages (e.g., marker-based or markerless, high or low accuracy, ability to detect and track in-air gestures in front of but not touching the screen).

3.3.4 Limitations and Practicalities

Our FACINGBOARD-II setup works well as a prototyping platform, but still has a ways to go before it could be considered a commercially deployable product.

First—and common across all transparent displays—the degree of transparency is greatly affected by various factors as already described in section 3.2.3. The following two chapters will outline a quantitative study that looked into these factors and their effects.

Second, the fabric used to construct FACINGBOARD-II is not ideal. The threads are not particularly reflective, which means that the projected image is not of the brightness and quality one would expect of modern screens. As was seen in Figure 3.9, there is a very small amount of bleed-through of bright image portions to the other side. However, this is not noticeable if the other side also contains a brightly projected image. We believe better fabrics or screens could alleviate these limitations. One possibility is to paint a small grid or series of reflective opaque dots onto both sides of a thin transparent surface. Section 6.3.2 will discuss further options for potential materials.

Third, as typical with all projection systems, image occlusion can occur when a person interposes part of their body between the projector and the fabric. We minimize occlusion by using downward-angled mirrors (Figure 3.8).

3.4 Designing FACINGBOARD-II Relaxed WYSIWIS

Our test-bed application is illustrated in Figure 3.10: an interactive photo and text label manipulation system. It includes a public area (top central), a private area (bottom), and a personal palette (left), all of which will be discussed below. Because we have independent control of both input and output on either side, we are able to realize the various relaxed-WYSIWIS (what-you-see-is-what-I-see) features as described in our Design Rationale section.

Selective image and text reversal. As mentioned, graphics displayed on a 'one-sided' traditional transparent display will appear mirror-reversed on the other side. For example, Figure 3.10 shows one person's view of the correctly oriented images and text in the public area, while in Figure 3.11 it appears in reverse to the person on the other side. We overcome this problem by selectively flipping images and text in place (Figure 3.12). Each image and text block is precisely aligned to display at the exact same location on both sides, but its

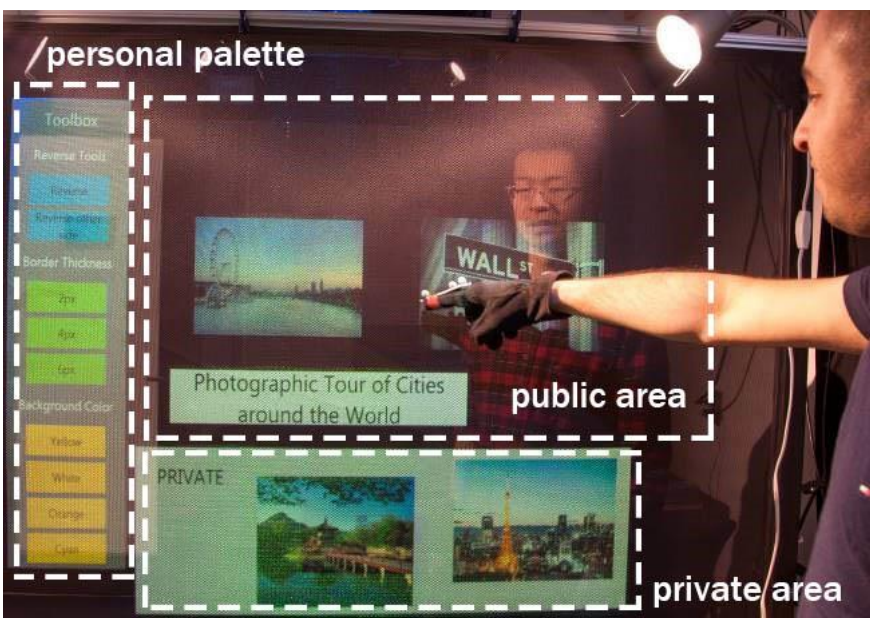

Figure 3.10: Our test-bed application with its public area, private area, and personal palette marked. The image also shows Person 1's view, where photos / text are correctly oriented

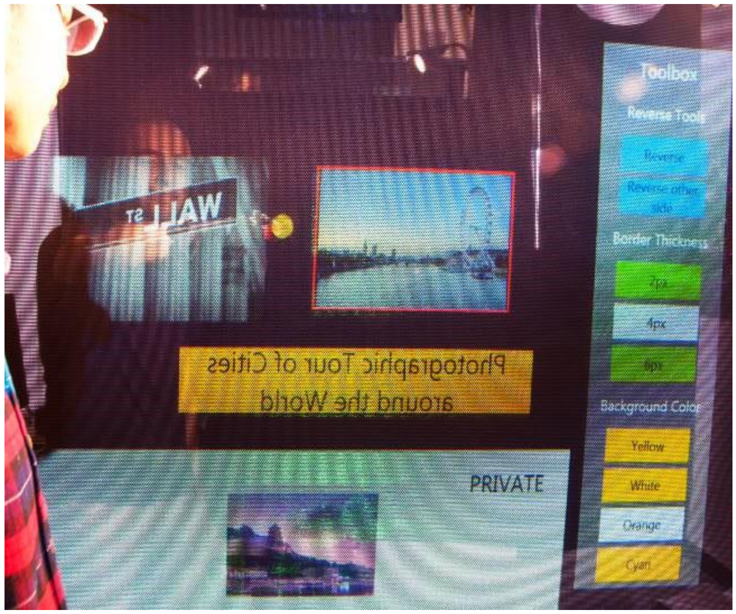

Figure 3.11: Person 2’s view on the other side, showing how photos and text would normally appear as reversed

content on one side are flipped to maintain the correct view orientation. Similarly, the text shown in the personal palette and private area are oriented to make them readable on either side.

Private work areas. While the public work area is visible to both people (albeit with flipped content), the content of the private area are distinct to the viewer. For example, Figure 3.10 shows how Person 1 has 2 photos in his private area, while 3.11, 3.12 shows

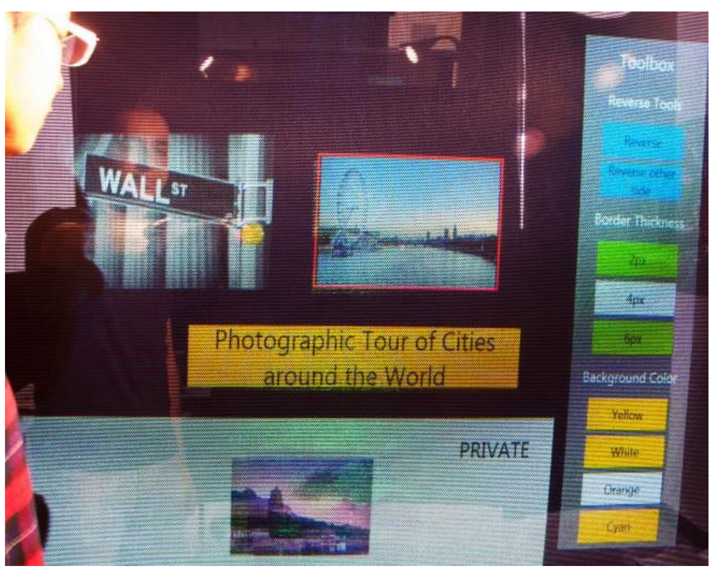

Figure 3.12: Person 2’s relaxed-WYSIWIS view; text/photos unreversed

how Person 2 has only 1 (different) photo. Each person can drag objects between their private and public areas, which causes them to disappear / reappear from the other person's view.

Semi-personal view of public objects. Each person is selectively able to modify the appearance of the text and images seen in the public view. Using the palette controls, they can reverse a selected object, add a red border to it, change the border thickness, as well as the background color of the text. These changes appear only on one side. For example, in Figure 3.11, Person 2 has reversed his image as he wishes to point to fine details of it: this makes its content identically aligned to what the other person sees. In Figure 3.11 and 3.12, he has added a red border to an image and has colored a text object in orange, which differs from what Person 1 sees in Figure 3.10.

Personal state. The palette controls, which are otherwise aligned on both sides, reflect their state on a personal basis, where selected radio buttons are shown in white. For example, we see in Figure 3.11, 3.12 that Person 2 has selected the '4px' border thickness and 'Orange' text block color, while in Figure 3.10 Person 1 has no options selected.

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

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

We enhance touch awareness by displaying a small visualization (a modest-sized dot) on the spot where the fingertip orthogonally projects onto the display. This dot only appears on the other side of the display, as it could otherwise mask the person's fine touch selections. For example, in Figure 3.10 Person 1 is touching a photo and no dot is visible

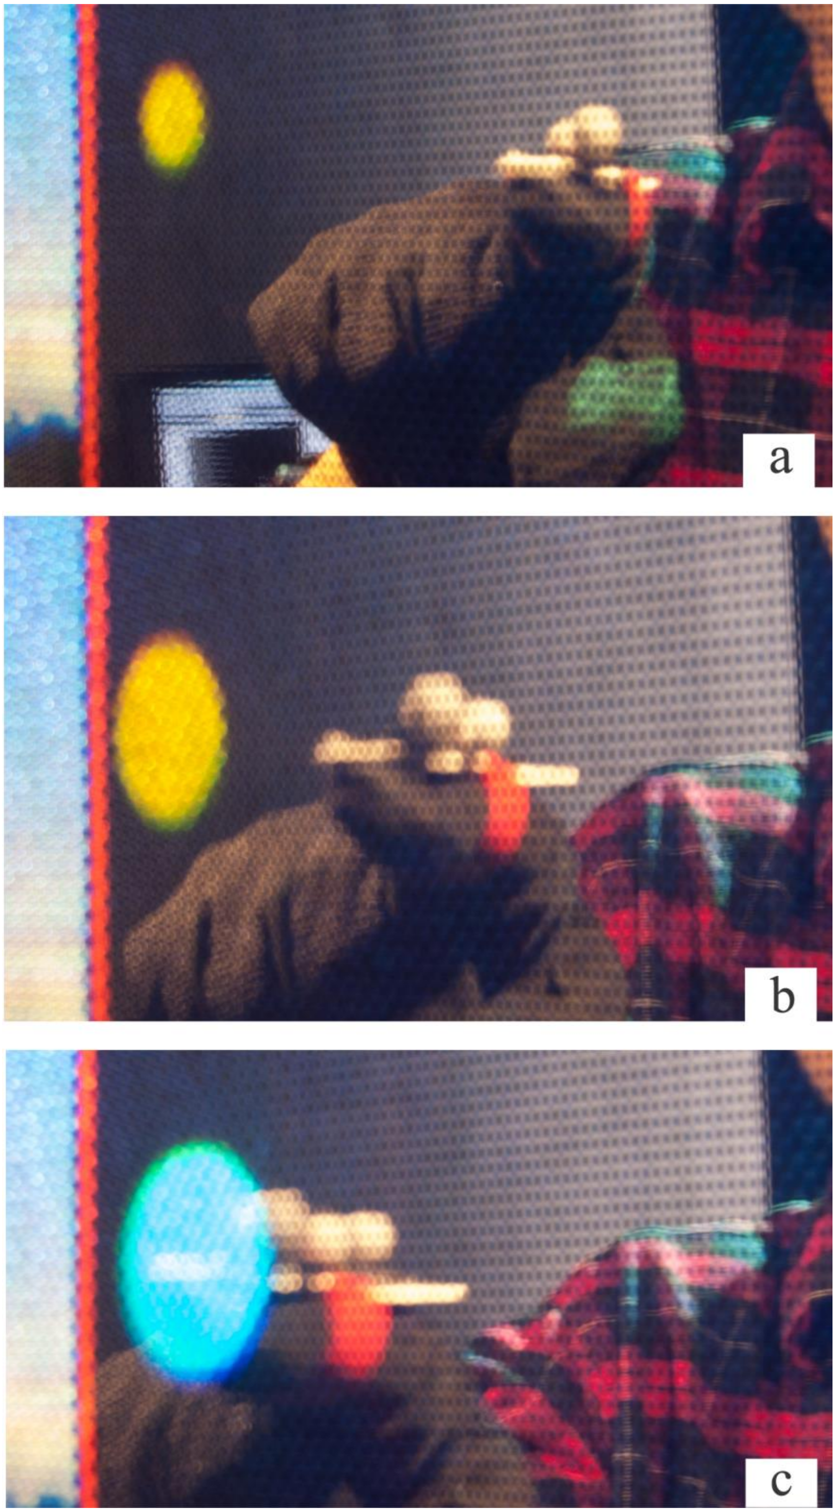

Figure 3.13: Using *touch augmentation* to enhance awareness of fingertip position (*a*) The dot is small, reflecting a distant finger. (*b*) The dot's size increases as the finger approaches. (*c*) The dot grows to its full size and changes color when the finger is making contact with the screen.

to him. However, Person 2 sees the dot on their side (Figure 3.11, 3.12). Figure 3.13a-c shows how the actual size of the dot varies as a function of the distance between the fingertip and the display, i.e., the dot is small when the finger is far from the surface (3.13a), gets increasingly larger as the finger moves towards the surface (3.13b) and is at its largest when touching the surface (3.13c). When a touch occurs, the dot's color also changes (3.13c).

We also use traces (Gutwin and Penner, 2002) to enhance gestural acts. As seen in Figure 3.14, an ephemeral trail follows a person's 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 named these two augmentation techniques *touch augmentation* (for the dot method) and *trace augmentation* (for the trace method), respectively. A study investigating their effectiveness will be outlined in the next chapter.

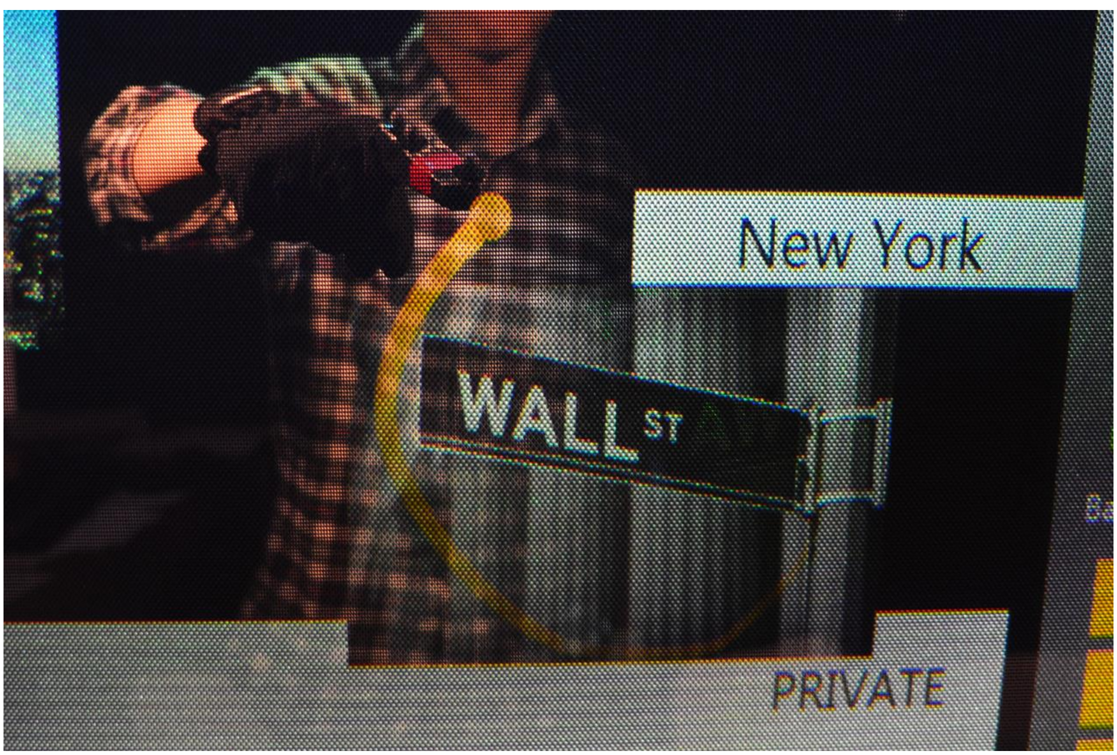

Figure 3.14: Augmenting gestures with *trace augmentation*

3.5 Summary

This chapter covered our design and prototyping efforts for collaborative transparent displays. It started with the implementation and informal evaluation of our first prototype, FACINGBOARD-I. Its purpose was to probe into the design space and inform our successive exploration. It was a one-sided LCD-based transparent display that supported simple cooperative applications. We then discussed three design requirements we determined to guide the design of collaborative transparent displays: *two-sided interactive input*, *different content on both sides*, and *augmenting human actions*. They were derived from synthesizing theories of workspace awareness and territoriality, and the lessons learned from FACINGBOARD-I. Finally, we described the setup and relaxed-what-I-see-is-what-you-see interface features of our second prototype, FACINGBOARD-II. It was a two-sided transparent display built to fulfill the suggested requirements.

FacingBoard-II embodied our design thinking so far on collaborative transparent displays. It also provided a platfom for further emperical study and evaluation. As introduced previously, the transparency of transparent displays is subject to various factors and it fundamentally influences usability. In the next chapter, we describe a study investigating the effect of changing transparency on workspace awareness in transparent displays, and the effecitiveness of our *touch* and *trace* visualization techniques for compensating potential awareness loss.

A video illustrating FacingBoard-II can be found at:

http://grouplab.cpsc.ucalgary.ca/grouplab/uploads/Publications/Publications/2014-TransparentDisplay.DIS.mp4

Chapter 4. Transparency *vs.* Awareness Enhancement Techniques

This chapter details a quantitative study we conducted that considered various factors that could affect collaboration on a two-sided transparent display. In particular the study primarily investigates how workspace awareness is affected by various level of display transparency, and whether particular action augmentation techniques can enhance awareness when transparency is compromised. The outcome of this study is highly relevant to the usability and practicality of not only the transparent display presented in the preceding chapters, but to collaborative transparent displays in general.

4.1 Independent Variables

A major premise behind the design of a collaborative transparent display is that transparency supports workspace awareness. Yet, as mentioned previously, transparency is not always guaranteed, where it can be affected by several factors (as discussed shortly). In the worst case, people on either side of the display will barely be able to see through it, which in turn means that they will have difficulty staying aware of the actions carried out by each other. This, of course, could defeat the purpose of such a display. We believe it critical for designers to understand these factors and the degree that they can affect collaboration. We also believe it is critical that designers consider workarounds that help augment awareness when transparency is compromised, which will also be discussed shortly.

