Waves:
A Collaborative Navigation Technique for Large Interactive Surfaces

by

Joseph Shum

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2014

AUTHOR'S DECLARATION

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Digital tables offer the possibility of performing collaborative tasks where two or more people can share artifacts in the same virtual space. However, most interactive methods of navigating through virtual space most commonly have the effect of changing the entire digital display simultaneously. In this thesis, I performed an exploratory study providing evidence for differences between two popular collaborative navigation techniques used in video games, split screens and single shared screen, in situational awareness, interference between collaborators, and difficulties with automatic view adjustment. Drawing inspiration from guidelines formulated from the results of the exploratory study, as well as previous work in interactive tabletops, collaboration, and navigation in information visualization, I designed and implemented Waves, a collaborative navigation technique for the tabletop. Waves simultaneously supports multiple personal workspaces, provides group workspace awareness, and mediates interference between workspaces.

Acknowledgements

It has been a long road. The fact that I was able to get here was due in no small part to kindness of many people to whom I owe my deepest gratitude.

First, I would like to thank my supervisor, Dr. Mark Hancock, for taking a chance on me and taking me in as a graduate student. I am extremely grateful for your guidance and teachings over the years. I will never forget your advice to me on “don’t try to write the Internet.” I truly tried. Hopefully, I did not disappoint. Most of all, I am most thankful for your tremendous patience and flexibility in letting me explore my career during these years.

Second, I would like to thank Dr. Stacey Scott, for first introducing me to Dr. Mark Hancock and agreeing to be my thesis reader. Thank you for opening my eyes to CSCW and tabletops. My thesis is built on a lot of your previously work.

I would also like to thank Dr. Michael Terry for introducing me to human computer interaction in your course CS489 and agreeing to be my thesis reader. Your insistence that I learn contextual design method properly greatly influenced my career direction. It was the primary reason why I pursued graduate school in this field.

Finally, I would like to thank my friends and family for encouraging me to never stop writing until I am done. I would like to thank my parents for their unconditional love, and my dear brother Wil, for running errands for my thesis and many other things when I am away.

Finally, to Xiaochen, you have been my constant source of encouragement, discipline, and emotional support throughout my entire graduate career. Thank you for being part of my life.

Dedication

To my grandmother, who had reminded me time and again that education is a privilege to be cherished.

Table of Contents

List of Figures

List of Tables

Chapter 1
Introduction

When two or more people gather to collaborate, each person's perception of the space will be different, because each person will occupy a different physical location. For example, two architects can stand on opposite sides of a building model, with one looking closely at the front entranceway, while the other stands back and tries to get a sense of the overall appearance from the rear. Both architects can benefit from the freedom to navigate independently around the environment, but also can easily grasp the connection between their own perspective and where their collaborator is working in that same environment.

Digital tables offer the possibility of performing collaborative tasks where two or more people can share artifacts in the same virtual space. However, most interactive methods of navigating through virtual space most commonly have the effect of changing the entire digital display simultaneously. For example, if two people are interacting with the same map, one person might pan the map to the left (e.g., by dragging their finger across the screen), and the entire screen moves, potentially moving the second person's work area off the edge of the screen. One application area that has begun to address navigation by two people in the same virtual world is video games, where multiple users can share a single screen in co-located gameplay. Very much like collaboration using the tabletop, video console game players engage in cooperative play scenarios and need to navigate the game environment using a single display.

In this thesis, I investigate two popular video game screen sharing techniques in an exploratory study, paying particular attention to their impact on people's understanding of the virtual space. The results show that providing separate control for navigation can improve people's understanding of their own presence in the space, but can hinder understanding of their partner's presence, while the shared control has the opposite effect. Adapting the findings from the study to the tabletop context, I designed and implemented Waves, a navigation technique that can allow independent control, while maintaining a connection between the collaborators' viewpoints. Waves works by allowing each person to make local changes that eventually propagate across the screen. The speed and area of effect of this propagation can then be adjusted to provide a variety of design alternatives that gain some of the benefits from the physical world of the combination of independent control with an awareness of a collaborator's perspective.

1.1 Motivation

There are many situations in which several collaborators need to maintain an awareness of each other and each other's work, yet perform activities individually within the shared workspace. For example, in emergency response situations, while the commanders need to focus on directing their own teams within their specific operational area, they would also need to be aware of the activities of teams managed by other emergency services commanders that are operating in the larger area in order to effectively collaborate in the overall relief effort. Similarly, in air traffic control settings, several air traffic controllers are each responsible for directing traffic inside a sector within a larger airspace. They must also have an awareness of the traffic situation in sectors outside of their own in order to smoothly handle aircraft travelling into their own sector from an adjacent sector. As such, there are two motivating factors behind this thesis: the need for interaction techniques that support collaboration and group situation awareness on the tabletop, and the lack of an existing technique that allows a variation of coupling (Tang et al., 2006) between different users —the degree to which collaborators interact with each other's work.