Before delving into these factors, it should be noted that the awareness provided by a two-sided transparent display may not be symmetrical across both sides. Figure 4.1 shows an extreme example of this effect. In this instance, side 1 (Figure 4.1a) is highly transparent as only one image is displayed on that side. In contrast, side 2 (Figure 4.1b) is almost opaque as more images have been projected on that side. That is, transparency symmetry (and thus awareness symmetry) is not a given.

More generally, at issue is how awareness is compromised as a function of transparency, and whether we can mitigate these effects by augmenting the interface with

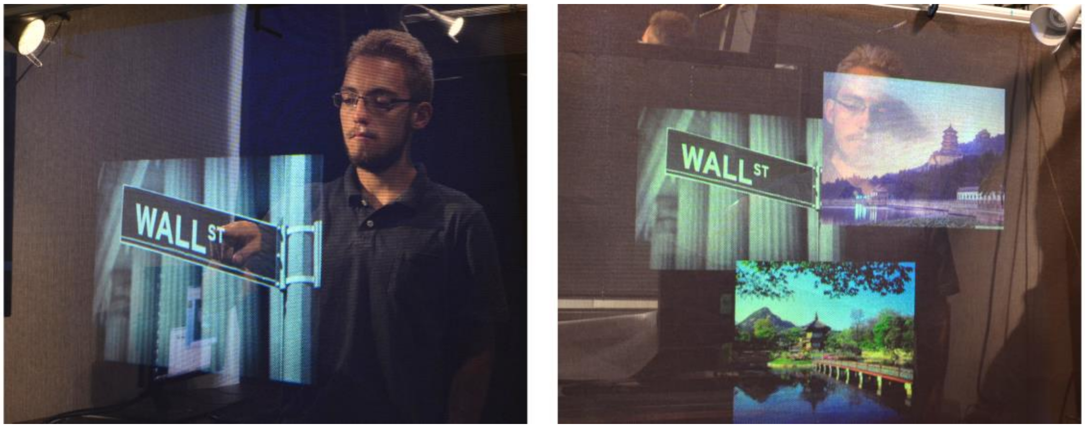

a) sparse graphics **b)** dense graphics

Figure 4.1: A person as seen through the display showing (a) sparse or (b) dense graphics

action visualization techniques that enhance what people can see. This study will measure how much the user on one side can monitor the activities of the other user on the other side of a collaborative transparent display, where we will consider two independent variables: various levels of transparency, and the presence or absence of augmentation techniques.

For terminology convenience, the '*viewer'* is the person who observes the activities (the '*actions'*) of the person (the '*actor'*) on the other side of the display.

4.1.1 The Transparency Factor: Screen Material, Graphics Density and Lighting

Our own observations of transparent displays along with our experiences creating the display described in prior chapters suggest three elements that may influence transparency and thus awareness.

1. The actual transparency of the screen material being used as part of the display.
2. The density of graphics being projected on the viewer's side of the display (which includes its brightness).
3. The lighting conditions on the actor.

Different materials have different transparency properties. Clear uncolored glass, for example, is usually considered fully-transparent. Yet manufactured screens often compromise full transparency to some extent, for example, by the sandwiching of emissive and conductive layers between glass plates in OLED displays. Our own idea of using fabric with large holes is also a trade-off: fabric with large holes is more transparent but less amenable to displaying graphics than fabric with smaller holes. We mentioned in previous chapters that we empirically compared various off-the-shelf fabrics, where we ultimately chose one that seemed to provide a reasonable trade-off between transparency and graphics display. Of course, other screen materials and technologies may exhibit different transparency properties. For example, the authors of the JANUS system mentioned in Section 2.1.2 describe their technology as being more transparent than ours (Lee et al., 2014).

To explain the other two properties, we saw that viewers found it increasingly difficult to see the actor's actions through the display when dense, bright graphics were projected onto it, and if the actor was poorly lit. Figure 4.2 (next page) illustrates this effect as a grid[2]. The top left image shows that a well-lit actor seen through a screen displaying sparse graphics is quite visible through the medium. In contrast, the bottom right image shows that a poorly lit actor seen through the same screen but displaying dense graphics is almost invisible. The top right images is a well-lit actor as seen through dense graphics, while the bottom left image is a poorly-lit actor as seen through sparse graphics. Comparing the images across each row and each column further illustrates that each factor by itself affects visibility.

In summary, we identify three properties affecting transparency. In the following studies, we keep the first property (fabric material) constant across all conditions. However, we vary *transparency* as a single independent variable. We will use four transparency levels, each created from a particular combination of: a) the density of graphics being projected on the viewer's side of the display, and b) different lighting conditions on the actor. We are particularly interested in how the different transparency levels affects the visibility of the actor and the actor's actions.

[2] Due to limitations of photographing our setup, the transparency is actually better than what is shown in Figure 4.2

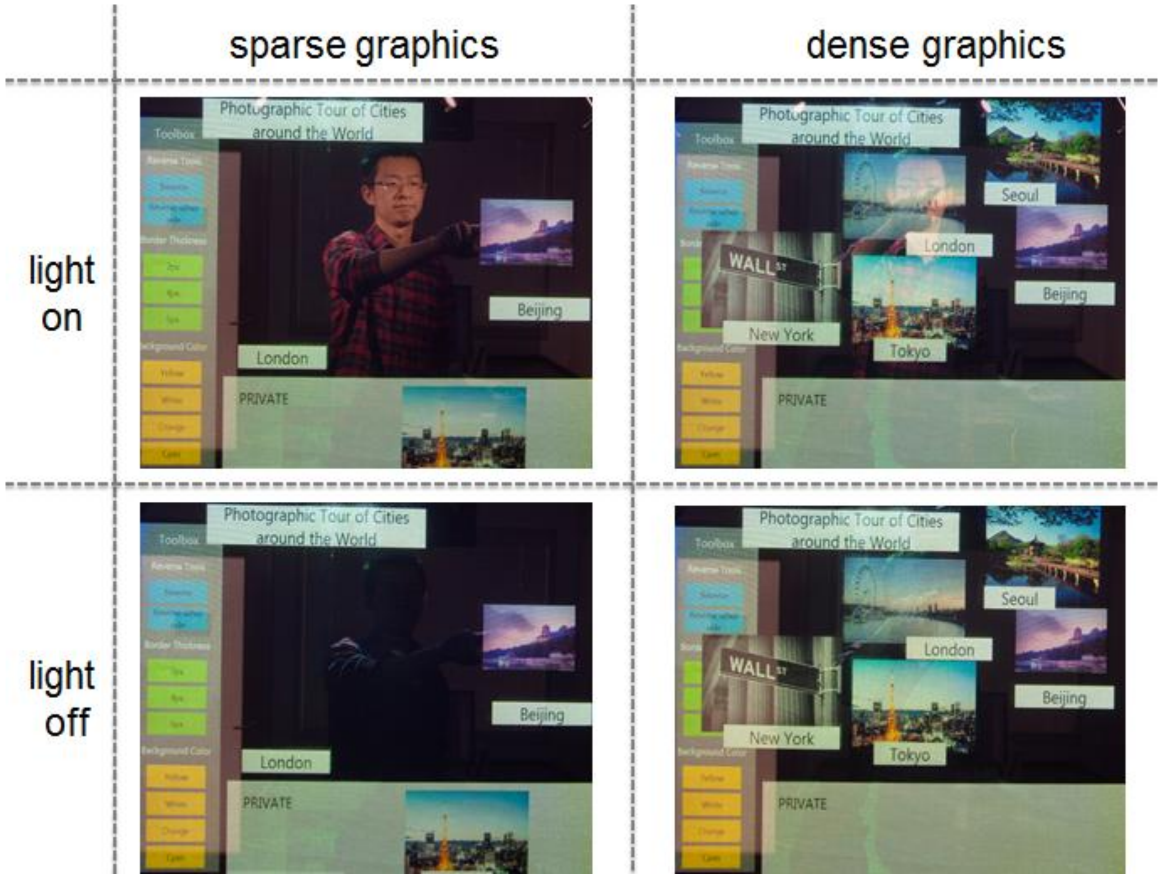

Figure 4.2: A person as seen through the display showing different graphical densities and lighting configurations

4.1.2 The Augmentation Factor: Enhancing Touch and Gestures

A key idea in a collaborative transparent display is that people can monitor each other's gestures, which in turn contributes to workspace awareness. As previously stated, gestures are critical. They are commonly used to communicate ideas. They mediate interactions in intentional communications. They contribute to consequential communication, where people observe others' actions to understand what they are doing, such as where they are focusing their attention and what actions they are performing (Tang, 1991).

When transparency is compromised, it may become difficult or impossible for the viewer to see the gestural actions of the actor, which in turn compromises workspace awareness. We previously described two action augmentation techniques in Chapter 3 that may mitigate this problem. First, the *touch* augmentation technique draws a dot on the screen location corresponding to the actor's finger. The dot becomes visually more intense as the actor's finger approaches the display, where the dot changes color when the display is actually touched. Second, the *trace* augmentation technique draws a fading line on the display, where the line follows the path of the actor's finger.

In the studies below, we treat *augmentation* as an independent variable, where it is either present or absent. As we will describe, the particular augmentation technique used (*touch* v.s. *trace*) will depend upon the particular task associated with each study.

4.2 Dependent Variables

The primary goal of each study is to determine if viewers can maintain workspace awareness under a variety of transparency conditions, and whether the use of augmentation across transparency conditions has a beneficial effect. To accomplish this goal, we developed several tasks. People's performance over these tasks are our dependent variables, where they serve as a reasonable measure of their ability to maintain workspace awareness.

4.2.1 Measuring Gesture Awareness: the shape and route task

Gutwin and Penner (2002) measured the effectiveness of using telepointer traces to improve gesture interpretation between remotely-separated participants as they collaborated over a shared workspace. Our gesture augmentation method is somewhat similar in spirit, as it also relies on traces (albeit of fingers rather than telepointers). Thus we developed our tasks and measures as variations of their methods.

Gutwin and Penner (2002) describe three types of gestural actions: *shapes*, *routes*, and *areas*.

- *Shapes* refer to finger movements that trace geometric shapes or symbols, which usually convey symbolic meanings. For instance, one draws a circle to indicate "a crowd moves in circular motion".
- *Routes* indicate paths that go through some objects in the workspace. Depending on the contexts, their meanings may be actual paths in the space, transitions between states, etc.
- *Areas* include gestures that outline a particular region or group of artifacts.

These gestures are only a subset of all communicative gestures. Still, they are indicative of many common gestures. To interpret these gestures, the viewer must correctly recognize gestural paths that are independent (shapes) and dependant (routes/areas) of workspace artefacts. Because routes and areas examine similar workspace-dependent gestures, we only use the shape and route gestural actions in our study.

We designed two tasks to see how well a person could recognize shape and route gestures: the *shape* task and the *route* task respectively. In both tasks, the experimenter as actor performed various gestural actions on his side of the display, while the participant as viewer was asked to interpret those gestures as seen through the viewer's side of the display.

Shape task. The actor used his finger to write, as a gesture, a horizontally reversed capital English letter just above the display surface (reversal made the letter correctly oriented on the viewer's side of the screen). The viewer's task was to say out loud what letter s/he saw. The answer was recorded as correct or incorrect, which is used in turn to calculate the *error rate* as the dependent variable. The error rate is calculated as the number of incorrectly recognized shapes (which included those events where the participant did not respond because they did not even notice the gesture) over the total number of shapes presented.

We note that this task also required the viewer to disambiguate those parts of the gesture that were not part of the letter (e.g., where the person's finger approached and left the display surface). For each trial, the actor chose a random letter from a 12-letter set (Appendix A.1), and chose one of four display quadrants to write that letter. Figure 4.3 on the next page outlines the task flow at particular transparency and gestural augmentation conditions. The top diagram shows the sequence of actions that the actor and the viewer each take. Subfigure A shows the actor writing the letter. Subfigure B and C show what a viewer may see with and without the *trace* augmentation, while subfigure D row shows the viewer saying what s/he has seen. Both subfigures B and C illustrates a condition with very dense graphics that significant compromises transparency.

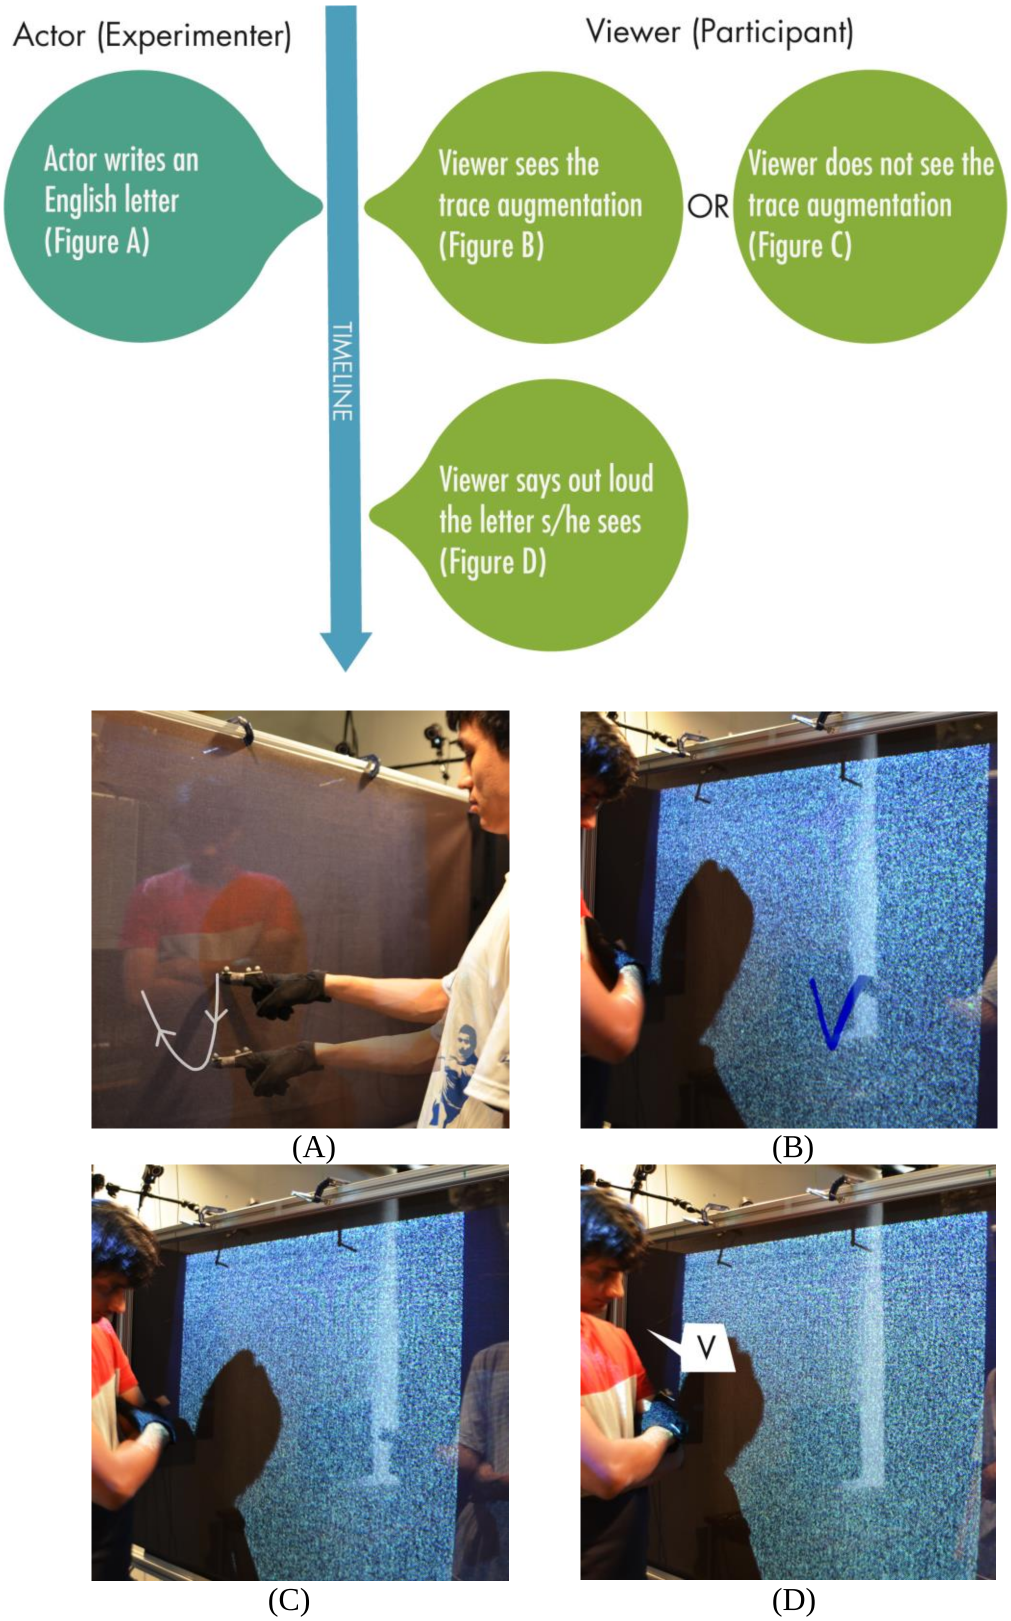

Figure 4.3: The flow of the *shape* task

Route task. A 16x10 grid of circles are aligned to appear on the same locations on both the actor's and viewer's sides of the screen, as illustrated in Figure 4.4. However, the grids' appearances differ: the circles are numbered on the actor's side (4.4a), and filled in red on the viewer's side (4.4b). The actor's task was to draw a path through a particular sequence of circles. The viewer's task was to reproduce the path by touching particular circles that the path went through. Results were logged as touch points corresponding to numbered circles, where they would be compared to the circles the route actually went through. This comparison was used to calculate the rate of correct responses by participants—the *accuracy* dependent variable—where they could exactly state which circle the gesture went through[3].

For each trial, the actor drew one of 8 different path shapes per trial as illustrated in Figure 4.5, where each path began at different circles. The actor referred to a cheat sheet that specified what route to draw where. Task difficulty across all paths were similar: all

a) actor's view b) viewer's view

Figure 4.4: The grid for the *route* task as displayed on either side

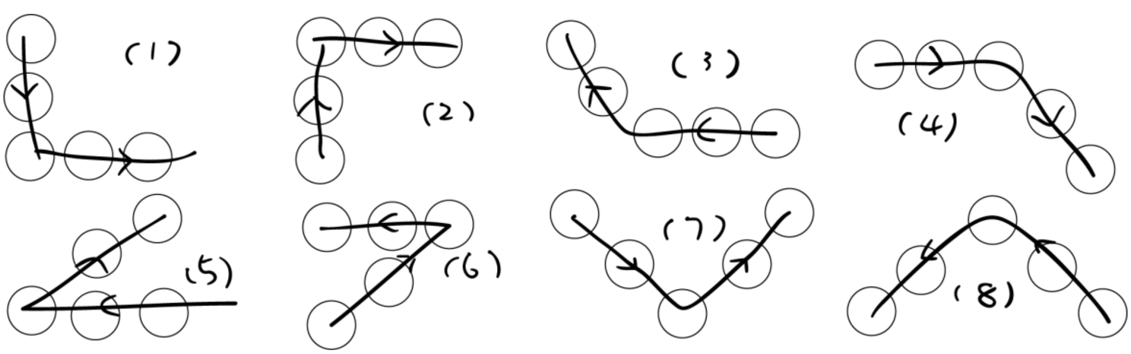

Figure 4.5: The gestures used in the *route* task

[3] For ease of analysis, we determine accuracy by a binary value: exactly right or incorrect. Our data, however, could allow for more subtle analysis, as participants sometimes notes some of the circles correctly, while still omitting or incorrectly identifying other circles as part of the route.

paths went through five circles, and had one turn in the middle. The viewer could indicate the path by touching the start, corner and end of it. Figure 4.6 (next page) outlines the task flow at particular transparency and gestural augmentation conditions. The top diagram shows the sequence of actions that the actor and the viewer each take. Subfigure A shows the actor drawing a path through the circle. Subfigure B and C show what a viewer may see with and without the trace augmentation, while subfigure D shows the viewer specifying the circles identifying that route.

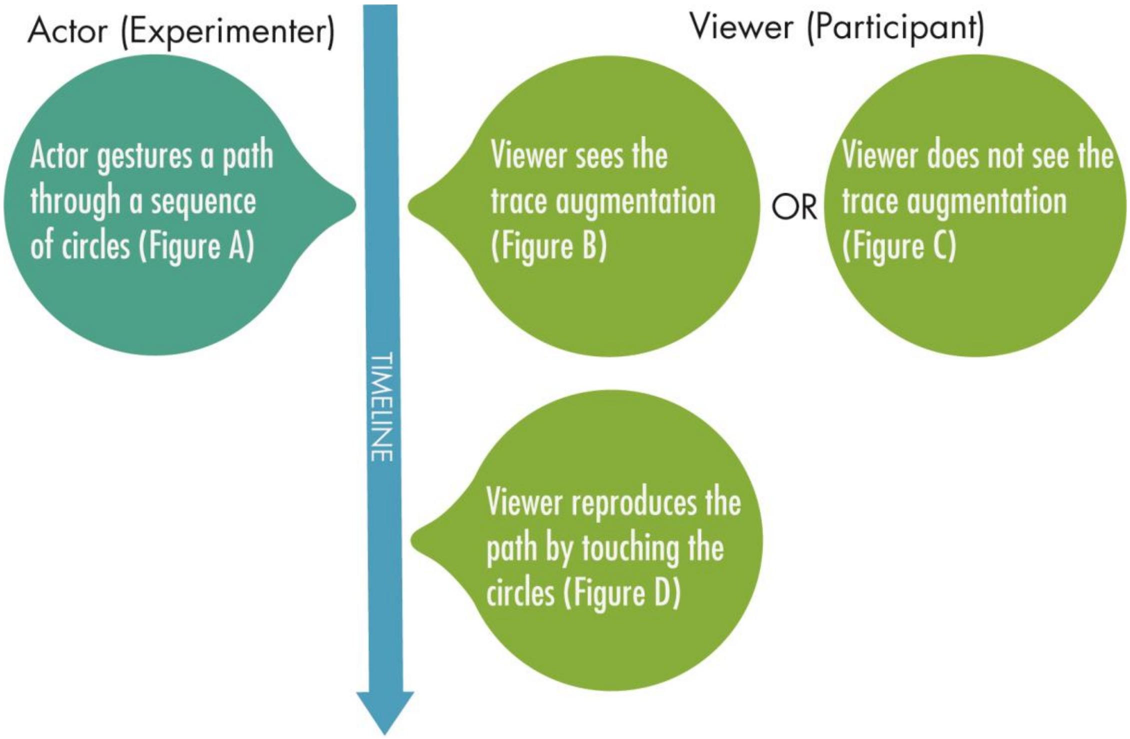

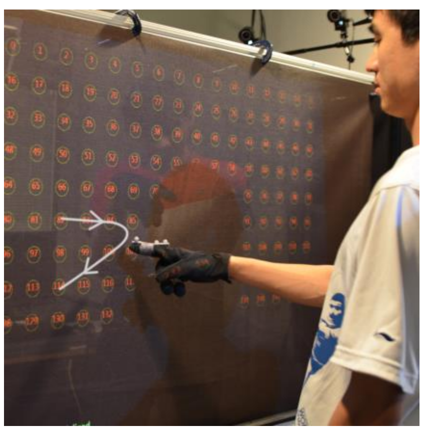
(A)

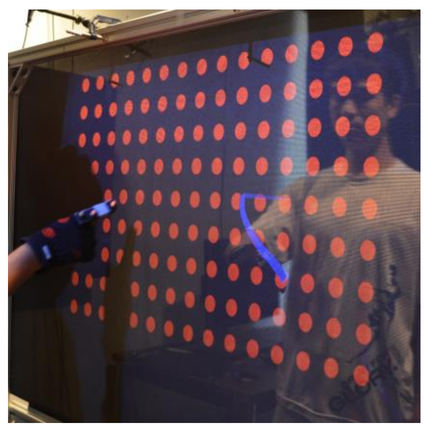
(B)

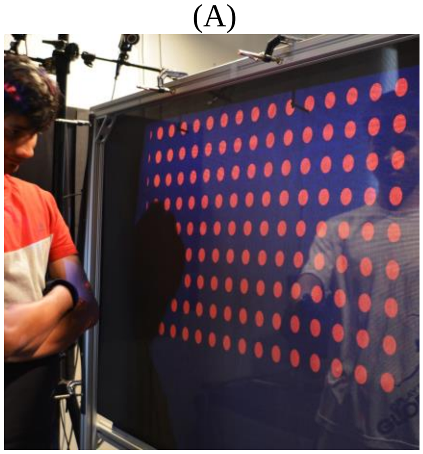
(C)

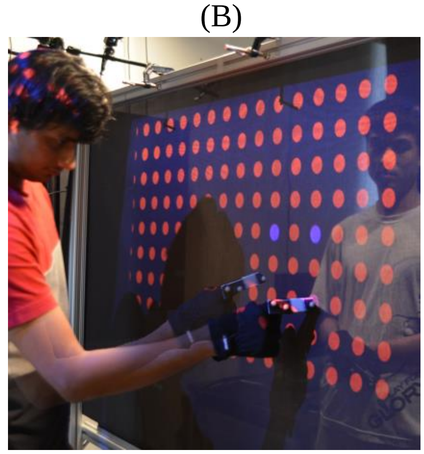
(D)

Figure 4.6: The flow of the *route* task

4.2.2 Measuring Touch Awareness: the *point* task

The *shape* and *route* tasks are examples of single-focus collaboration, where both actor and viewer are focusing their attention on the gesture as it is being performed. Yet many activities on a shared workspace are characterized as mixed-focus collaboration (see Section 2.2.1), where participants frequently switch between individual and group tasks and where attention is divided (Gutwin and Greenberg, 1998). Workspace awareness is particularly important for mediating the shift from individual to group work during mixed-focus collaboration, as seeing what others are doing helps (amongst other things) create opportunities to collaborate and helps coordinate mutual actions. When performing individual tasks, people still need to maintain an awareness of the whole workspace to coordinate their actions. Consider the importance being able to notice consequential communication via an actor's touches over a shared workspaces. Most contemporary interaction methods require the actor to touch the display in order to manipulate the workspace artefacts or change the workspace state. In mixed-focus collaborative activities, we expect a viewer pursuing individual work to either notice the touch actions of the actor in their peripheral vision, or occasionally glance around the workspace to see what the actor is doing.

Our third *point* task measures, in part, a viewer's ability to stay aware of the actor's touch actions during mixed focus collaboration. The viewer, while perform his/her individual work, had to indicate when s/he saw the viewer touch the work surface.

The actor's task was to tap randomly-positioned circles that only appeared on the actor's side of the display. After the circle was touched, it would disappear and a new one would appear shortly afterwards elsewhere. The viewer had two tasks. The *individual task* was similar to the actor's task, where the viewer was asked to tap solid squares as they appeared on the viewer's side of the display as quickly and as accurately as possible. In the *follower task*, the viewer was asked to tap those spots that s/he had noticed were touched by the actor (for convenience, we call this the *following touch*). The viewer was told that the follower task took precedence over the individual task, where s/he had to react as quickly and as accurate as possible to indicate what where the actor had touched.

The presentation of circles and squares was designed to split the attention of the viewer, creating a mixed-focus scenario for touch awareness. The software initially generated a random presentation sequence of viewer squares and actor circles, with a ratio of three squares for every one circle presented (60 squares and 20 circles per session). That is, we could consider a viewer's touch episode as comprising four touches: three of a square and one following touch. The timing was irregular to make this somewhat unpredictable to the viewer. Figure 4.7 on the next page illustrates the flow of the *point* task. The top diagram shows the sequence of actions that the actor and the viewer each take. Subfigure A shows the viewer touching the square. Subfigure B shows the actor touching the circle. Subfigure C and D show what the viewer may see with and without the touch augmentation when the actor touches the circle. Subfigure E shows the viewer responding by touching where the actor has touched.

Three metrics were used to measure awareness. *Response time* is a dependent variable calculated by measuring the elapsed time between the touch from the actor and the following responding touch from the viewer. *Response error* is a dependent variable that measures the distance between the location touched by the actor and the location touched by the viewer. Finally, the *miss rate* is a dependent variable that measures the rate where participants failed to react to a touch by the actor, e.g., because the viewer did not notice the touch or because the viewer failed to see where the touch occurred. The lower the values of these three metrics, the better the touch awareness.

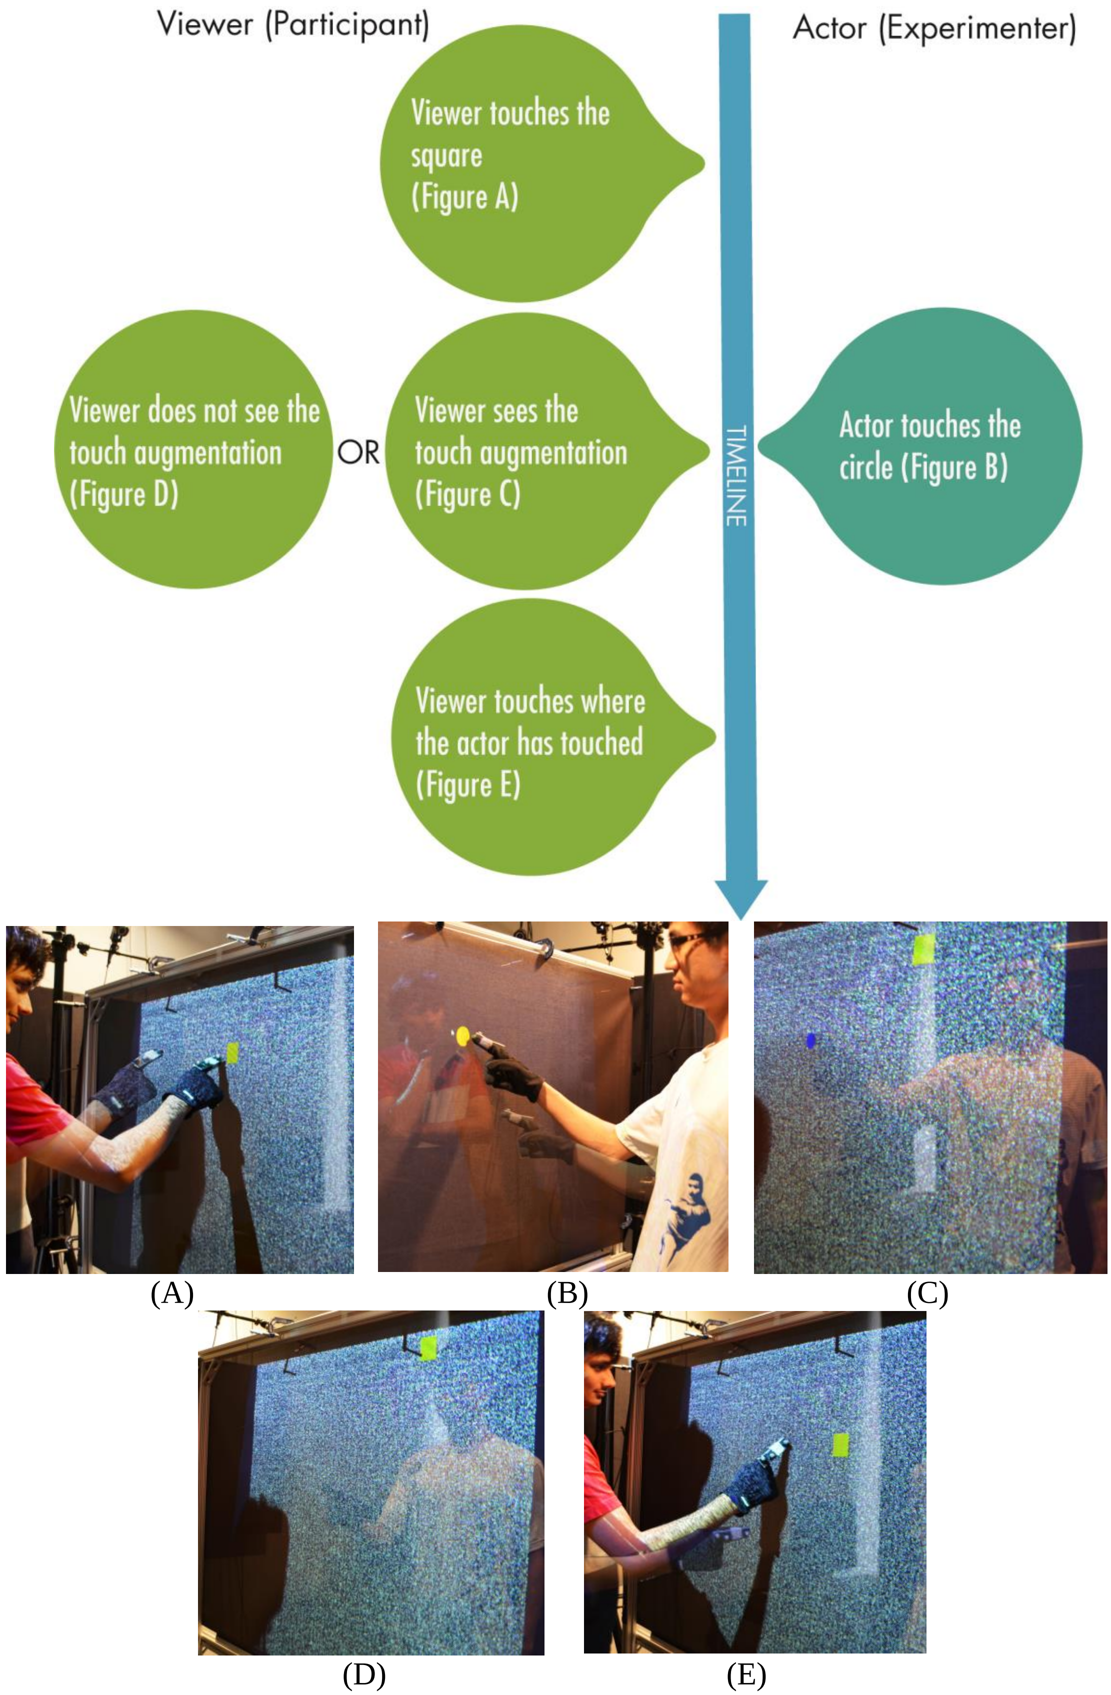

Figure 4.7: The flow of the *point* task

4.3 Study Design

We designed a series of three experimental studies, each corresponding to the tasks described above. All are based upon a within-subject (repeated measures) ANOVA factorial design. All used the same participants as viewers, where each participant did all three studies in a single session. Each study is somewhat similar, except that participants performed a different task (*shape, route* and *point*), each with their own dependent variables, which discriminated between studies.

4.3.1 Study factors

All studies included cross-combinations of transparency (4 levels), and augmentation (2 levels) as generally described above and detailed below.

Transparency is represented by four levels comprising a particular mix of graphics density and actor lighting. To explain, Figure 4.8 illustrates three different graphical density patterns that could be projected on the viewer's side of the display. 4.8a is all black, which affords the best transparency. 4.8b,c was generated by randomly drawing a given ratio of pixels white, and leaving the rest black, where the ratio of white to all pixels are 0.2 and 0.4. For convenience we call these *density low, density medium,* and *density high* for the black, 0.2, and 0.4 white pixels respectively. In general, the more white pixels, the more difficult it became to see through the display. Next, lighting indicates if the gloved hands of the actor was either illuminated or not illuminated by external lights. To control this lighting, we first isolated our study setup from external light sources using blackout curtains and other materials. We installed one 25W bulb as the constant ambient light

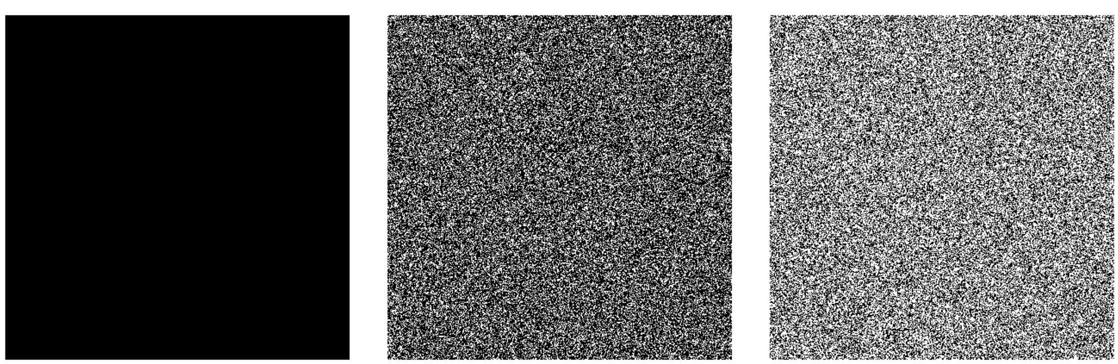

(a) (b) (c)

Figure 4.8: Three levels of graphics density: a) all black, b) 0.2 white, and c) 0.4 white

Figure 4.9: The same person under the lighting on (*left*), and the lighting off (*right*) conditions

source and two directional 25W incandescent lamps on the ceiling. Both directional lamps pointed downward in the region where the actor would be performing his hand gestures. We only used two levels of lighting, where the directional lamps were either switched on (*lighting on)* or off (*lighting off*). Figure 4.9 above shows the same person in both lighting levels behind a blank display. Our four transparency levels were created from a combination of density and lighting conditions as illustrated in Figure 4.10 on the next page. From most to least transparent (named transparency 1 – 4), these are low, medium, and high density with the light on, and high density with the light off.