Studies in the past have looked into how the tabletop should support collaboration. For example, there are guidelines to help the design of tabletops systems for collaboration (Scott et al., 2003). These guidelines call for shared access to physical and digital objects, as well as to support simultaneous user actions. There are techniques in information visualization that enable manipulation (zoom, pan and rotate) of various parts of a contiguous piece of visual data (Mikulecky et al, 2011) while maintaining the connection between the areas of interest and their surrounding space. What is missing is the exploration of how these individual spaces in the shared screen on the tabletop can interact in a multiuser environment such that members of the group can perform individual tasks while maintaining an awareness of the perspective and progress of their peers.

As tabletops are being used for board games (Chang, 2012) and real-time strategy games (Tse et al., 2006), an interaction technique that addresses both personal navigation and group awareness will enable the multiuser applications that designers have envisioned.

1.2 Research Question and Objectives

The overall research question that is addressed in this thesis is:

How can we bridge the gap between the shared single-screen mode, which offers situational awareness, and the split-screen mode, which offers autonomy?

The three research objectives in this thesis are to:

- **Identify existing interaction techniques for multi-user shared screen applications.** A review of previous literature is conducted to find existing solutions to this problem. This work is presented in Chapter 2.
- **Explore two screen sharing technique used in video games to develop an understanding of how each affects interference between collaborators and situation awareness.** I performed an exploratory study to find differences between two popular collaborative navigation techniques used in video games, split screens and single shared screen, in situational awareness, interference between collaborators, and difficulties with automatic view adjustment. This work is presented in Chapter 3.
- **Design an interaction technique that bridges the gap between single screen and split screen in terms of group awareness and performance.** The major work in this thesis is the development of a multi-user interaction technique for the tabletop called Waves. The differing perspective between users propagate against one another to create group situation awareness between workspaces while maintaining individual perspectives. This work is presented in Chapter 4.

1.3 Scope

Group work involves the collaboration of multiple individuals towards the completion of a common goal. Groupware is a computer system that supports multiple users in this effort. Computer Supported Cooperative Work (CSCW) is the research discipline that studies the nature of cooperative work and the design of computer-based technology that supports it (Schmidt & Bannon, 1992). CSCW systems can be classified depending on two major characteristics: interaction and location (Rodden, 1991). Figure 1 shows the classification space and example activities. This thesis focuses on the study of multi-user interaction techniques on the digital tabletop. Collaboration in this setting is co-located and synchronous.

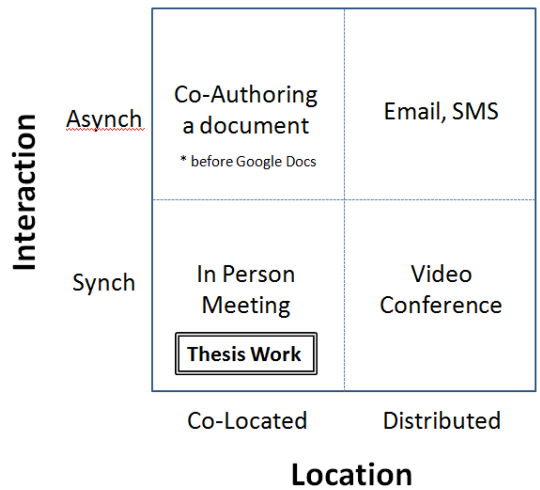

Figure 1: Classification space for CSCW (Rodden, 1991) and the focus of this thesis

This thesis will limit itself to co-located synchronous collaboration as defined in CSCW (Rodden, 1991). Specifically, this thesis concentrates on multi-user touch interaction on large horizontal interactive surfaces, and orientation of the surface is not explicitly considered. I assume that the starting state of the system is a display showing a single top-down perspective of a contiguous virtual space. The major contribution of the thesis will be to address the challenge of dividing the workspace between the users when competing navigation commands are received, showing the discontinuity of the divided workspace, the interaction between the divided workspace, and the ultimate reintegration of the divided spaces.

1.4 Thesis Contributions

In this thesis, I:

- Performed an exploratory study investigating the differences between two popular collaborative navigation techniques used in video games, split screens and a single shared screen. The study provided evidence to that there are differences in situational awareness, interference between collaborators, and difficulties with automatic view adjustment (Chapter 3).
- Designed and implemented Waves, a collaborative navigation technique for the tabletop based on guidelines formulated from the results of the exploratory study; Waves

simultaneously supports multiple personal workspaces, provides group workspace awareness, and mediates interference between workspaces (Chapter 4).

1.5 Thesis Organization

This thesis is presented in the following manner:

- **Chapter 1, Introduction** – presents the motivation and the research question of this thesis.
- **Chapter 2, Related Works** – shows existing research and how they relate to this work.
- **Chapter 3, Exploratory Study** – presents the summary of a preliminary study on the interaction methods used by two popular games in the PlayStation3 video game console.
- **Chapter 4, Design of Interaction Technique** – describes the design and rationale of a multi-user interaction technique for the digital tabletop that implements Waves, the proposed multiuser navigation technique of this thesis.
- **Chapter 5, Conclusion and Future Work**– summarizes the accomplishments of this thesis and how they relate to the completion of research objectives.