Augmentation is represented by two levels, where it is either present or absent (named *augmentation on* and *augmentation off* respectively). The method used to augment gestures are as previously described, where—depending on the task—either the *touch* augmentation or *trace* augmentation are used as the augmentation technique.

4.3.2 ANOVA design

We combined the independent variables of transparency and augmentation to yield a 4 (*transparency*) x 2 (*augmentation)* design, or 8 different conditions per task. Because we had three separate studies (one study for each task), this resulted in 24 different conditions per subject. For each condition, subjects underwent many repeated trials. While this

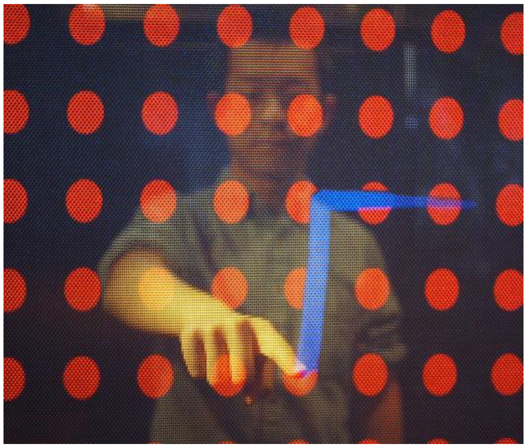

Transparency 1: density low, light on

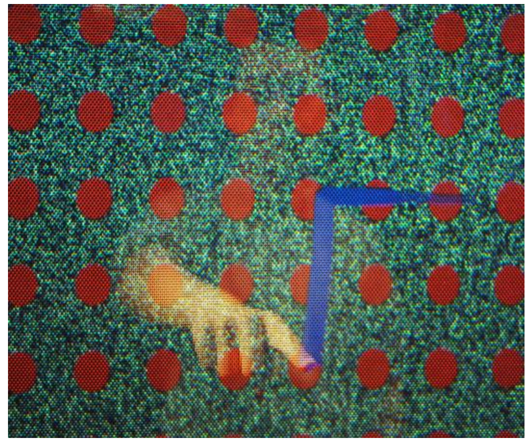

Transparency 2: density medium, light on

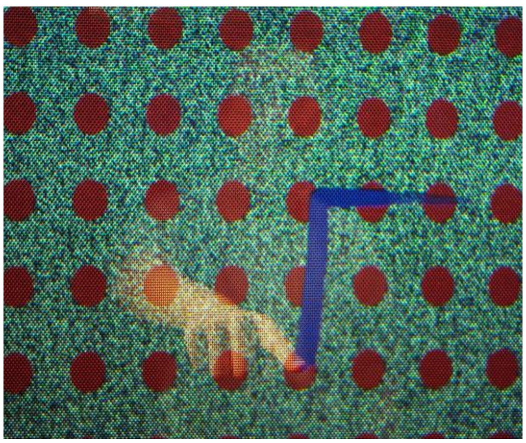

Transparency 3: density high, light on

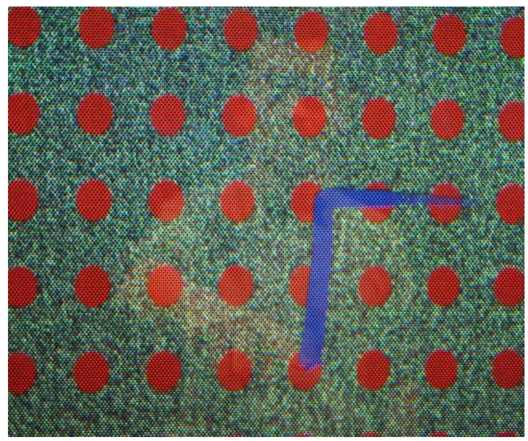

Transparency 4: density high, light off

Figure 4.10: The four transparency levels as a combination of density and lighting conditions. Note that the photos do not accurately portray what was actually visible through the screen due to difficulties taking photos in low-light conditions, and the print quality of these photos.

resulted in a considerable number of trials, trials were quick to do. Thus participants could complete the entire study within a single 90 minute session.

Our rationale is that the cross-combination of *transparency* x *augmentation* provide a reasonable spread from best case transparency (lighting on/ density low) to worst case (lighting off/ density high), and how augmentation can mitigate negative effects of decreasing transparency. Because provision of lighting is easy to do, we could reasonably expect that an installation would do so, but we wanted to include the case where lighting was insufficient as a point of comparison. In contrast, it is very difficult to mitigate graphics density, as doing so would affect what would be projected on the display. Thus understanding the effects of various ranges of graphical density under reasonable lighting conditions is important.

4.3.3 Hypotheses

The null hypothesis is inherent in the experimental design, and can be summarized as follows.

Null hypotheses: There is no difference in participant's ability to:

- recognize the shape (for the *shape* task) as measured by the error rate,
- to trace a route (for the *route* task) as measured by the accuracy, and
- to observe touches (for the *point* task) as measured by the response time, the response error, and the miss rate,

across the four transparency levels (constructed from a mix of graphics density and actor lighting) and the presence or absence of augmentation.

Our main predictions are as follows:

- Decreasing transparency reduces awareness and thus performance across all measures.
- Augmentation techniques can mitigate awareness loss (and thus increase performance) when awareness is compromised, but have little or no effect when awareness is not compromised.
- When performances across transparency levels are compared in *augmentation off* conditions,
  - denser bright graphics on the viewer's side (with lighting) progressively reduces awareness and thus performance.
  - dense graphics with no lighting reduces performance even further.

4.3.4 Materials

The study was conducted on the two-sided transparent display prototype, whose technical details have been described in Chapter 3. The two-sided transparent display had an effective display area of 57 cm by 36 cm. Two NEC GT2150 projectors were used to project visual contents on either side. Finger tracking was detected by 10 OptiTrack Flex

13 motion capture cameras, which tracked markers placed on the index finger of a glove that was worn by participants. All software was implemented in C++ and C#/WPF. Software modules displayed screen contents, tracked finger movements, and collected data about user actions. Two computers, each responsible for a projector, communicated via a 10Mb Ethernet to synchronize their activities. One computer was also responsible for tracking: a Xeon 2.4 GHz processor and 4GB memory. The other computer used a Core 2 2GHz processor and 2GB memory.

4.3.5 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. While all participants were experienced in touch screens interactions (e.g., daily use), only 8 of them had experience interacting with large interactive displays. All were right-handed. Each participant received a $15 payment for their contribution.

4.3.6 Procedure

After being briefed about the study purpose, the participant filled in a pre-study questionnaire that collected demographic information. The participant then performed the various tasks in sequence, beginning with the *shape* task, then the *route* task, and finally the *point* task. For each task, the experimenter instructed participants on what they had to do, after which participants did one practice block (see below for the number of trial per block). 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.

For each task, participants went through eight blocks, where each block corresponded to one of the eight conditions mentioned above in Section 4.3.2. The number of trials per block varied per task. For *shape* and *route* task, each block comprised 8 trials. For the *point* task, each black comprised 80 trials (60 individual task trials and 20 follower task trials; only follower task performance was measured). The *trace* augmentation was used for the *shape* and *route* task, while the *touch* augmentation was used for the *point* task. The presentation order of these two conditions was counterbalanced across the participants. The ordering of conditions per task was counter-balanced.

Throughout the study tasks, the experimenter and the participant stood facing each other on the two sides of the display prototype. The experimenter was responsible for switching the software between conditions (a brief computer interaction), and for presenting the gesture to the participant (using a prescribed protocol, see Appendix A.4). In the *shape* task, the experimenter also recorded participant's responses.

4.4 Summary

This chapter documented the studies we conducted to understand the influence of reduced display transparency on workspace awareness provided by collaborative transparent displays, and the efficacy of visually augmenting user actions to neutralize such influence. We had participants perform three tasks—*shape*, *route*, and *point*—over four degrading display transparency levels, with augmentation techniques (the *trace* augmentation for *shape* and *route*; the *touch* augmentation for *point*) absent or present. In the *shape* task, the viewer recognized the letter the actor wrote. In the *route* task, the viewer reproduced the path the actor drew. In the *point* task, the viewer monitored and responded to the actor's touch actions while performing a separate independent task. We measured participants' performance in each task as indicators of their awareness of touch and gestural actions. A video illustrating our study design can be found at:

http://grouplab.cpsc.ucalgary.ca/grouplab/uploads/Publications/Publications/2014-TransparentStudy.Report2014-1065-16.mp4

The next chapter detail the study results and our analysis.

Chapter 5. Study Results and Discussion

Chapter 4 outlined the three studies we conducted to investigate the effects of transparency and augmentation on a viewer's awareness of an actor's actions, as captured by various measures. We detailed the purpose of the study, the independent and dependent variables, the particular tasks, the experimental design, and the methodology. In this chapter, we will present our results on a study by study basis, where we will see that most null hypothesis are rejected and our predictions confirmed. This will be followed by our interpretation of those results, where we highlight the implications of these findings for future collaborative transparent display designers. We will close by discussing the limitations of both the studies and their findings.

5.1 Statistical Analysis Method

To summarize, our study involved three separate tasks (*shape*, *route*, *point*). Each task was performed using a cross-combination of the same conditions: 4 (transparency) x 2 (augmentation) as explained in Chapter 4 and illustrated in Table 5.1. The actual measures depended on the particular task.

For our analysis[4], we ran a two-way repeated measures analysis of variance (ANOVA) for each of the measures obtained from the three tasks (see Table 5.1), with sphericity assumed. For sphericity-violated cases, we used Greenhouse-Geisser corrections. For the post-hoc tests, we

			Augmentation	
			1-off	2-on
Transparency	1	Density low, lighting on		
	2	Density medium, lighting on		
	3	Density high, lighting on		
	4	Density high, lighting off		

Table 5.1: The 4 (transparency) x 2 (augmentation) experimental design common to all tasks.

[4] Because our statistical expertise is limited, we consulted with a statistician to verify our analytic approach.

used the test of simple main effects with Bonferroni corrections. All statistical calculations were done using the SPSS package. The level of significance was set a-priori at $p<0.05$.

5.2 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 (see Section 4.2.1). The experimenter (actor) then recorded if the viewer's response was correct or incorrect. The *error rate* of the *shape* task was then calculated as the ratio of misrecognized letters in each condition for each participant. Our null hypothesis is that there is no difference in participants' error rate in the *shape* task, regardless of the transparency and/or augmentation condition.

5.2.1 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 ($F_{3, 69} = 14.73$, $p < 0.05$). Given the significant effect of interaction, we conducted a post-hoc test of simple main effects on both independent variables. Figure 5.1 graphically illustrates the means of the error rate plotted by condition. The marginal means for each condition are reported in Table 5.2, with the results of the post-hoc test reported in Table 5.3a+b.

5.2.2 Discussion

The null hypothesis is rejected. As seen in the tables and as illustrated by the graph, we interpret the results as follows.

First, without augmentation, there is a notable increase in the error rate as display transparency decreases (see the blue line in Figure 5.1). Most of the pairwise differences between these means are statistically significant (Table 5.3b, top 6 rows). Differences are practically significant as well, where the error rate of

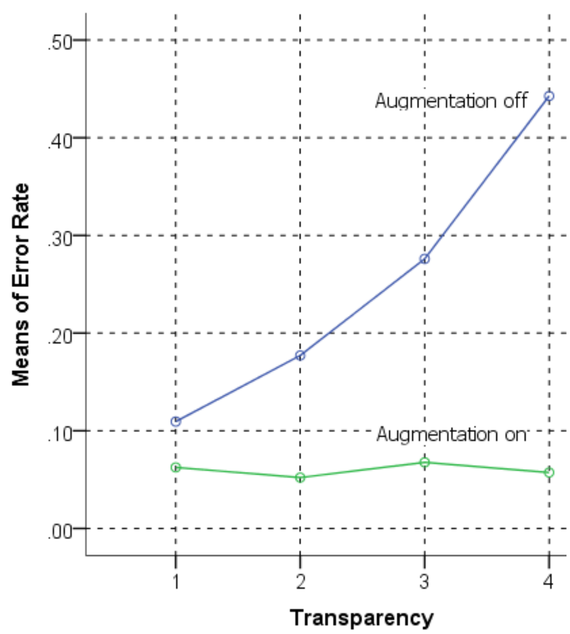

Figure 5.1: Means of error rate for the *shape* task, plotted by condition

~6% in the most transparent condition increases to ~44% in the least transparent condition (see the data points in the blue line in Figure 5.1, and Table 5.2 augmentation-off column).

Second, with augmentation, the error rate is constant regardless of the transparency level. That is, there is no significant difference in the error rate across any of the transparency levels when augmentation is used (green line in Figure 5.1; see Table 5.3b, bottom 6 rows). Notably, the error rate is low (at ~6%). This sharply contrasts with conditions without augmentation, where we saw the error rate increase as transparency is compromised (see Table 5.3a, rows 2-4).

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 (compare the first data points in the green and blue lines in Figure 5.1; see Table 5.3a, top row and Table 5.2 augmentation-on column). That is, there is no significant difference between the error rate.

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

Shape Task: Error Rate

			Augmentation	
			1-off	2-on
Transparency	1	Density low, lighting on	0.109 (0.027)	0.063 (0.020)
	2	Density medium lighting on	0.177 (0.040)	0.052 (0.018)
	3	Density high, lighting on	0.276 (0.034)	0.068 (0.018)
	4	Density high, lighting off	0.443 (0.056)	0.057 (0.020)

Table 5.2: *Shape* Task error rate: the marginal mean and standard error for each condition.

		Comparison between augmentation levels	Absolute value of mean difference	Std. error	Sig.
Transparency	1	ON vs. OFF	0.047	0.032	0.153
	2	ON vs. OFF	0.125	0.040	0.005*
	3	ON vs. OFF	0.208	0.037	0.000*
	4	ON vs. OFF	0.385	0.062	0.000*

a) Comparisons between *augmentation* levels at individual *transparency* levels.

		Comparison between transparency levels	Absolute value of mean difference	Std. error	Sig.
Augmentation	Off	1 vs. 2	0.068	0.046	0.947
	Off	1 vs. 3 *	0.167	0.046	0.008*
	Off	1 vs. 4 *	0.333	0.062	0.000*
	Off	2 vs. 3	0.099	0.042	0.162
	Off	2 vs. 4 *	0.266	0.051	0.000*
	Off	3 vs. 4 *	0.167	0.056	0.041*
	On	1 vs. 2	0.010	0.026	1.000
	On	1 vs. 3	0.385	0.062	1.000
	On	1 vs. 4	0.005	0.027	1.000
	On	2 vs. 3	0.016	0.019	1.000
	On	2 vs. 4	0.005	0.016	1.000
	On	3 vs. 4	0.010	0.018	1.000

b) Comparisons between *transparency* levels at individual *augmentation* levels.

Table 5.3: *Shape* Task post-hoc test on error rate data. '*' and peach-colored rows denotes a statistically significant difference at the p<.05 level.

5.3 The Route Task

In the *route* task, the actor drew a path through a particular sequence of circles shown on the display with his finger. The viewer's task was to reproduce the path by touching particular circles that the path went through. The circles the viewer selected were recorded and compared to those in the original path. For each participant, the *accuracy* of the route task was then calculated as the ratio of correctly reproduced paths to the total paths in each condition. Our null hypothesis is that there is no difference in participants' accuracy in the *route* task, regardless of the transparency and/or augmentation condition.

5.3.1 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 ($F_{3, 69} = 4.515$, $p < 0.05$). Given the significant effect of interaction, we conducted a post-hoc test of simple main effects on both independent variables. Figure 5.2 illustrates the means of the accuracy plotted by condition. The marginal means for each condition are reported in Table 5.4, with the results of the post-hoc test reported in Tables 5.5a+b.

5.3.2 Discussion

As seen in the tables and as illustrated by the graph, we interpret the results as follows.

First, without augmentation the accuracy decreases noticeably as display transparency deteriorates (the blue line in Figure 5.2). There are statistically significant differences between the accuracy of transparency level 1 and the other levels (see Table 5.5b top three rows). The differences are also practically significant: the ~91% accuracy in the most transparent condition degrades to ~62% in the least transparent condition (Table 5.4, augmentation-off column).

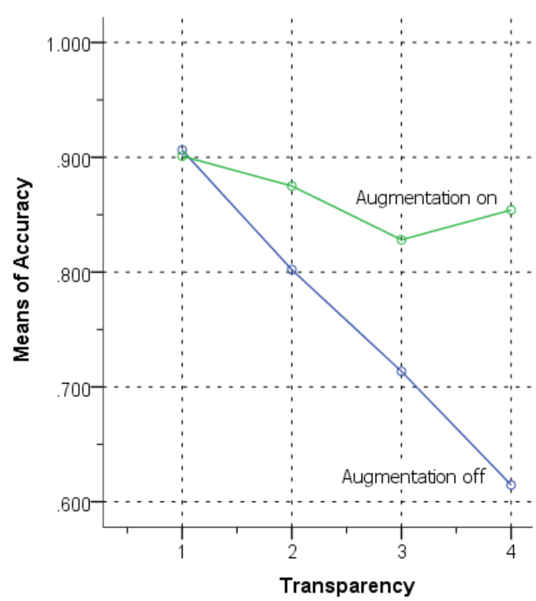

Figure 5.2: Means of accuracy rate for the *route* task, plotted by condition.

Second, the accuracy in *augmentation on* conditions is largely constant across the transparency levels. No significant difference is found between the means of the accuracy in these conditions (see Table 5.5b, bottom six rows) although there is a slight downward (albeit non-significant) trend as transparency declines (see the green line in Figure 5.2). In particular, in *augmentation on* conditions the accuracy is quite high (~87%, SD = 0.019). For transparency level 4, the accuracy with augmentation is significantly higher than the value without augmentation (see Table 5.5a, last row).

Third, in higher transparency conditions, there is no significant difference between the accuracy between augmented and un-augmented conditions. That is, the use of augmentation does not negatively affect accuracy when it otherwise may not be required (compare the first few data points in the green and blue lines in Figure 5.2; see Table 5.5a, first few rows).

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 method alleviates these difficulties. That is, the *trace* augmentation supports people's ability to perceive the other's path drawing gestures as transparency deteriorates.

Route Task: Accuracy Rate

			Augmentation	
			1-off	2-on
Transparency	1	Density low, lighting on	0.906 (0.044)	0.901 (0.023)
	2	Density medium lighting on	0.802 (0.047)	0.875 (0.033)
	3	Density high, lighting on	0.714 (0.061)	0.828 (0.038)
	4	Density high, lighting off	0.615 (0.052)	0.854 (0.026)

Table 5.4: *Route* Task accuracy rate: the marginal mean and standard error for each condition.

		Comparison between augmentation levels	Absolute value of mean difference	Std. error	Sig.
Transparency	1	ON vs. OFF	0.005	0.046	0.911
	2	ON vs. OFF	0.073	0.047	0.134
	3	ON vs. OFF	0.115	0.059	0.065
	4	ON vs. OFF	0.240	0.051	0.000*

a) Comparisons between *augmentation* levels at individual *transparency* levels.

		Comparison between transparency levels	Absolute value of mean difference	Std. error	Sig.
Augmentation	Off	1 vs. 2*	0.104	0.027	0.004*
	Off	1 vs. 3*	0.193	0.066	0.044*
	Off	1 vs. 4*	0.292	0.074	0.004*
	Off	2 vs. 3	0.089	0.066	1.000
	Off	2 vs. 4	0.188	0.067	0.063
	Off	3 vs. 4	0.099	0.063	0.793
	On	1 vs. 2	0.026	0.038	1.000
	On	1 vs. 3	0.073	0.043	0.628
	On	1 vs. 4	0.047	0.037	1.000
	On	2 vs. 3	0.047	0.040	1.000
	On	2 vs. 4	0.021	0.030	1.000
	On	3 vs. 4	0.026	0.044	1.000

b) Comparisons between *transparency* levels at individual *augmentation* levels.

Table 5.5: *Route* Task post-hoc test on accuracy rate data.
'*' and peach-colored rows denotes a statistically significant difference at the p<.05 level.

5.4 The Point Task

In the *point* task, the viewer was asked to: (a) carry out a separate 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. The *response time* was calculated as the average elapsed time between the touch from the actor and the responding touch from the viewer. The *response error* measured the distance between the location touched by the actor and the location touched by the viewer. Finally, the *miss rate* is the rate where participants failed to react to a touch by the actor.

Our null hypothesis is that there is no difference in participants' response time, response error, and miss rate in the *point* task, regardless of the transparency and/or augmentation condition.

5.4.1 Results: Response Time

Our analysis revealed a significant main effect for transparency ($F_{3, 69} = 20.731$, $p < 0.05$), augmentation ($F_{1, 23} = 4.517$, $p < .05$), and the interaction between them ($F_{3, 69} = 4.620$, $p < 0.05$). Given the significant effect of interaction, we conducted a post-hoc test of simple main effects on both independent variables. Figure 5.3 shows the means of the response time plotted by condition. The marginal means for each condition are reported in Table 5.6, with the results of the post-hoc test reported in Table 5.7a+b.

5.4.2 Discussion: Response Time

The null hypothesis is rejected. As seen by the Tables 5.6 and 5.7 and as illustrated by Figure 5.3, we interpret the results as follows.

First, without augmentation, response time tends to increase as display transparency decreases (see the blue line in Figure 5.3). There are statistically significant differences between means of the response time in

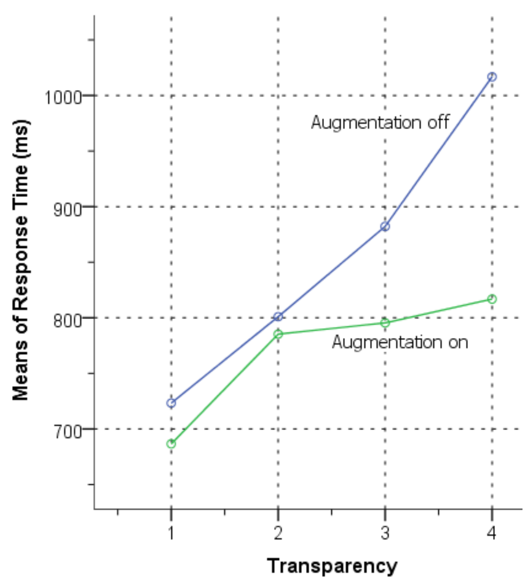

Figure 5.3: Means of response time for the *point* task, plotted by condition

transparency level 1 and level 3, level 1 and level 4, and level 2 and level 4 (see Table 5.7b top six rows). The differences are practically significant, with ~700 ms in the most transparent condition surging to ~1000ms in the least transparent condition (Table 5.6, augmentation-off column).

Second, with augmentation the response time exhibits a statistically significant increase from transparency level 1 (~700ms) to level 2 (~800ms), but did not increase afterwards (Table 5.6, augmentation-on column). That is, there was no significant difference is found between the response times of transparency level 2, level 3, and level 4 (see Table 5.7b, bottom six row). Notably, there is a statistically significant difference between *augmentation on* and *augmentation off* for transparency level 3 and level 4 (see Table 5.7a the last two rows).

Third, for the most transparent conditions (levels 1 and 2), adding augmentation neither increases nor reduces the response time (compare the first two data points in the green and blue lines in Figure 5.3; see Table 5.7a, first two rows). That is, there is no significant difference in the response time between both pairs of conditions.

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.

Point Task: Response Time

			Augmentation	
			1-off	2-on
Transparency	1	Density low, lighting on	723.259 (41.211)	686.613 (41.238)
	2	Density medium lighting on	801.011 (47.050)	785.333 (34.690)
	3	Density high, lighting on	882.205 (40.746)	795.408 (25.777)
	4	Density high, lighting off	1016.835 (68.608)	816.841 (28.716)

Table 5.6: *Point* Task response time (in milliseconds): the marginal mean and standard error for each condition.

		Comparison between augmentation levels	Absolute value of mean difference (ms)	Std. error (ms)	Sig.
Transparency	1	ON vs. OFF	36.646	35.576	0.314
	2	ON vs. OFF	15.677	49.434	0.754
	3	ON vs. OFF	86.797	41.584	0.048*
	4	ON vs. OFF	199.994	73.039	0.012*

a) Comparisons between *augmentation* levels at individual *transparency* levels.

		Comparison between transparency levels	Absolute value of mean difference (ms)	Std. error (ms)	Sig.
Augmentation	Off	1 vs. 2	77.751	34.043	0.192
	Off	1 vs. 3 *	158.946	38.052	0.002*
	Off	1 vs. 4*	293.576	60.377	0.000*
	Off	2 vs. 3	81.195	32.907	0.129
	Off	2 vs. 4*	215.825	51.573	0.002*
	Off	3 vs. 4	134.630	55.761	0.145
	On	1 vs. 2*	98.720	33.271	0.041*
	On	1 vs. 3*	108.795	34.486	0.027*
	On	1 vs. 4*	130.229	31.770	0.003*
	On	2 vs. 3	10.075	27.246	1.000
	On	2 vs. 4	31.508	22.715	1.000
	On	3 vs. 4	21.433	19.640	1.000

b) Comparisons between *transparency* levels at individual *augmentation* levels.

Table 5.7: *Point* Task post-hoc test on response time data.
'*' and peach-colored rows denotes a statistically significant difference at the p<.05 level.

5.4.3 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, $p < 0.05$). Given the significant effect of interaction, we conducted a post-hoc test of simple main effects on both independent variables. Figure 5.4 depicts the means of the response error plotted by condition. The marginal means for each condition are reported in Table 5.8, with the results of the post-hoc test reported in Table 5.9a+b.

5.4.4 Discussion: Response Error

The null hypothesis is rejected. As seen by the Tables 5.8 and 5.9, and as illustrated by the graph in Figure 5.4, we interpret the results as follows.

First, without augmentation the response error increases as display transparency deteriorates (see the blue line in Figure 5.4). Statistically significant differences in the response time are found between transparency level 1 and the other three levels, and between level 2 and level 3 (see Table 5.9b top six rows). The differences are practically significant, where the response error of ~28mm in the most transparent condition increases threefold to ~99mm in the least transparent condition (see Table 5.8, augmentation-off column).

Second, with augmentation the response time is constant regardless of the transparency levels. That is, no significant difference is found in the response error across any of the transparency levels in *augmentation on* conditions (see the green line in Figure 5.4 and Table 5.9b, bottom six rows). Furthermore, the response error stays low (at ~33mm) when augmentation is present; this contrasts dramatically to the statistically significant increase in response error without augmentation when display

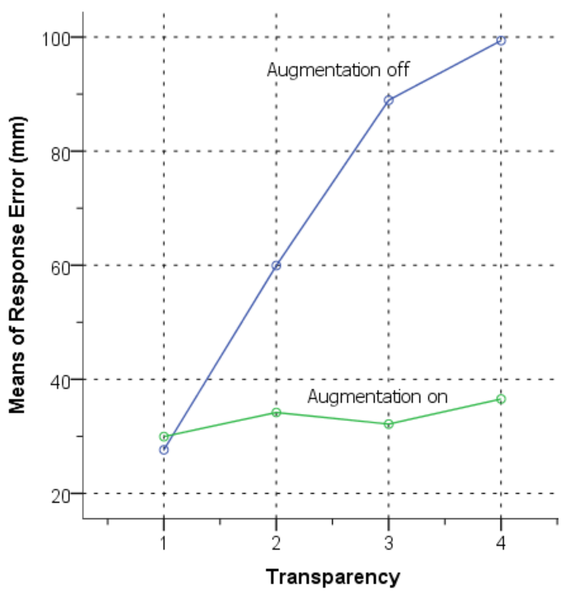

Figure 5.4: Means of error rate for the *point* task, plotted by condition.

transparency is compromised (see Table 5.9a).

Third, in highly transparent level 1 conditions, there is no significant difference in response error when augmentation is on or off (compare the first data points on the blue line and the green line of Figure 5.4; see Table 5.9a, the first row). That is, augmentation has no negative effect on the response error.

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, where they are equally precise across all transparency conditions.

Point Task: Response Error

		Augmentation	
		1-off	2-on
Transparency	1 Density low, lighting on	27.636 (2.053)	29.955 (3.759)
	2 Density medium lighting on	59.945 (9.286)	34.187 (4.988)
	3 Density high, lighting on	88.939 (9.972)	32.149 (3.932)
	4 Density high, lighting off	99.383 (13.067)	36.551 (5.234)

Table 5.8: *Point* Task response error (in mm): the marginal mean and standard error for each condition.

		Comparison between augmentation levels	Absolute value of mean difference (mm)	Std. error (mm)	Sig.
Transparency	1	ON vs. OFF	2.319	3.464	0.510
	2	ON vs. OFF	25.758	7.564	0.002*
	3	ON vs. OFF	56.789	9.930	0.000*
	4	ON vs. OFF	62.832	11.960	0.000*

a) Comparisons between *augmentation* levels at individual *transparency* levels.

		Comparison between transparency levels	Absolute value of mean difference (mm)	Std. error (mm)	Sig.
Augmentation	Off	1 vs. 2*	32.309	8.606	0.006*
	Off	1 vs. 3*	61.303	9.607	0.000*
	Off	1 vs. 4*	71.747	13.470	0.000*
	Off	2 vs. 3*	28.993	9.682	0.039*
	Off	2 vs. 4	39.438	15.381	0.104
	Off	3 vs. 4	10.444	13.914	1.000
	On	1 vs. 2	4.233	3.730	1.000
	On	1 vs. 3	2.195	4.301	1.000
	On	1 vs. 4	6.597	5.910	1.000
	On	2 vs. 3	2.038	3.627	1.000
	On	2 vs. 4	2.364	6.116	1.000
	On	3 vs. 4	4.402	5.401	1.000

b) Comparisons between *transparency* levels at individual *augmentation* levels.

Table 5.9: *Point* Task post-hoc test on response error data.
'*' and peach-colored rows denotes a statistically significant difference at the p<.05 level.

5.4.5 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 ($F_{3, 69} = 15.434$, $p < 0.05$). Given the significant effect of interaction, we conducted a post-hoc test of simple main effects on both independent variables. Figure 5.5 illustrates the means of the miss rate plotted by condition. The marginal means for each condition are reported in Table 5.10, with the results of the post-hoc test reported in Table 5.11a+b.

5.4.6 Discussion: Miss Rate

The null hypothesis is rejected. As seen in Tables 5.10 and 5.11 and as illustrated by the graph in Figure 5.5, we interpret the results as follows.

First, without augmentation the miss rate increases sharply as transparency is reduced (the blue line in Figure 5.5). There are statistically significant differences in the miss rate between transparency level 4 (the least transparent condition) and the other three levels (see Table 5.11b first six rows). The differences are practically significant, where the miss rate jumps from ~6% in the most transparent condition to ~43% in the least transparent condition (Table 5.10 augmentation off column).

Second, with augmentation the miss rate remained invariably low at ~8% (see the green line in Figure 5.5, and Table 5.10 augmentation-on column). No significant differences are observed between any pairwise comparisons of transparency levels (see Table 5.11b, bottom six row).

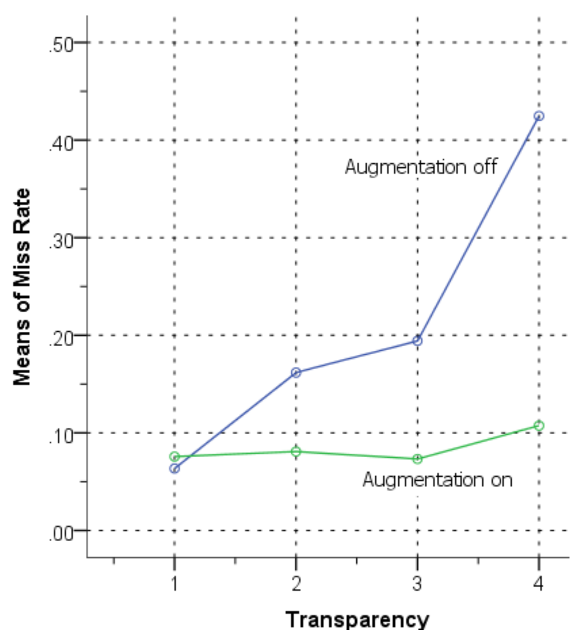

Figure 5.5: Means of miss rate for the *point* task, plotted by condition

Third, in the highly transparent level 1 condition, there is no significant difference between the augmentation on / off conditions. That is, the use of augmentation does not negatively affect the miss rate when augmentation may not be required (compare the first data points on the blue line and the green line in Figure 5.5; see Table 5.11a, the first row). However, when display transparency is compromised the miss rate in *augmentation on* conditions is significantly less than that in *augmentation off* conditions.

In summary, the results indicate that people 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: their miss rate stays low under all transparency conditions.

Point Task: Miss Rate

			Augmentation	
			1-off	2-on
Transparency	1	Density low, lighting on	0.064 (0.013)	0.076 (0.017)
	2	Density medium lighting on	0.162 (0.036)	0.081 (0.019)
	3	Density high, lighting on	0.194 (0.051)	0.073 (0.023)
	4	Density high, lighting off	0.425 (0.062)	0.107 (0.021)

Table 5.10: *Point* Task miss rate (in milliseconds): the marginal mean and standard error for each condition.

		Comparison between augmentation levels	Absolute value of mean difference	Std. error	Sig.
Transparency	1	ON vs. OFF	2.319	3.464	0.510
	2	ON vs. OFF	25.758	7.564	0.002*
	3	ON vs. OFF	56.789	9.930	0.000*
	4	ON vs. OFF	62.832	11.960	0.000*

a) Comparisons between *augmentation* levels at individual *transparency* levels.

		Comparison between transparency levels	Absolute value of mean difference	Std. error	Sig.
Augmentation	Off	1 vs. 2	0.098	0.037	0.084
	Off	1 vs. 3	0.131	0.055	0.162
	Off	1 vs. 4*	0.361	0.064	0.000*
	Off	2 vs. 3	0.033	0.030	1.000
	Off	2 vs. 4*	0.263	0.046	0.000*
	Off	3 vs. 4*	0.023	0.037	0.000*
	On	1 vs. 2	0.005	0.020	1.000
	On	1 vs. 3	0.002	0.018	1.000
	On	1 vs. 4	0.032	0.018	0.515
	On	2 vs. 3	0.008	0.020	1.000
	On	2 vs. 4	0.027	0.020	1.000
	On	3 vs. 4	0.034	0.023	0.894

b) Comparisons between *transparency* levels at individual *augmentation* levels.

Table 5.11: *Point* Task post-hoc test on miss rate data.
'*' and peach-colored rows denotes a statistically significant difference at the p<.05 level.

5.5 Discussion and Implication

Our discussion first focuses on the direct results of our study, and then somewhat more broadly on the implications of this work to collaborative transparent displays.

5.5.1 Overall discussion

The results presented above support our first prediction: decreasing display transparency reduces people's awareness of the other person's actions on the other side of a transparent display. 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.

The results also support the second prediction: 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 also saw 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 is clearly visible. Across all tasks and for four of the five measures, the presence or absence of augmentation had little effect on participants' performance on the highly transparent conditions. On the other hand, we also saw that augmentation usually 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 Figure 5.3 and the top two rows of Table 5.7a, there is no significant difference between augmentation on/off conditions across transparency level 1 and across level 2. Yet as the 7th row of Table 5.7b shows, 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 seem to help. Our post-study interviews of participants performing the point task 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 made contact with the screen—raising the arm, reaching out the hand, and the finger approaching—as a signal that a touch was soon to occur (i.e., that person's consequential communication). As display transparency decreased, those clues were less visible. They found it increasingly difficult to see the actor's pre-cursors to the touch action, and consequently they reacted more slowly. For example, in the case of transparency level 2, people found it more difficult to see those pre-cursor actions as they were further from the screen (such as arm movements), but they could still see the hand during its later close approach. 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. Thus for the transparency level 2 conditions, augmentation 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 (Figure 5.3. blue line). In those cases, augmentation helped signal that approach, thus enabling people to react faster as compared to no augmentation (Figure 5.3. 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.