Chapter 2
Related Work

This thesis focuses on the development of an interaction technique that supports simultaneous navigation by multiple users around an interactive tabletop, informed by an exploratory study of navigation in multiplayer games. This related work section discusses existing work that has been done in the research areas of interactive tabletops, collaboration, and navigation in information visualization. Much research has already been done in each of these individual areas. However, it is the combination of these three areas that makes the topic of this research unique. The section follows an incremental approach which starts with the introduction of interactive tabletops. This will then be followed by how the interactive tabletop facilitates collaboration. Then finally, how collaboration is supported on the tabletop by information visualization techniques.

2.1 Interactive Tabletops

Digital tabletops are large horizontally-oriented interactive surfaces. Examples include the Microsoft Surface, DiamondTouch (Dietz & Leigh, 2001), and SMART Table (www.smarttech.com). One of the first interactive tabletops was Wellner's DigitalDesk (Wellner, 1991). Wellner combined what he described as the two desktops: the virtual desktop on the computer screen where electronic documents are read, and the physical desktop where paper documents are laid out horizontally. The DigitalDesk also experimented with direct touch interaction where virtual objects such as documents were manipulated using physical touch on screen as opposed to being manipulated indirectly by a mouse.

Benko et al. (2009) describe the benefits of horizontal interactive surfaces as "ease of use and intuitiveness due to multi-touch direct manipulation, quick learning time for novice users and support for collaborative work". While the DigitalDesk was built with individual work in mind, advances in multi-touch sensing made simultaneous input from multiple users possible. This further enhances the digital tabletop as a collaborative tool.

The recent surge in technology allows for the use of multiple hands and fingers by multiple people around the display (Dietz & Leigh, 2001; Han, 2005) , and research on digital tables has been largely focused on the design of multi-touch interaction techniques, its evaluation as a collaborative medium, and applications of the technology in specific domains. I will discuss this research in more detail in the following subsections.

2.1.1 Multi-touch Techniques

With Windows-Icons-Menu-Pointer (WIMP) graphical user interfaces (Bowman, 2004, p.91), users interact with the virtual environment indirectly through a physical device such as a mouse. The introduction of touch devices such as large interactive surfaces, smartphones and tablets in the post-WIMP world now allows users to directly touch and manipulate the virtual content on-screen. While the mouse is still effective for single point interaction, direct-touch inputs are found to be more effective and more preferred by users in bimanual tasks on surface devices such as the tabletop (Forlines et al., 2007). The multiple touch points of multi-touch interfaces afford additional degrees of freedom in 3D interaction and allow the user to perform navigation such as pan, zoom and rotation simultaneously (Westerman, 1999; Hancock et al., 2006).

Over the years, a number of multi-touch interaction techniques have been developed and proposed. The incorporation of physics such as friction, collision, and multiple touch points into touch interaction affords the user the ability to interact with the virtual world in a natural way that they are accustomed to in the real world (Wilson et al., 2008). An example of this is Sticky Tools developed by Hancock et al. (2009), an interaction paradigm that provides six degrees of freedom (DOF) of manipulation through force-based interaction. A specific challenge in the development of touch interaction techniques on a surface is the use of coordinates generated from *n* touches on a 2D surface for the navigation or manipulation of objects in a 3D virtual environment (Ortego et al., 2013). The question is how separate touches should work together to perform navigation commands such as pan, zoom, and rotate, or in the case of object manipulation, translation, rotation, and scaling in all 3 axes.

At the time of writing this thesis, the Rotate, Scale, Translate (RST), introduced by Westerman (1999), is the de-facto method for the interaction with 2D data. Adopted by popular applications like Google Maps for Android (maps.google.com), this one to two-touch technique works by keeping the content under the touch points sticky as the user's fingertips move. To perform navigation or manipulation, the user uses two fingers to twist around a centre to rotate, two fingers to pinch or stretch to scale, and one finger to drag to translate. As this technique utilizes up to two touches, it is limited to 4DOF. Hancock et al. (2007) presented shallow-depth 3D interaction using one-, two- and three-touch techniques that uses the displacement of up to three touch points over time to determine the rotation as well as the displacement in all three axes for 6DOF.

While research on multi-touch interaction focuses largely on the manipulation of objects, less work has explored navigation of 3D space on touch surfaces, especially how the above manipulation techniques can be used to support multiple users wishing to interact with the tabletop at the same time. Using the single-person 3D touch manipulation techniques discussed above as the basis, this thesis explores additional techniques that allow the simultaneous navigation by multiple users in the context of interactive tabletops.

2.1.2 Applications

Since its introduction more than 20 years ago (Wellner, 1991), the interactive tabletop has received keen interested from the Computer Supported Collaborative Work (CSCW) community. Its large horizontal form factor facilitates co-located collaborative work. People typically gather around the tabletop and use artifacts displayed on the tabletop to establish grounding (Clark & Brennan, 1991) in their communication. Over the years, studies were carried out to evaluate the effectiveness of the interactive tabletop in a variety of settings. Scott et al. (2010) investigated the use of a digital tabletop application in a naval command and control setting. Hancock et al. (2010) looked into how a tabletop could support sandtray therapy. Antle et al. (2011) created a digital tabletop game to explore collaborative learning. Common in these particular applications is the display and interaction of geospatial data from a bird's-eye-view perspective.

My particular research interest is the application of tabletop displays in scenarios where multiple people work together to make decisions, such as command-and-control or emergency response situations, when collaborative use of 2D data, such as a map with geographical information, can help provide situational awareness. An example of this is the investigation conducted by Scott et al. (2010) on the role of digital tabletops in a naval command-and-control environment. Figure 2 shows the top-down view of the prototype where it allowed multiple users around the table to interact with a windowed view of the area of interest as well as other data windows through a pen interface.

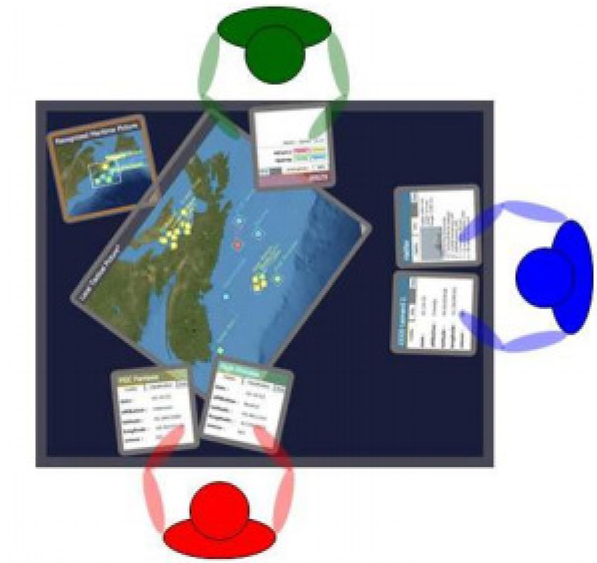

Figure 2: Application of table in naval command and control scenario (Scott et al., 2010, with permission)

Command and control involving the navigation of geospatial data is also found in entertainment. Real-time strategy (RTS) is a genre of video game that resembles in many ways the command and control environment. The player interacts with his/her units through a map interface with a bird's eye view of the area of operation. In order to effectively manage his/her army, the player must navigate to different parts of the map to give commands to different units. Tse et al. (2006) ported Warcraft III, a popular RTS game by Blizzard Entertainment, onto the DiamondTouch and used a software wrapper to enable multiple players to collaboratively command the same army. While this concept shows how multiple players can use touch gestures simultaneously to command units on the map (i.e., simultaneously give commands to units on the screen), and collaborate when the view is stationary, it does not provide individual workspaces for each user. This thesis aims to provide these workspaces.

More recently, at Electronic Entertainment Exposition 2009 (www.e3expo.com), an annual video game conference where game developers demonstrate their new games for the year, game developer Ubisoft (www.ubisoft.com) showed R.U.S.E., an RTS game originally built for mouse and keyboard interaction, running on an interactive tabletop. The demonstrator used touch to both navigate the map and manipulate units on the screen[1]. In the video, it can be seen that the presenter was able to use

[1] https://www.youtube.com/watch?v=vf9csiGVqgs

touch gesture to navigate the battlefield and issue commands to units. However, despite the tabletop form factor, it lacks the support for collaboration. Only one user can interact with the game at any one time.

While all of the applications above provide a collaborative environment in which multiple users can interact with the workspace, none of them provides the ability for users to navigate independently of one another in the same shared space. This thesis aims to address this issue by proposing a design that provides each collaborator with their own personal, but connected, workspace.

2.2 Collaborating with the Tabletop

Other research has explored collaboration around a table, and this work builds upon this literature by creating a technique that bridges collaboration and navigation. The Workspace Awareness Framework for Real-Time Groupware (Gutwin & Greenberg, 2002) was created to help the development of software for distributed groups. The underlying concepts of workspace awareness also apply to co-located settings, such as tabletops. Workspace awareness (WA) is defined as the-up-to-the-moment understanding of how other people are interacting with a shared workspace (Gutwin & Greenberg, 1996, p.208-209). While situation awareness (SA) refers to a person's awareness of what is happening around him/her (Endsley, 1995), WA is a specialized kind of SA that involves the state of the workspace and the domain task (Gutwin & Greenbergy, 2002, p.417). In a collaborative situation, WA further includes the collaboration and the state of the shared workspace. Workspace awareness supports collaboration by reducing effort, increasing efficiency, and reducing errors for the activities of collaboration. We hypothesize that an interaction technique that supports workspace awareness will better serve the user. This thesis intends to develop a multiuser interaction technique that facilitates WA.

One of Scott et al.'s (2003) system guidelines for tabletop displays calls for designs to support simultaneous user interaction. This is especially important, as the large size of the tabletop affords simultaneous use by multiple users. Research on territoriality (Scott et al. 2004) showed that tabletop users keep separate spaces for personal work and group work. Social conventions discourage users from intruding on another user's personal workspace. This is related to the protection activity, a subtask in the Mechanics of Collaboration (Gutwin & Greenberg, 2000), which break down the activities of collaboration into: communication, coordination, planning, monitoring, assistance, and protection. They suggest that users of groupware should have ways to safeguard their own work from being changed by others using the groupware. Scott et al.'s (2004) territoriality study involves manipulation of pieces of discrete data/objects in a shared space. Territoriality is shown through placement of mobile virtual objects. Collaborators all share the perspective of the workspace while at the same time manually maintained their border through placement of virtual objects. This study shows that people around the tabletop implicitly created personal territories. This thesis builds on the idea of territoriality on the tabletop by using the system to allow individual users to manipulate their own view without interfering those of others, thus allowing territoriality.