Providing necessary workspace awareness is crucial for the utility and usability of collaborative transparent displays. Therefore, the hardware and software interface design for them should guarantee a reasonable capability to support workspace awareness. Knowing the effects of display transparency on the awareness provided, and the effectiveness of the augmentation methods, we suggest a few implications for addressing the awareness requirement.

5.5.2 Implication 1: Controlling Transparency

Most research and commercial works on transparent displays to date have based their interface designs on the assumption that the displays are truly transparent. They are often portrayed this way in advertisements, research figures, and even futuristic visions of technology. We suspect that the graphics density and lighting are tuned to show such displays at their best. But as our experiences point out, a transparent display is not invariantly transparent. The consequence (as our results clearly show) is that degrading transparency can greatly affect its ability to provide awareness information in collaborative situations. For example, we saw participants make ~3 times as many errors recognizing *shape* gestures between the least and most transparent condition (Table 5.2, first and last row). If designers of a collaborative transparent display want to support workspace awareness, they first need to recognize both the problem and possible solutions.

One partial solution is to control the display transparency as much as possible. Our experimental setup and study confirmed that high graphics density on the screen and dim lighting on the actor can reduce display transparency. A designer can control both factors to a reasonable level to ensure high transparency. For lighting, the system could incorporate illumination sources (perhaps integrated into the display frame) that light its users well. While we don't know what comprises 'good' illumination, we suspect this will be a combination of lighting position, intensity, and color. Graphics density will depend on the particular application and people's interaction with it. While this would be difficult to control for general purpose computing, it can be designed into custom applications. The custom application should distribute graphics sparsely on the screen, or have enough clear space between graphical elements to permit one person to see through those spaces. Other graphical attributes—colors, brightness and texture—can be adjusted to find a balance between seeing display graphics and seeing through those graphics.

Another partial solution considers other factors affecting display transparency. As previously introduced in Chapter 3 and 4, the ability of a person to see another person through the screen is affected by the combined influence of a number of environmental and personal factors. One environmental factor would control the ambient light in the room that may reflect off the display. Another environmental factor is the color of surrounding walls and furniture. For example, in our experimental installation, we surrounded the display with blackout curtains both to block out light and to provide a dark background color. A personal factor includes the color of the actor's clothes (bright colors are more reflective than dark colors) and how that color differs from the surrounding walls (contrast). It also includes special purpose clothes, such as reflective gloves that would better illuminate hand movements. For example, in our experimental installation we had participants wear gloves with reflective markers positioned on the finger trip (for the tracking system). We noticed that these markers made fingers much more visible through the display, and actually had to control for that in our experiment.

Another partial solution rests 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 can afford more light transmission, such as the JANUS system (Lee et al., 2014). However, we should not expect technical miracles, as we believe that all technologies will be affected by the other factors.

In principle, all the factors above can be controlled to find a good balance between seeing display graphics on the screen and the actor through the screen. In practice, however, some are easier to adjust and/or more impactful on display transparency. This, of course, is highly dependent on context. If designers do know the installation context, they may be able to devise (or recommend) specific transparency modulation mechanisms according to where the display is used and what task people are carrying out on it. For example, consider a large outdoor transparent display installation used as a public entertainment facility. Here, a designer may include artificial light sources that dynamically adapt to ambient light intensity, which can neutralize the change in display transparency due to variable weather and sunlight. In contrast, consider a display used in special purpose meeting room environment. Here, the wall colors and lighting may be tuned to the situation, and the color and brightness of the graphics can be carefully chosen to maximize display transparency.

Of course, there is a limit to how designers can control the installation. For example, overly strong lighting can disturb people. Reducing graphics density can affect both what application designers may want to do, and what users may want to see. We cannot expect users to change their clothes or gloves in normal situations. Demanding environmental changes is unreasonable for most situations. Fortunately, we can still enhance workspace awareness by augmenting user actions, as discussed next.

5.5.3 Implication 2: Augmenting User Actions

Our study results showed that augmentation techniques can mitigate awareness loss when display transparency is compromised. In spite of their simple forms, the *trace* and *touch* augmentation techniques proved effective for providing the awareness of the actions they represent across all tasks. We saw that, for most augmentation cases, people's performance remains relatively stable at high levels even when transparency degraded. What makes this particularly useful is that the information provided by the augmentation methods were somewhat crude: i.e., as simple as a moving dot or a winding trace of a single fingertip. Thus a partial and highly beneficial solution to the transparency problem is to provide a visual augmentation of a person's finger movements, which in turn signals pointing gestures, the focal point of input interaction (although this depends on the system), and hints at where the actor is directing their gaze.

However, we can do even better. While seeing finger movement is helpful, body language is far richer. We need to develop augmentation techniques that capture that richness. We previously stated how the *touch* augmentation could signal the occurrence and location of touch actions, but how people's response times were still slower in moderate transparency conditions because they could not see pre-cursor arm movements. In daily face-to-face activities, we maintain workspace awareness by observing movements of multiple body parts (including gaze awareness) and interpret those sequences o movements in relation to the workspace. In contrast, the *trace* and *touch* augmentations indicate only the movement of a single body part, the fingertip. While effective and efficient for particular tasks, we believe that more nuanced augmentation will be even more 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., 2006), or even that show the entire body (e.g., perhaps as a silhouette (Krueger et al., 1985)).

Of course, new techniques must be carefully crafted and tested. Technical challenges include tracking. Graphical challenges include finding an easily understood representation that does not occlude, distract, or otherwise interfere with a person's view of the workspace (as workspace awareness involves a view of the participant, the workspace artifacts, and the participant's actions relative to those artifacts).

One design consideration is to choose an augmentation technique that highlight 'core' aspects of a user's input or gestural actions essential to the task at hand. For example, highlighting touch via the *touch* technique may suffice for touch-screens, while the *trace* technique may be better for gesture-based input systems. The choice of augmentation may also consider the likely 'worst case' display transparency. Because awareness information is lost progressively as transparency deteriorates, the augmentation may need to supply only that critical information which is lost.

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

5.6 Limitations

Our study was (to our knowledge) the first of its kind. As typical with such controlled studies, it has several limitations as discussed next.

Study conditions. We used only four transparency levels, formed as a mix of different graphical densities and lighting conditions. While these were chosen to capture a range from what we considered highly transparent to barely transparent, it does not cover the full transparency spectrum. For example, we could have used brighter lighting on the actor to see its effect, or tested a broader combination of lighting on/off *vs.* graphical density. Our manipulation of graphical density was also artificial, where we used a random pixel pattern containing a well-defined ratio of bright *vs.* dark pixels to mask the display. Real world graphics are different. For example, we could have tested how people maintain awareness through (say) a document editor, a photo-viewing application, and/or a running video. Although we are confident about our conclusions, future works needs to investigate a broader spectrum of display transparency conditions.

Tasks. The three tasks of our study only examined a small set of various tracing gestures and touch actions that people perform during cooperative work. While we consider these tasks reasonable representatives of things that people do during collaboration, these tasks do not cover all interaction nuances. For example, the tasks necessarily require people to track a single finger movement, but not other bodily movements.

Our tasks were also artificial. They did not test people's awareness of others' actions on transparent displays with real tasks, where people may exhibit more complex interaction and gestural patterns. Alternately, people may change how they do their actions to compensate for loss of transparency, e.g., by shadowing their actions with verbal alouds ("I'll move the object at the top of the screen to the screen's bottom"). On the other hand, measuring workspace awareness in real tasks has long been a challenging research problem because of the complicated communicative and cognitive mechanisms involved in cooperative activities (see Hornecker et al, 2008; Morris et al., 2006). We will leave this for future work.

5.7 Summary

This chapter presented the results of the study described in Chapter 4 and our interpretation of the results. The study investigated the effect of display transparency on people's awareness of others' actions, and the effectiveness of action augmentation techniques. The statistical analysis validated our predications that people's awareness is reduced when display transparency is compromised, and that augmentation techniques can mitigate awareness loss. Based on our findings, we suggested a few implications for collaborative transparent display designers. The chapter closed with a few study limitations.

Chapter 6. Conclusion

This thesis presented our exploration of facilitating face-to-face collaboration with transparent displays. We motivated our research by arguing that transparent displays offer the unique collaborative benefit of supporting workspace awareness. Drawing on CSCW theories of workspace awareness and workspace territoriality, and reflecting on our first prototype, FACINGBOARD-I, we proposed three design requirements for collaborative transparent displays. In order to operationalize these requirements, we created a second prototype, FACINGBOARD-II, a two-sided transparent display affording touch input and independent, thus possibly different, graphical output on its two sides, along with a palette of supportive interaction techniques. We conducted empirical studies to understand how the capability of a two-sided transparent display to provide workspace awareness can be compromised by reduced display transparency, and the efficacy of action augmentation techniques to compensate for this potential awareness loss. We hope that our efforts will inform and inspire further research on supporting collaboration with transparent displays.

In this closing chapter, we revisit our research contributions, reflect upon the limitations of this work, and discuss possible paths for future research beyond those already presented.

6.1 Research Contributions Revisited

In Chapter 1 we outlined the five contributions of our work on interaction design for collaborative transparent displays. In the following section we revisit these contributions, clarify how our efforts led to them, and explain their significance.

Contribution #1, **proposing design guidelines for collaborative transparent displays**. Drawing upon workspace awareness theories (see Section 2.2), workspace territoriality theories (see Section 2.3), and the lessons we learned from prototyping FACINGBOARD-I (see Section 3.1), we set out three design requirements—*two-sided* *interactive input, different content on both side, and augmenting human actions*—as guidelines for the design of collaborative transparent displays (see Section 3.2). Our efforts to realize these design guidelines are reflected in all the other contributions described below. We believe that our design guidelines transcend the specifics of our prototype and can be applied to and inform the design of future collaborative transparent displays.

Contribution #2: **building arguably the first interactive two-sided transparent display prototype that can present different content on both sides**. In Section 3.3, we described the implementation details of FACINGBOARD-II, a fabric-based two-sided transparent display that shows independent, thus possibly different, graphics on its two sides. Our approach was built upon a projection film providing 'mixed transparency' (see Section 3.3.2), and explored a technological solution which is distinct and different from other efforts to implement two-sided transparent displays (Olwal et al., 2008; Lee et al., 2014).

Contribution #3, **exploring interaction techniques supporting collocated collaboration on two-sided transparent displays**. In Section 3.4, we illustrated a palette of interaction techniques that embodied the relaxed what-I-see-is-what-you-see interaction theme we proposed for collaborative transparent displays. These included enabling cooperative tools (selective image reversal, personal state, private work areas, and feedthrough) and visualization techniques that augment tracing and touch actions. The techniques we designed attempted to leverage the benefits of showing different content on both sides of the display in order to promote workspace awareness, which supports effective collaboration within groups.

Contribution #4: **investigating how transparency of transparent displays can be compromised, and how this in turn can severely affect workspace awareness**. Our work was the first (to the best of our knowledge) to take into account the variability of display transparency and its effect on workspace awareness as provided by collaborative transparent displays. In particular, we discussed the factors that affect display transparency (see Section 3.2.3, 3.3.2, and 4.1.1). Our control studies, which measured people's performance in perceiving *route* and *shape* gestures and touch actions in different transparency conditions, revealed that reduced transparency can compromise the workspace awareness provided by two-sided transparent displays (see Chapter 4 and 5).

Contribution #5, **evaluating the efficacy of using visualization techniques that augment actions to compensate for awareness loss resulting from reduced display transparency**. As outlined in Contribution #3, we proposed using augmentation techniques that visualize human actions to compensate for awareness loss when the display transparency is compromised. Our studies demonstrated that augmentation techniques can enhance workspace awareness when display transparency is compromised, and has no negative effect when the display is highly transparent (see Chapter 4 and 5).

6.2 Limitations

While we believe that our work provided significant contributions to the design of collaborative transparent displays, we also recognize the limitations of our efforts and believe that our research could be improved in the following three aspects:

Technology. Both FACINGBOARD-I and II were prototypes. While they functioned, they were both suboptimal in terms of display and touch input quality when compared to non-transparent interactive surfaces: low display resolution, low display contrast, slightly distorted color, less precise and accurate finger tracking, etc. We expect that advancements in related technologies can gradually address these technical challenges and help create much better collaborative transparent displays.

Design and evaluation context. Although targeted at supporting practical tasks, our interaction designs for collaborative transparent displays did not include a complete practical application. Subsequently, they have not been evaluated 'in-the-wild', within a real-world context, but rather with the artificial tasks described in Chapter 4. This limitation was discussed in more details in Section 5.6 and we will revisit it in Section 6.3.1.

Scope of design. Our collaborative transparent displays design was to a large extent motivated by supporting workspace awareness. Although workspace awareness affects many aspects of collaboration in workspaces, it does not capture all nuances related to human cooperative activities. For example, it is not related to how people behave in response to awareness information, such as how to manage information overload or make timely reactions. We believe that other design motivations will enrich and diversify the design of collaborative transparent displays beyond the guidelines and prototypes presented in this thesis.

6.3 Future Work

Our work had only scratched the surface of collaborative transparent display research. Below we discuss some of the future research threads that could emerge from our efforts.

6.3.1 Design for Real, Evaluate in Real

Our FACINGBOARD-I and II, and other recent similar systems (see Olwal et al., 2008; Heo et al. 2013; Lee et al., 2014) were all research prototypes that have never been evaluated in a real-world cooperative work setting. Although our design for collaborative transparent displays were based on considerations for practical applications, its appropriateness has never been validated for practical tasks, nor by people who would use it for practical purposes. We believe that for collaborative transparent displays to be truly useful, such evaluations and following redesign efforts are of critical importance. There is always a discrepancy between pure research thinking and practical needs.

In the short term, laboratory control studies that measure real-world task performance can be conducted to investigate the effect of low transparency, to evaluate interaction techniques, and to compare collaborative transparent displays with other platforms. Researchers have developed several methods to quantify awareness of others in the workspace. Some of them recorded performance measures specific to tasks (e.g. Ha et al., 2006), some counted the number of conflicts (e.g. Nacenta et al., 2007), and some focused on people's behaviors as indications of their awareness statuses (e.g. Hornecker et al., 2008). The unique challenge to evaluation of collaborative transparent displays is to find appropriate tasks where users can fully leverage the power of the platform, providing more valid findings and reflections. Some options to be explored are multiplayer games, design sessions with specified goals, and data visualization analysis. Results from these control studies can inform early design decisions, solving low-level usability and utility issues.

In the long run, evaluation of collaborative transparent displays will leave the laboratory and move into the natural setting. In such studies, high-fidelity prototypes will be deployed in offices, or other workplaces where they will be of practical use, to solicit feedback from real users doing real tasks in real environments. These prototypes can be first loaded with applications that support free sketching and writing, the type of work common in many workplaces. More features can be added as people suggest or require them, and they will be co-designed by designers and real users. A package of low-level application programming interfaces (APIs) that allows people to use their domain applications on transparent displays can be provided to users with related technical skills, as it would be interesting to see how people appropriate the unfamiliar platform for their own purposes. Participants will be regularly interviewed about their comments, suggestions, and typical interactions with the display, and the display will evolve in the course of the study according to their feedback. Such field studies will definitely be costly and lengthy, but long-term observation and the insights they can induce will be beneficial to the evolution of this new interactive medium.

Besides requiring resources and time, field studies of collaborative transparent displays will require technological advancements that go far beyond the current prototypes, supporting high quality graphics, precise sensing, and fluid interaction. The next section outlines our visions of future collaborative transparent display implementation research, focusing particularly on display technology.

6.3.2 Future Transparent Display Technology

Transparent display technology is still at its early stage of development. We believe it will keep evolving and offer better graphics quality, and better ways of showing different content on the two sides of the display.

The main limit on the graphics quality of current transparent displays is blending: that is, a viewer sees both the light from the graphics on the display and the light from real-world objects behind it. This mixture can distorts the visual appearance of both light sources. This phenomenon was examined in Chapter 4 and 5 of this thesis, where we saw that high density of graphics reduces the display transparency. Researchers have begun working towards solving this problem. For example, Sridharan et al. (2013) modeled color blending on transparent displays and proposed correction algorithms that find the alternate color that preserves the original color of the displayed content. However, for collaborative transparent displays, displayed content and people seen through the display are of equal importance, requiring new optical and software techniques to optimize and balance their visibility and visual fidelity.

A unique challenge for collaborative transparent displays is maintaining high graphics quality while presenting different content on both sides of the display. Although the fabric-based implementation in this thesis suffices for prototyping, the resolution, color fidelity, and transparency it affords are not ideal for deployment in natural settings. We believe that these limitations can be addressed by using new display materials. We foresee customized display panels that follow the principle of fabric displays (see Chapter 3), but are fine-tuned to have optimal hole size, 'thread' size, color, and reflection may be able to replace the projection film in FACINGBOARD-II. Further, materials that selectively diffuses light (see Hsu et al., 2014) can be used to make display panels that are opaque to light at particular wavelengths, but transparent to others. As such, they can display different images on both sides generated by light at those selected wavelengths while being highly transparent to ambient light.

6.3.3 Merging into Transparent Separators around Us

In previous chapters, we mostly considered a collaborative transparent display as a standalone setup, where it plays the role similar to that of a whiteboard in supporting group activities. However, another class of applications are possible if such displays can be integrated into existing transparent space separators, as described below.

People commonly construct and use transparent separators in order to partition a space while still leaving the view unobstructed. For example, glass walls have gained popularity in architectural design. Because they are transparent, they introduce a feeling of openness to the space. They allow physically separated people to view what others are doing. They offer a degree of security and privacy by defining territorial boundaries. They create an opportunity for people to communicate across the barrier. For example, glass walls in offices and meeting rooms allow others to see who is present and roughly what they are doing, which in turn increases the opportunity for casual interaction. Public spaces such as museums use glass walls to separate areas, while still bringing in a feeling of openness. In some banks and kiosks, clerks are protected by transparent barriers while customers can still talk through them. In prison visiting rooms, prisoners and their visitors are separated by transparent panels for security.

If augmented with collaborative transparent displays, existing transparent barriers can move from space separators to active communication facilitators. For example, people in offices can quickly set up ad-hoc discussions with people walking by in corridors using display-enabled transparent walls (e.g., to leave a note if someone appears busy inside). Bank transactions can be expedited if non-private information and actions can be shown or performed on protective panels with transparent displays. As illustrated in Chapter 1, surgeons in a sterile operation room can study medical imagery of patients with their colleagues in the adjacent non-sterile room on transparent walls. As related technologies advance, we anticipate that collaborative transparent displays can quietly merge into existing see-through separators, providing spatially separated people with additional workspaces for cooperative work.