2.3 Information Visualization Techniques

Information visualization (InfoVis) is a branch of research that studies the presentation of data. A known challenge in HCI is displaying a large amount of available information in the limited screen space. Despite the large size of the tabletop's display space, it is not always possible to place every piece of data on it. In the case of collaboration, not only should the tabletop support the exploration (in the form of panning, zooming, and rotating) of the information as a group, but also that of individual users around the table. The connection between InfoVis techniques and collaboration was evident in Tang et al.'s study (2006) where the effects of InfoVis techniques on collaborative coupling, namely, the degree to which the participants interact with each other, in the exploration of fixed spatial data were investigated. The study's goal was to find out whether different visualization techniques would change the degree to which partners collaborate with each other. Participants were given a pathfinding task on a large map and found that different visualization techniques affect collaboration styles. For example, collaborators that were provided with localized lenses spent more time working on the problem individually. Those who used filters, graphical overlays that cover that entire map, spent more time working in a more tightly coupled manner. Instead of looking at the amount of collaboration, this thesis will instead focus on InfoVis techniques that improve the quality of collaboration by reducing interference between user workspaces and facilitating workspace awareness between collaborators. In the following subsections, I discuss specific InfoVis techniques that are relevant to addressing the collaborative exploration of spatial data on the tabletop.

2.3.1 Detail-in-context InfoVis Techniques

When exploring spatial data such as maps and large graphs, Shneiderman's Information-Seeking Mantra can be applied: overview, zoom and filter, details-on-demand, relate (Shneiderman, 1996). It can be useful to see an *overview* of the entire data set, to gain an understanding at a global level, such as an entire route on a map, or a global view of how nodes are connected in a large graph. At the same time, it may also be necessary to drill down or *zoom* to specific areas of interest and see the data in finer detail. While in this *detail* view, some sort of connection between this view and the larger global view provides the user with the awareness of how this detail *relates* to its surrounding larger data set. This particular kind of information display where users can see details within a larger set of data is called detail-in-context presentation (Spence & Apperley, 1982; Furnas, 1986; Sarkar & Brown, 1994).

Spence and Apperley (1982) were some of the first researchers to present a detail-in-context display with their Bifocal Display. They describe a segmented touch screen in which the left most pane shows a "demagnified" graphical view of the contents of a database. The user could see the details of the contents by moving a selector onto an item. The contents of the item are then expanded onto the second pane. Furnas (1986) presented the generalized fisheyes which showed the current area of focus and its immediate surroundings in detail. For things further away, only major "landmarks" are shown. The example in that paper showed a generalized text-based example where the code in the lines surrounding the cursor is shown in detail. Further away from the cursor, only lines with major keywords such as "if", "else", or "while" are shown.

Sarkar and Brown (1994) then adapted Furnas' idea for graphs and the "graphical fisheyes". The example in the paper shows how this detail-in-context method enlarges a particular item in a network graph and displays the items connected by surrounding edges in decreasing magnification. Other methods make use of 3D manipulation and perspective viewing to show the connection between the area in focus and the surrounding area. Perspective Wall (Mackinlay et al., 1991), Document Lens (Robertson & Mackinlay, 1993) and Pliable Surfaces (Carpendale et al., 1995) all show detail by bringing the area of focus closer while leaving the surrounding area further away. The Manhattan Lens also follows the same idea where the area of interest appears protruded from the map's surface (Carpendale & Montagnese, 2001).

The many detail-in-context visualization techniques proposed were seen as distinct methods. Furnas and Bederson's (1995) space-scale diagram provided an analytic framework to visualize and characterize different techniques. Drawing on the space-scale diagram framework, Carpendale & Montagnese's (2001) Elastic Presentation Framework (EPF) presented a unified mathematical description of both distortion and non-distortion methods. The EPF describes each technique based on differences in the characteristics in transition between the area of focus and the surrounding map. The EPF made it possible to incorporate a number of techniques in the same interface.

A highly emphasized part of EPF is the idea of resilient deformability. It is the display space's ability to be stretched and the ability to return to its original shape. In collaboration, it is likely that personal workspace needs to be created and reintegrated with the larger workspace. This thesis will leverage the detail-in-context visualization techniques and workspace deformability discussed above to facilitate the creation and re-integration of personal views that will result from simultaneous navigation of multiple users.

2.3.2 Multi-Focus Interaction

To facilitate collaboration, the tabletop could integrate multiple detail-in-context views such that more than one user can explore the content being displayed. A specific technique that addresses this is the Mélange space folding method for multi-focus interaction (Elmqvist et al., 2008). This method was designed to have four requirements in mind.

1. *Guaranteed focus visibility*. Multiple areas of focus must be simultaneously visible.
2. *Surrounding context visibility*. The area surrounding the areas of focus should be visible.
3. *Intervening context awareness*. The space between focus regions should be shown to give a frame of reference.
4. *Distance awareness.* Users should be able to tell how far the focus regions are apart.

To accommodate multiple focus areas in a single display, the Mélange technique distorts the intervening regions between the two view areas in the form of folded space. The depth of the folded space gives the notion of distance between the two spaces. The compact folded region also provides context between the view areas. While this method supports multiple focus regions and has the potential to support collaboration, Elmqvist et al. only conducted a user study on single-user tasks. The search and estimation tasks call for dual focus regions. However, at any point, only one region was user controlled while the other was fixed to a primary focus view. It was unclear how the technique would behave when both views are manipulated simultaneously in a collaborative setting.

It can be argued that even though Mélange was created with single user, multiple-focus scenario in mind, the same requirements that it set out to satisfy, namely providing guaranteed focus visibility, surrounding context visibility, intervening context awareness, and distance awareness, are also applicable to multiple-user, multiple focus scenarios in collaborative navigation. As such, a Melange-like distortion between workspaces is incorporated in the design of the proposed interaction technique proposed in this thesis to provide awareness between collaborators of the relations between their workspaces.

2.3.3 Incorporation of Physical Metaphor

Mélange borrowed from the physical metaphor of folding maps. For example, in the real world, when a person needs to see two different areas on a paper map at the same time, the unneeded parts can be folded away. Another technique that borrowed heavily from physical metaphors is the Information Cloth (Mikulecky et al., 2011). Here data such as a geographical map is rendered as the texture of a malleable cloth that can be stretched, pinched, and folded. To create focus areas, various resizable shapes can be placed below the cloth. As the cloth is placed on top of the shapes, the physics engine deforms the cloth along with its data texture according to the shapes below. Various distortion techniques like Perspective Wall (Figure 3), Graphical Fisheye Lenses (Figure 4) and, to a certain degree, Manhattan Lenses (Figure 5) can be created by placing the appropriate object under the cloth.

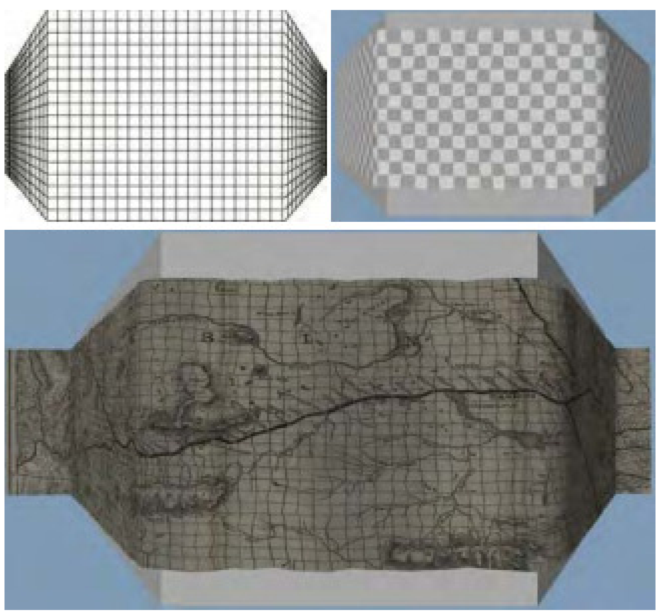

Figure 3: Information cloth implementation of Perspective Wall (Mikulecky et al., 2011, with permission)

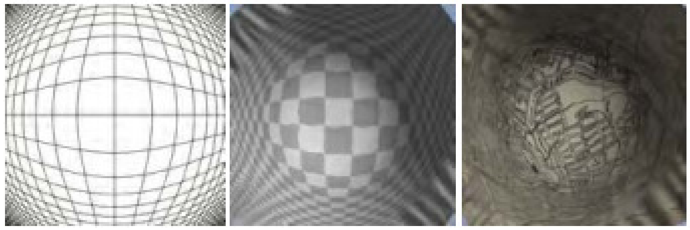

Figure 4: Information cloth implementation of Graphical Fisheye (Mikulecky et al., 2011, with permission)

Figure 5: Information cloth implementation of Manhattan Lens (Mikulecky et al., 2011, with permission)

Other distortions like the fold in the Mélange technique can also be reproduced by draping the cloth over two shapes, leaving the middle hanging.

Figure 6: Information cloth implementation of Malange (Mikulecky et al., 2011, with permission)

The Information Cloth technique has the advantage of allowing multiple collaborators to interact with the data simultaneously. The cloth's physical behaviour naturally addresses the issues of continuity between focus areas and provides context. When different focus areas are manipulated simultaneously, the cloth simply reacts by stretching or folding as defined by the physics engine. In many ways, this information visualization technique is very close to solving the issue for multiple focus areas in collaborative settings.

While this work uses a cloth-like interaction technique similar to the one described in this thesis, their work does not have a mechanism to dynamically separate individual workspaces. For example, when multiple users simultaneously pan away in different directions, the entire table view becomes stretched and deformed as seen in Figure 7.