6.4 Closing Remarks

This thesis was dedicated to exploring the interaction design of collaborative transparent displays, a new medium that affords a rather unconventional but ostensibly profitable form of collaboration. While we still cannot demonstrate their true usefulness in the real world, we are confident that collaborative transparent displays will find a niche in people’s cooperative activities. We believe that our research will provide part of the foundation and building blocks for future collaborative transparent display research, and hopefully will inform and inspire later researchs in this domain.

Reference

Azad, A., Ruiz, J., Vogel, D., Hancock, M. and Lank, E. (2012). **Territoriality and Behaviour on and around Large Vertical Publicly-Shared Displays**. In *Proceedings of ACM DIS '12*, 468-477.

Corning, Inc. (2011, 2012) **A Day Made of Glass I and II**. Youtube: v=6Cf7IL_eZ38 and v=jZkHpNnXLB0. Retrieved December 31, 2013.

Dispair (2014) **Displair—Air screen, Air display**. Retrieved August 5, 2014 from http://displair.com/

DiVerdia, S., Ismo R., Tobias H., and Alex O. (2006) **A Novel Walk-through 3D Display**. In *Proceedings of SPIE 2006 Electronic Imaging*, 428-437.

Dix, A. (1994) **Computer Supported Cooperative Work: A Framework**. In *Design Issues in CSCW*, 23-37.

Dourish, P., and Bellotti, V. (1992) **Awareness and Coordination in Shared Workspaces**, In *Proceedings ACM CSCW '92*, 107-114.

Downing, E., Hesselink, L., Ralston, J. and Macfarlane, R. (1996) **A Three-color, Solid-state, Three-Dimensional Display**. In *Science*, 273, 1185–1189.

EarlyTechNews. (2014) **Samsung Transparent OLED Laptop**. Retrieved September 21, 2014 from http://earlytechnews.wordpress.com/tag/samsung-transparent-oled-laptop/

Endsley, M. (1995) **Toward a Theory of Situation Awareness in Dynamic Systems**, In *Human Factors*, 37(1), 32-64, 1995.

FogScreen (2014) **FogScreen Projection Screen**. Retrieved August 5, 2014 from http://www.fogscreen.com/

Genest, A., Gutwin, C., Tang, A., Kalyn, M. and Ivkovic, Z. (2013) **KinectArms: a Toolkit for Capturing and Displaying Arm Embodiments in Distributed Tabletop Groupware**. In *Proceedings of ACM CSCW '13*, 157-166.

Gessler, N. (2014) **Military War Room Command and Control Center Situation Maps**. Retrieved September 29, 2014 from https://web.duke.edu/isis/gessler/topics/warrooms.htm

Greenberg, S., Bohnet, R., Roseman, M. and Webster, D. (1992). **GroupSketch**. *In ACM SIGGRAPH Video Review: Special Edition of the ACM CSCW '92 Technical Video Program*, 87.

Gutwin, C. (1997) **Workspace Awareness in Real-Time Distributed Groupware**. PhD thesis, Department of Computer Science, University of Calgary, Calgary, Alberta, Canada, December 1997.

Gutwin, C. and Greenberg, S. (1998) **Design for Individuals, Design for Groups: Tradeoffs between Power and Workspace Awareness**. In *Proceedings of ACM CSCW '98*, 207-216.

Gutwin, C. and Greenberg, S. (2002) **A Descriptive Framework of Workspace Awareness for Real-Time Groupware**. In *Journal of CSCW*, 11(3-4):411-446.

Gutwin, C., Greenberg, S. and Roseman M. (1996) **Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation**. In *Proceedings of HCI '96*, 281-298.

Gutwin, C., Penner, R. (2002). **Improving Interpretation of Remote Gestures with Telepointer Traces**. In *Proceedings of the ACM CSCW '02*, 49-57.

Ha, V.; Inkpen, K.M.; Mandryk, R.L.; Whalen, T., (2006) **Direct Intentions: the Effects of Input Devices on Collaboration around a Tabletop Display**, In *Proceedings of IEEE TABLETOP '06*, 177-184.

Hall, E.T. (1966). **The Hidden Dimension**. Doubleday.

Heo, H, Kim, S., Park, H., Chung, J., Lee, G., and Lee, W. (2013). **TransWall**. In *Proceedings of ACM SIGGRAPH '13* Emerging Technologies. Article No. 14.

Hornecker E., Marshall, P., Dalton, N. S., and Rogers Y. (2008) **Collaboration and Interference: Awareness with Mice or Touch Input**. In *Proceedings of ACM CSCW '08*, 167-176.

Hsu, C. W., Zhen, B., Qiu, W., Shapira, O., DeLacy, B. G., Joannopoulos, J. D., & Soljačić, M. (2014) **Transparent Displays Enabled by Resonant Nanoparticle Scattering**. In *Nature Communications*, 5, 3152.

Humphries, M. (2010). **Samsung's transparent 46″ LCD on display at SID 2010**. Retrieved September 10, 2014 from http://www.geek.com/chips/samsungs-transparent-46-lcd-on-display-at-sid-2010-1259424/

I2O Technologies. (2014). **Overview**. Retrieved August 5, 2014 from http://www.io2technology.com/io2-overview.htm

Ishii. H., and Kobayashi, M. (1992) **ClearBoard: a Seamless Medium for Shared Drawing and Conversation with Eye Contact**. In *Proceedings of ACM CHI '92*, 525 - 532.

Kent Optronics. (2014) **Transparent Display**. Retrieved August 5, 2014 from http://www.kentoptronics.com/transparent.html

Krueger, W. M., Gionfriddo, T., and Hinrichsen, K. (1985) **VIDEOPLACE—an artificial reality**. In *Proceedings of ACM CHI '85*, 35-40.

Kuo, H., Hubby, L. M., Naberhuis, S. and Birecki, H. (2013) **See-through Display**. US Patent 8,462,081 B2. Filed Mar 9. 2011. Issued Jun 11. 2013.

Lee, H, Cho, S, Hong, J, Lee, G and Lee, W. (2014) **JANUS**. In *Proceedings of ACM SIGGRAPH 2014* Emerging Technologies Article No. 15.

Li, J, Sharlin, E., Greenberg, S., and Rounding, M. (2013) **Designing the Car iWindow: Exploring Interaction through Vehicle Side Windows**. In *ACM CHI EA '13*, 1665-1670.

Li, J., Greenberg, S., Sharlin E., and Jorge, J. (2014) **Interactive Two-sided Transparent Displays: Designing for Collaboration**. In *Proceedings of ACM DIS '14,* 395-404.

Liu, J., Sun, X. (2006) **System and Method for a Transparent Color Image Display Utilizing Fluorescence Conversion of Nano particles and Molecules**. US Patent 7090355 filed 18 May 2004, and issued 15 Aug 2006.

Lumineq. (2014) **Transparent Electroluminescent Displays, TASEL**. Retrieved August 5, 2014 from http://lumineq.com/en/products/tasel

Lysikatos, G. (2012). **Innovative Technologies**. Retrieved September 10, 2014 from http://typographyininformationdesign.blogspot.ca/2012/05/innovative-technologies.html

Microsoft. (2014). **Kinect for Windows**. Retrieved September 27, 2014 from http://www.microsoft.com/en-us/kinectforwindows/

Morris, M. R., Huang, A., Paepcke, A., Winograd, T. (2006) **Cooperative Gestures: Multi-User Gestural Interactions for Co-located Groupware**. In Proceedings of the *ACM CHI ’06*, 1201-1210.

Morris, M.R., Paepcke, A, and Winograd, T., (2006) **TeamSearch: Comparing Techniques for Co-present Collaborative Search of Digital Media**. In *Proceedings of IEEE TABLETOP ’06,* 97-104.

Nacenta, M.A., Pinelle, D., Stuckel, D., Gutwin, C. (2007). **The Effects of Interaction Technique on Coordination in Tabletop Groupware**. In *Proceedings GI'07*, 191-198.

Olwal, A. Lindfors, C., Gustafsson, Kjellberg, T. and Mattsson, L. (2005) **ASTOR: an Autostereoscopic Optical See-through Augmented Reality System**. In *Proceedings of ISMAR ’05*, 24-27.

Olwal, A., DiVerdi S., Rakkolainen, I., and Höllerer, T. **Consigalo: Multi-User Face-to-Face Interaction on Immaterial Displays**. In *Proceedings of ACM INTETAIN ’08.* Article No.8.

phoneArena.com. (2012). **Transparent-display Explay Crystal is a Phone Merging the Past and the Future**. Retrieved September 10, 2014 from http://www.phonearena.com/news/Transparent-display-Explay-Crystal-is-a-phone-merging-the-past-and-the-future_id30667

Planar Systems, Inc. (2014) **Planar LookThru Series**. Retrieved August 5, 2014 from http://www.planar.com/products/transparent-displays/

Polhemus. (2014). **Polhemus is the Premier Precision Motion Tracking Company**. Retrieved September 27, 2014 from http://polhemus.com/

Rekimoto, J. and Saitoh, M. (1999). **Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments**. In *Proceedings of CHI '99,* 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/LH22NLBVTVC/EN

Scott, S. D., Carpendale, M. S. T., and Inkpen, K. M. (2004) **Territoriality in Collaborative Tabletop Workspaces**. In *Proceedings ACM CSCW '04*, 294-303.

Shen, C., Everitt, K.M., and Ryall, K. (2003). **UbiTable: Impromptu Face-to-Face Collaboration on Horizontal Interactive Surfaces**. In *Proceedings of UbiComp'03*, 281-288.

Sridharan, S.K., Hincapié-Ramos, J.D., Flatla, D.R. and Irani, P. (2013) **Color Correction for Optical See-through Displays Using Display Color Profiles**. In *Proceedings of ACM VRST '13*. ACM, 231-240.

Sugimoto, M, Hosoi, K., and Hashizume, H., (2004) **Caretta: a System for Supporting Face-to-Face Collaboration by Integrating Personal and Shared Spaces**. In *Proceedings of ACM CHI '04*, 41-48.

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, A., Boyle, M. and Greenberg, S. (2004) **Display and Presence Disparity in Mixed Presence Groupware**. In *Proceedings of AUIC '04*, 73-82.

Tang, A., Neustaedter, C., and Greenberg, S. (2006). **VideoArms: Embodiments for Mixed Presence Groupware**. In *Proceedings HCI 2006*, 85-102.

Tang, J. and Minneman, S. (1990) **Videodraw: A Video Interface for Collaborative Drawing**. In *Proceedings of ACM CHI '90*, 313-320.

Tang, J. and Minneman, S. (1991) **VideoWhiteboard: Video Shadows to Support Remote Collaboration**. In *Proceedings of ACM CHI '91*, 315-322.

Tang, J. C. (1991) **Findings from Observational Studies of Collaborative Work**. In *International Journal of Man-Machine Studies*, 34, 143-160.

Tse, E., Greenberg, S., Shen, C. and Forlines, C. (2006) **Multimodal Multiplayer Tabletop Gaming**. In *Proceedings Third International Workshop on Pervasive Gaming Applications (PerGames'06)*, in conjunction with 4th Intl. Conference on Pervasive Computing, 139-148.

Wu, M. and Balakrishnan, R. (2003). **Multi-finger and Whole Hand Gestural Interaction Techniques for Multi-user Tabletop Displays**. In *Proceedings of UIST'01*, 193-202.

Appendix. A Study Materials

This appendix contains supplementary materials for the study described in Chapter 4 of this thesis. It includes:

- The letter set used in the shape task (see Section 4.2.1)
- Approval from the University of Calgary's Conjoint Faculties Research Ethics Board to perform the study in question.
- The informed consent form given to study participants.
- The experimental protocol, which describes the actions taken by the administrator during the study.
- The pre-study questionnaire issued to participants before they began to perform the tasks.

A.1 Letter Set Used in the Shape Task

The 12-letter set from which the letters presented to the viewer (the participant) in the *shape* task were selected is as follows:

A C D G I J L M N O S U V W Z

We selected these letters as they all comprises relatively fewer strokes among the alphabet, making it easy for the actor to write and for the viewer to recognize.

A.2 Informed Consent Form

UNIVERSITY OF CALGARY

Name of Researcher, Faculty, Department, Telephone & Email:

Jiannan Li
M.Sc. Student
Department of Computer Science
E-mail: jiannali@ucalgary.ca
Phone: 403 399 8791

Ehud Sharlin
Associate Professor
Department of Computer Science
E-mail: ehud@cppsc.ucalgary.ca
Phone: 403 210 9404

Saul Greenberg
Professor
Department of Computer Science
E-mail: saul.greenberg@ucalgary.ca
Phone: 403 220 6087

Title of Project:

Correlation of head position to viewing direction

This consent form, a copy of which has been given to you, is only part of the process of informed consent. If you want more details about something mentioned here, or information not included here, you should feel free to ask. Please take the time to read this carefully and to understand any accompanying information.

The University of Calgary Conjoint Faculties Research Ethics Board has approved this research study.

Purpose of the Study:

This project aims to study some interaction techniques which will be used on a new interactive system, the see-through collaborative display. This system, currently at the early prototyping stage as you may see, allows people facing each other to work on it while being able to see each other. The material of the display enables it to show different contents on both sides and still to be transparent to some extent. We are particularly interested in how to help people see each other's action and intentions even when the display is full of contents and no more transparent. So we designed some techniques to enhance the display. In this study, we would like to evaluate the efficacy of these techniques as well as how much visual clutters on the display affects its transparency. We invite you to participate in this study and to help the design of this promising new medium.

What Will I Be Asked To Do?

We will first ask you to about some basic demographic information and your experience with some computer interfaces. You will be performing a series of brief tasks on a semi-transparent display. Generally speaking, for each task you will see something on or through the display, and will be told to perform some actions (such as touching a particular part of the display) depending upon what you see. You may also be asked to do some simple collaborative tasks with another participant on the display. Low level details of what you are expected to do will be provided verbally before and during the study.

During the task, the camera will track the markers on your hand and record when you touch and where you touch. After these tasks, you will be asked to fill in a questionnaire about your opinion on this system.

Any questions with regards to the study are welcomed. Though we may not be able to answer them during the study, they will be answered in the break or after the study.

This experiment is expected to take about 60 minutes. We do not foresee any risks from participating in this study.

Keep in mind that your participation in this study is completely voluntary. You are free to withdraw from the experiment at any time without any kind of penalty. Your data collected to that point will be deleted. If you decide to withdraw, the experiment will be interrupted immediately. Participating in the experiment will grant you a total of $15. Should you choose to withdraw, you will still be allowed to keep the $15 if you have completed half of the study.

What Type of Personal Information Will Be Collected?

Should you agree to participate, you will be asked to provide basic demographic information and experience with computer interfaces. We will also use a tracking system to record your bodily movement. After the study, we would like to hear your opinions about the prototype and some techniques used on it.

We will take notes as you interact. Aspects that we might write down include: problems occurring during the interaction, such as when the application does not respond as expected, or some opinions that you may state. We may also record video and audio of this session with your explicit consent (see below).

Are there Risks or Benefits if I Participate?

There are no known harms or risks associated to the participation in this study. If you participate you will receive a compensation of $15 for your time. You will also have the opportunity of using a state-of-the-art interface facility.

What Happens to the Information I Provide?

The researchers will record your interaction with the prototype using software, your responses to the interview and/or the questionnaire. Only the researchers will have access to the full recordings and the responses that you provide.

This information will be kept in a secure location (locked cabinets and password-protected drives). The information that we collect will not be associated to you personally. However, the researchers will publish the results of their analysis of your data in anonymized form in academic journals and conference papers.

The researchers might quote the responses in the questionnaires or any of your comments in anonymized form. Please check the boxes below to confirm that you understand this use of your data.

____I agree that the researchers may use any written or verbal comments and answers I may provide in research presentations and publications.

____ I consent for this session to be photographed.

____ I consent for this session to be video and /or audio recorded.

____I agree that the researchers may use some of the photographs, video and/or audio content of my actions to illustrate their findings in research presentations and publications. I realize that once the video has been released through publications or presentations it is not under the control of the researchers who has access to it.

All the collected data will be kept by the investigators for at least a year, where it will be destroyed after it is no longer required.

Signatures (written consent)

Your signature on this form indicates that you 1) understand to your satisfaction the information provided to you about your participation in this research project, and 2) agree to participate as a research subject.

In no way does this waive your legal rights nor release the investigators, sponsors, or involved institutions from their legal and professional responsibilities. You are free to withdraw from this research project at any time. You should feel free to ask for clarification or new information throughout your participation.

Participant's Name: (please print) ____

Participant's Signature ____________________________________Date: ____________

Researcher's Name: (please print) ____

Researcher's Signature: ________________________________Date: ____________

Questions/Concerns

If you have any further questions or want clarification regarding this research and/or your participation, please contact:

Jiannan Li
403 399 8791, jiannali@ucalgary.ca
or
Ehud Sharlin
403 210 9404, ehud@cpsc.ucalgary.ca
or
Saul Greenberg
403 220 6087, saul.greenberg@ucalgary.ca

If you have any concerns about the way you've been treated as a participant, please contact an Ethics Resource Officer, Research Services Office, University of Calgary at (403) 220-3782; email cfreb@ucalgary.ca

A copy of this consent form has been given to you to keep for your records and reference. The investigator has kept a copy of the consent form.

A.3 Experimental Protocol

Instruction

The following description will be read to each participant at the beginning of the study to inform participants of the procedures prior to giving consent. Italicized text are instructions to the investigator.

Note: The questionnaires and verbal protocol included are indicative of what we will ask and say. Minor modifications may be made to smooth the process. The task in this instruction serves as an example of the tasks we would like participants to perform. Additional questions may be asked depending upon particular comments and / or actions observed as the study progresses.

<The side of the display where the experimenter works is named Side A. The other side, where the participant works is named Side B.>

<Before the participant enters the room, put the homepage of iLab on side B and put a "welcome" slide on side A>

<Before the study starts, start the software. Guide the participant to Side A of the display>

Hello, welcome to the iLab. My name is *<experimenter>* and I am running this study for our transparent display prototype. I'll guide you through the study and feel free to ask me any questions at any time.

Before starting I would like to let you know about your rights as a participant.

- If you don't want to proceed you may quit at any time. The data that we have collected up to that point will be permanently removed.
- No data will be used without your explicit consent.

First, I'd like to introduce the purpose of our study.