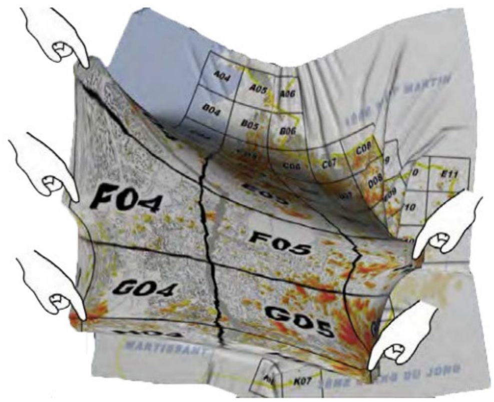

Figure 7: Distortion affects the entire area in Information Cloth (Mikulecky et al., 2011, with permission)

The distortion can interfere with individual exploration. To address this, the interaction technique proposed by this thesis provides an undistorted view within each user's workspace by limiting the distortion to the borders between workspaces. In addition, these borders adjust dynamically according to the movements within each individual workspace.

2.4 Summary

Much existing research has been done in the areas of multi-touch 3D navigation, collaboration using the tabletop, and information visualization. Important insights from each research area are:

- 3D multi-touch manipulation techniques such as Westerman's RST (1999) and Hancock et al.'s shallow-depth 3D interaction (2007) allow multiple users around the tabletop to simultaneously manipulate objects on a shared workspace.
- Taking the subtask for protection from Gutwin & Greenberg's Mechanics of Collaboration (2000) and applying it to Scott et al.'s concept of territoriality (2004) for the tabletop, suggests that a collaborative navigation technique should include affordances for creating and isolating personal workspaces within the shared space.
- The idea of malleable workspaces like the information cloth (Mikulecky et al., 2011) and multi-focus interaction techniques like Mélange (Elmqvist et al., 2008) can be leveraged in the interaction between individual work spaces as a mechanism to facilitate workspace awareness by representing spatial relations between each individual's view.

Moving forward, this thesis will build on the work above with the lessons learned from the exploratory study of popular screen-sharing techniques already in use in video games to inform the design and implementation of a novel collaborative navigation technique for tabletop displays.

Chapter 3
Exploratory Study

As very little research on collaborative navigation on the tabletop exists, my research effort began with an exploratory study on existing collaborative techniques in console video games featuring single screen cooperative play. While it could be argued that the players do not interact directly with the screen using touches as people do when using they use tabletops, the underlying screen sharing issues are very similar. In both cases, two collaborators need to navigate a virtual space while sharing a common display. Two games that employed very different screen-sharing techniques were selected. The first is the split screen technique used by Call of Duty: BlackOps, a first-person shooter video game. Players in this game mode simultaneously navigate the game's virtual environment in their own halves of the split view. This is similar to each player having his/her own personal workspace. The second technique is LEGO Star War's shared-view. In this adventure game, two players navigate the game's virtual world via a single common perspective that adjusts continuously to keep both players' avatars in view.

3.1 Motivation for Studying Collaboration in Video Games

There are two reasons why studying collaboration in video games is a good idea. The first reason was quality. The video game industry is competitive. Vast amounts of effort and resources have been spent by game studios to make their games fun and the user interfaces efficient. The most popular games were ones that had been refined over many iterations and had large fan bases. The two games chosen represented the state-of-the-art the gaming industry offered at the time of the study. The second reason was availability. Video games were one of the places where matching applications were readily available. While the concept of collaborative navigation on the table top was relatively rare, the same concept could be found in multiple console video game titles.

Two very popular video games had in them a built-in collaborative navigation component that would be examined in this study. The first game was called Call of Duty (COD) Black Ops. It was a first-person shooter in which players navigated through the game world in a first-person perspective. The second game was called LEGO Star Wars. Players here navigated the game world looking down on their avatars in a third person perspective. In both of these games, players had to collaborate with their partners in order to complete game objectives. In COD Black Ops, players were trapped in a building and must work together to defend against wave after wave of invading zombies. In LEGO

Star Wars, players took on different characters of the Star Wars movie trilogy and worked together to advance through the storylines of the movie.

More important was the fact that the two games sported different screen sharing techniques for collaborative play. Very much like Tang et al.'s study (2006) where different InfoVis techniques had an effect on the participants' level of collaboration, this exploratory study aimed to find out how the different screen sharing modes affected how the players navigate collaboratively, and the interaction with their partners. To the author's knowledge, there are no existing techniques for collaboratively navigating on a tabletop display. Thus, in order to inform the design of a collaborative tabletop navigation technique, I look to existing collaborative navigation techniques in video games, of which there are several. The expected end result of this exploratory study was a set of implications that will inform the design of a collaborative navigation technique for the tabletop.

3.2 Methodology Overview

Three research techniques were used in this study: contextual inquiry (Holtzblatt et al., 1993), critical incident technique (Flanagan, 1954), and situation awareness analysis (Endsley, 1988). Figure 1The rationale and the application of specific techniques will be discussed in respective sections below.

Figure 8 illustrates how the methodologies were applied and the sequence of associated research tasks. A total of eight people in groups of two were recruited in this study. Participants were male, between the ages of 25 and 30. Six of the eight participants were graduate students at the time of the study. The two others were adults with undergraduate degrees and were working full-time in the technology industry.