We are now designing the interfaces for a collaborative transparent display. *<Point to the display>*Here you can see a prototype of the display we are working on. The projectors above you project visual contents on the display in front of you, and the visual contents on both sides are not necessarily the same. Two people can stand on the opposite of the display *<go to Side B>* and as you can see, you can see me through this piece of semi-transparent material. We think that being able to see your co-worker while doing collaborative work can bring some interesting opportunities. Now we would like to ask to you to help us with studying the properties of this system and also evaluating some interfaces on it. *<Go to Side A>*

Now please read this consent form carefully, as it explains your rights as a participant and the conditions of the study, and sign it if you agree with these terms. *<hand form to the participant, give them time to read on their own>*.

Now, I would like to let you know that as long as you finish half of the study you will receive the $15 payment.

Although I don't know of any reason for this to happen, if you should become uncomfortable or find this study objectionable in any way, you are free to quit at any time. Also if you would like to take a break just let me know and I will pause the study.

You may have a copy of the consent form for your own records.

Before we start the actual task, please fill in a short pre-study questionnaire.

Now I'm going to introduce you to your tasks. The study is consisted of 3 separate tasks. In all these tasks, you will stand on that side *<side B>* and I will stand on this side *<side A>*. From now on, I will call this side "my side" and that side "your side". If you find the tasks hard to perform, don't worry. It's not your problem and that's exactly where we learn. Let's start from the first task, named "character recognition". *<Guide the participant to Side B, stand on Side A>*

I am going to write a CAPITAL ENGLISH LETTER near the display with my finger, and you'll tell me what you think that letter is. The letter will be written in the orientation that's readable for you. For example, *<write a mirrored "B">*, what letter is this? Yes, you're right. You may see various backgrounds, and a trace showing where my fingertip is in some rounds of the study. In each rounds, I will write 12 letters.

Do you understand what your task will be? Do you have any questions? OK, let's do a trial with three letters first.

Are you ready for the actual study? Good, let's start.

<Press key "D", start the task. After each round, press "T" to terminate>

Ok, the first task is done. *<Go to side B>*Now let me ask you a few questions.

<After the interview, let the participant take a break>

Before moving on to the second task, please wear this glove on your right hand. The shiny balls on the glove allow the camera to track your finger position, so don't cover them.

Now let's move on to the second task, named "path recognition". *<Press "A" to start the task>*

Can you see some circles on your side? Each of these circles is called a target. In the task, I will draw a path near the display that go through some of these targets following a certain path. *<Show a path to the participant.>* Like this. And you'll observe this path and select the targets that are on the path. *<Go to Side B>* Note first you only need to touch the start, the end and the corner of the path. *<Show an example>* And second, the touch input works slightly differently from that on a mobile phone or a tablet. You don't have to press the screen, instead, just get close to the target *<demonstrate it>* and it will change its color. You can try it yourself. *<Let the participant do a few more trials>* When you've selected the start, the end and the corner, let me know. *<Go back to Side A>* Like in the first task, there will still be different backgrounds, and the trace visualization. This subtask also has eight rounds, and in each round I'll do 8 gestures.

Do you understand what your task will be? Do you have any questions? OK, let's do a trial with 2 paths.

<*Do the trial, terminate it and restart*>

Ok, the second task is done. <*Go to side B*>Now let me ask you a few questions.

<*After the interview, let the participant take a break*>

<*Go to side A*>Now let's move to the last task, named "mixed focus". <*press "I" initiate the task*> In this task you need to pay attention to both sides of the display. You have some work to do on your side, but also need to look at my actions. On your side, you can see a dark green square on the display. It is your target. Touch it when you see one. When it's touched, it will be repositioned at a random spot on the display <*press "S" let the participant touch the target*>.

At the same time, I will touch some spots on my side. When I touch, you should touch where I have touched immediately, and then touch your square target. I would like you to follow my touch as fast as possible, by observing my actions and make prediction, just like you reaching out your hand in a handshake. <*let the participant experience for a while*> <*Go to side B*>Note that there is a countdown timer on each of your target. You can only touch the target when the countdown goes to 0 and disappear. <*Let the participant try it.*>

<*Go back to Side A*> During some rounds of the study, I'll turn on this visualization <*turn on the visualization*> and it indicate where my fingertip is. When the dot changes its color to orange, it means the finger is very close to the display. Like in the first and the second task, you will see different backgrounds.

Do you understand what your task will be? Do you have any questions? OK, let's do a trial first.

<*After the trial, start the actual task*>

Ok, the last task is done. <*Go to side B*>Now let me ask you a few questions.

There comes the end of our study, thanks so much for participating!

<*Hand the participant the $15 payment, and let them sign the form*>

<*End of the study*>

A.4 Pre-study Questionnaire

Pre-study Questionnaire

Help us with designing collaborative see-through displays

Participant #

Project Title: Interaction Mechanisms for Two-Sided Transparent Displays

Demographic Information

1. Gender ____F ____M
2. Age ____
3. Professional ____________

Computer Interface Experience

1. How often do you use touch interfaces (such as the touch screen on smartphones)?

   ____Everyday ____Several times a week ____Several times a month

   ____Less than all above

2. Have you ever interacted with large displays (not just watching)?
   ____I often interacts with large displays
   ____I've interacted with large displays for several times
   ____I've seldom or never interacted with large displays

3. Have you ever interacted with motion sensing based interface (such as Microsoft Kinect)?
   ____I often interacts with these interfaces
   ____I've interacted with these interfaces for several times
   ____I've seldom or never interacted with these interfaces

Appendix. B Augmented Reality for Vehicle Passengers Using Transparent Displays

This appendix documents iWindow, a side project which sought to enhance vehicle side windows with transparent displays, providing location-based augmented reality information to passengers. Although not closely related to the topic of this thesis, it presents an interesting instance where transparent displays can be used to augment our daily lives. In addition, building iWindow helped us better experience the characteristics of transparent displays.

B.1 Abstract

Interactive vehicle windows can enrich the commuting experience by being informative and engaging, strengthening the connection between passengers and the outside world. We propose a preliminary interaction paradigm to allow rich and un-distracting interaction experience on vehicle side windows. Following this paradigm we present a prototype, the Car iWindow, and discuss our preliminary design critique of the interaction, based on the installation of the iWindow in a car and interaction with it while commuting around our campus.

B.2 Introduction

Automobiles have served humans for more than a century and are continuing to be important in modern transportation. Drivers and passengers are holding expectations for richer in-vehicle experiences as they spend significant amount of their daily time commuting in vehicles. Over the years, various improvements have been made to turn automobiles from merely transportation tools to a livable and comfortable space. Examples include high fidelity radios and media players, video consoles, and even refrigerators.

However, most cars still lack interactivity and information richness. This void has been filled somewhat by passengers using smart phones or tablets.

Yet the car provides a specific context that could be leveraged as part of the interactive experience. Commuting continuously provides passengers with new stimulus and visual scenes as their travel unfolds, as viewed through the car windows. These scenes often provoke interest and perhaps a desire for related information. Thus it is common to see passengers seeking information about a landmark they have seen via the car window, e.g. a community, a waterfall, a restaurant. This is usually done through their smart phone. But what if the car itself could become the information appliance, where it could show that information in context? Like others, we see vehicle windows as a natural medium to provide such contextual information to passengers, where these windows provide location-aware interactive display capabilities.

In practice, front windows of aircrafts have been used to show information to pilots, which assists aviation and target-aiming. These are commonly known as head-up displays (HUDs). There are attempts to transplant this technology to automobiles. However, due to the risk of driver distraction, many HUDs in automobiles are basically passive digital representations of existing dashboards and GPS navigators, where they offer little in the way of interactivity [1]. Thus it is the passenger – rather than the driver – that has been considered as the end user. Several commercial visions have been produced that simulate an interactive side windows [2][3], through which the passenger views and interacts with the world while commuting in a vehicle, a car, taxi, tour bus, a mass public transportation such as a train, or an airliner. This paper describes our efforts of trying to bring these visions closer to reality.

Figure B.1: The setup of the Car iWindow

We believe that an interactive vehicle window should be informative but not distracting. We propose a simple 3-phase interaction paradigm to realize rich and un-distractive interaction on side windows. We then present our prototype, the Car iWindow (Figure B.1), whose design follows our side window interaction paradigm and is implemented using a transparent LCD display installed in a car. Using the iWindow prototype we ran a Wizard of Oz (Woz)-operated design critique, where we reflected on a child's interaction experience during a drive around our university campus. We hope that our effort can highlight some of the challenges and promises of this interaction design problem, and serve future explorations of interactive side windows.

B.3 Related Work

There are two impressive future-envisioning videos that imagine enriched in-vehicle spaces equipped with interactive side windows.

In 2011 Microsoft released a video envisioning a future in which a travelling businesswoman can see the current time and the highlighted hotel where her meeting is going to be held via her taxi side-window [2]. The video briefly illustrates possible opportunities enabled by interactive vehicle windows in a combination of location-based applications.

In 2012 General Motors introduced their Window of Opportunity concept [3] in cooperation with Israel's Bezalel Academy of Art and Design. The video illustrates four creative applications for interactive side windows constituting a spectrum of novel riding experiences. In addition, a static car-like prototype is built to demonstrate the concept, using two external projectors. One simulates the outdoor scene, and the other projects the content on the window-screen.

These envisionments [2,3] conceptualize the interactive side window but stop short of actually implementing, installing and evaluating the user experience in-situ, i.e., a passenger in a car driving in the physical world.

In a related effort, Olwal [4] evaluated various interaction techniques for transparent displays, including touch, mobile device control, hand gestures and eye-tracking with a prototype named ASTRO (not necessarily in a car setting). The results indicate that hand gestures and eye-tracking are overall less preferable than touch, a conclusion that informed our design of the iWindow prototype.

B.4 Design

The key question which drives our design process is: what is the purpose of a digital vehicle side window? We believe that most passengers would like to remain intrigued by the rich physicality of the outside world, and by the changing environments they view through the vehicle's window. Our answer to the question is to use the digital side window to tighten the connection between the physically isolated passenger and the outside environments, rather than to create more disengagement and separation. We are aware that the information superimposed an interactive window is likely to be distractive, or that any presented information may obscure real scenes. Thus our interaction design approach attempts to mitigate distraction caused by the iWindow visuals, while still maintaining its informative goal. Our iWindow design approach is based on three interaction phases, each with a different distraction potential: active notifications, ever-present widgets and information conjuring (Figure 2).

Our design pursues equilibrium between the information the user seeks about the scene viewed through the iWindow, and the potential for disturbances. Thus, the interaction phase containing higher risk of distraction is designed to provide less information, and vice versa. Figure B.2 is a schematic diagram relating our three interaction phases, the probability of distracting in each of them, and the information volumes associated with them.

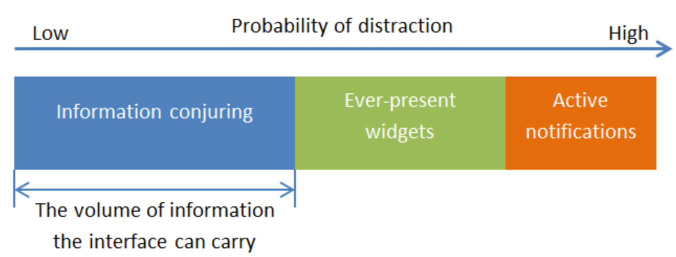

Figure B.2: A schematic diagram of the 3 interaction phases. The probability of distraction increases from left to right and the phase areas represent the volume of information the interfaces in this phase can carry.

Active notifications pop up on the window to inform passengers of pre-defined types of events which they cannot easily perceive. One interesting possibility is supporting serendipitous finding. For example, if the user expresses an interest in "churches built before 1800" the churches in this category will be highlighted when passed. Pop-up notifications are the most distractive amongst the three phases, so they are only activated for passenger-defined events and are designed to convey the least amount of information. *Ever-present widgets* are information sources always visible on the window. They indicate simple and general information such as the time and temperature. An ideal widget should be presented in an unobtrusive, even ambient way, for example, hidden in the lower bottom corner of the iWindow, thus lowering the potential for distraction.

Information conjuring refers to displaying information in response to a passenger's explicit request. For instance, if the passenger touches the window where an old bridge is seen, information related to the bridge is shown. Since they are response to expressed request, interfaces in this phase allow the passenger to browse much richer content than in the other two phases. In order to determine which target the user is specifying, the iWindow approximates his/her line-of-sight with a line from the estimated head position to the touching fingertip. Combining this with map databases, the area being pointed at can be identified and related information is then revealed.

These 3 phases together form an interaction space in which passengers benefit from a comfortable balance between augmentation and reality.

B.5 Implementation and Critique

Following the above design approach, we implemented a prototype we call the car iWindow. We install a Samsung 22'' transparent LCD display panel connected to a control PC in a Kia Sorrento SUV as an interactive side window (Figure B.1). We used the iWindow in a Woz design critique session, where a 6 year old participant was sitting in the 2nd row of the SUV and interacting with the iWindow as the car was driven around our campus. The experiment administrator sat in the 3rd row of the SUV, and operated the iWindow via Woz. Head position estimation and touch sensing is not realized in the current prototype.

Our participant was given brief explanation about the basic functions of the iWindow, the role of the ever present widgets, the active notifications and told that she will need to touch the iWindow when she saw a building invoking her curiosity (initiating the conjuring phase). After this brief explanation, the car was driven around campus with its actual side window all the way down, and the iWindow visually replacing it (although physically not covering the entire window space). The Woz administrator sitting behind the passenger generated and manipulated all the information displayed on the iWindow using a basic Woz iWindow software tool we prototyped.

The three *ever-present widgets* indicating the time, temperature and the direction to which the window is facing were located at the top-right corner of the display (Figure B.3a). The direction was presented by a rotating 3D compass visualization, pointing to the north.

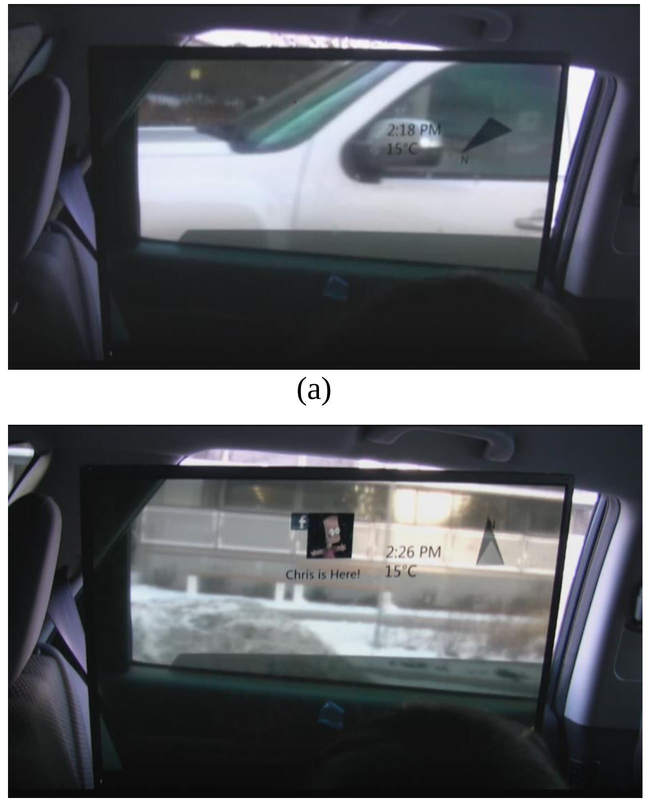

(a)

(b)

Figure B.3: (*a*) The *ever-present widgets* present the time, temperature and the direction. (*b*) The *active notification* indicating that Chris is in this building

When the car passed by a certain building, the iWindow showed a cartoon avatar along with texts saying “Chris is Here!” superimposed on the building to show the user’s friend’s hypothetical location as an *active notification* (Figure B.3b). After the building was out view of the iWindow (and thus out of the passenger’s view) that notification was turned off.

As the passenger held her fingertip on a particular building seen through the iWindow, a text block expanded from where she touched and eventually revealed its name to complete an information conjuring process. The expanding interface, as being “conjured” by the touching finger, was designed to confirm the user’s intent for responses (and visual obscures at the same time) and to avoid unwanted disturbance caused by casual contacts. In addition, the user could hide the interface simply by dragging it aside. We note that, in our design process, we were unsure about whether touching a window would appear natural to passengers. However, in a limited space like a private car, touch input uses space more efficiently as compared to pointing or gesture. Olwal’s evaluation [4] proves that touch is still welcomed in interactions with transparent displays and, in our tour through the campus, touch as a input method was learned and performed without issue by the young participant.

B.6 Conclusion and Future Work

Inspired by visions of more interactive and informative in-vehicle environments, the iWindow explores the interaction space of vehicle side windows. In order to allow future interactive side windows to enhance riding experiences we proposed an interaction paradigm aiming at creating a strong and balanced information connection between passengers and outside environments. This paradigm, consisting of active notifications, ever-present widgets and information conjuring, tries to offer considerable interactivity and information while minimizing visual disturbance. Based on this interaction model we designed the Car iWindow prototype and presented its Wizard of Oz design critique in a car.

Our current iWindow prototype and its evaluation are very preliminary and still need considerable improvements. First, to evaluate the design more thoroughly, a high-fidelity prototype should be built. These could incorporate location-aware sensors such as GPS systems, touch sensing, and algorithms linking vehicle positions and passenger inputs to the adjacent environment. Second, the information content needs to be expanded beyond the extremely simple information available in our prototype. Third, a larger study involving more participants from diverse age groups needs to be conducted to find answers to some important questions about iWindow usability and user experience. Questions include: are people comfortable with the 3-phase interaction paradigm when moving in fast-changing environments? What is the best input method for interacting with interactive car side windows? Are superimposed texts and images capable of transmitting location-based information clearly, especially in urban areas crowded by dense buildings which make an ununiformed clutter background?

We would also like to explore the possible application of our simple 3-phase interaction paradigm, although originally formulated for ensuring undisturbed viewing experiences through interactive side windows, in a broader range of displays, and whether it could be extended to serve as a model for analyzing cognitive loads of elements comprising other interactive systems.

B.7 References

[1] Wu, W., Blaicher, F., Yang, J., Seder, T., Cui, D. A Prototype of Landmark-based Car Navigation Using a Full-windshield Head-up Display System. In Proc. AMC 2009, ACM Press (2009), 21-28.

[2] Productivity Future Vision.

http://www.youtube.com/watch?feature=player_embedded&v=a6cNdhOKwi0.

[3] 2012 GM Advanced Tech Window.
https://www.youtube.com/watch?v=9G4cYYlIaCc.

[4] Olwal, A. Unencumbered 3D Interaction with See-through Displays. In Proc. NordiCHI 2008, ACM Press (2008), 527-530.