The study began with the observation portion of the contextual inquiry. Each group was allowed to play the first game for up to 20 minutes. At some point during the 20-minute session, the game play was suspended and a blank test in which the game display was abruptly shut off and the players were asked to draw the details of the environment in which they were navigating was conducted to test the participants' situation awareness. After the game play, an open-ended interview was conducted as part of the interview portion of the contextual inquiry. The participants would also be asked to recall any memorable moments in the game. This was to gather additional data for the critical incident technique. This process would be repeated for the second game. The order in which the groups played the games were counterbalanced between groups.

At the end of the study, over six hours of video was recorded. The data were then reviewed and compiled. As well, contextual inquiry models were constructed to determine any breakdowns during the observation. Specific issues brought up by the participants as critical incidents are analyzed in more details. Design implications were then drawn from the analysis of the data collected.

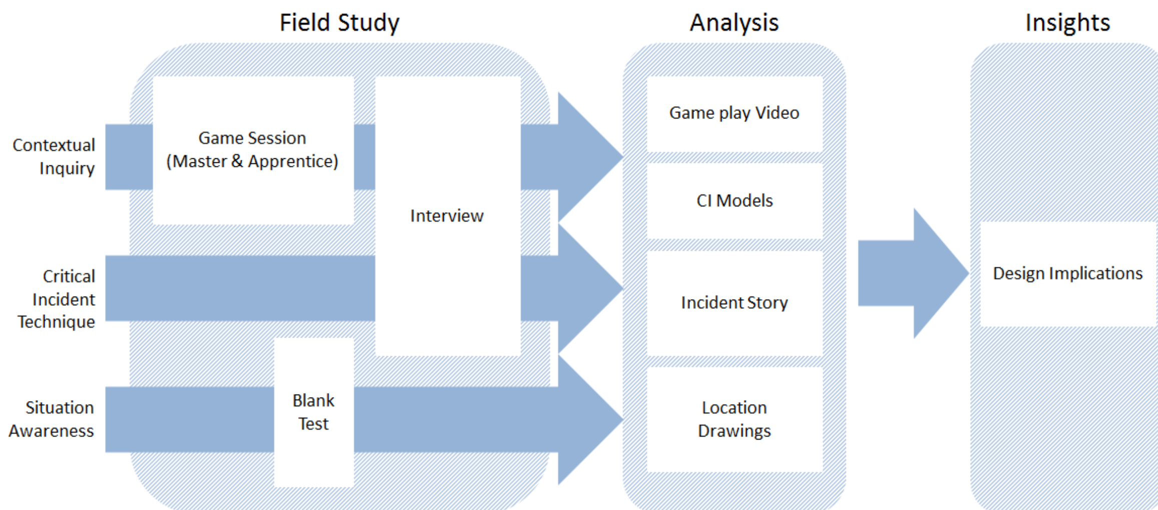

Figure 8: Exploratory study overview

3.3 Contextual Inquiry

Contextual inquiry (CI) was chosen to jump start this exploratory study because its master/apprentice model (Beyer & Holtzblatt, 1997) provided the researcher (the apprentice) with the much needed exposure to the users' (master) environment. As video gaming was largely an experiential activity, being there to observe and interact with the users first hand during game play served as a good starting point for the researcher to begin learning about the users' tasks.

3.3.1 Method

According to Beyer & Holtzblatt (1997), CI's core principles are context, partnership, interpretation, and focus. These principles were employed to provide systematic mechanisms with which to study the work task. The following explains how each of these principles was addressed in this study.

3.3.1.1 Context

This principle called for the research to be physically present in the users' native work environment. This was so that the researcher could be immersed in the richness of the user environment to gather what Bayer called on-going experience rather than summary experience, and concrete data instead of abstract data (Beyer & Holtzblatt, 1997).

To address this, the researcher was physically present during all of the user game plays. Due to logistical limitations, it was not always possible to visit each participant in their native play environment. As such, the researcher initially conducted observations and interviews at a pair of users' home where they normally play video games. Special attention was paid to the placement of the following: display, console, seating arrangements. A physical arrangement similar to the living room was then reproduced at the lab where subsequent participants were invited to be observed and interviewed. Participants were free to rearrange the setting as needed.

3.3.1.2 Partnership

This principle called for the adherence to the master/apprentice relationship between the researcher and the participants during the inquiry.

To do this, the researcher made an attempt to maintain a collaborative partnership by avoiding the use of a questionnaire or using technical jargon when asking questions. This was to avoid falling into the role of the interviewer/interviewee and the expert/novice.

3.3.1.3 Interpretation

This principle called for the researcher to confirm his interpretation of the on-going action with the participants. As the participants were engaged in the activity, they could directly use the knowledge in their working memory at the time to fine tune the researcher's understanding.

To do this, the researcher in this study employed the watch and probe technique as described by Bayer & Holtzblatt (1997). Probing questions or statements were posted to the participants as the action occurred to encourage explanations. An example of this occurred in the following dialog between the researcher and Participant 1, an expert COD Black Ops player.

P1 was holding his aim at an empty space.

Researcher: Hey P1, what's happening down there?

P1: Just waiting.

Researcher was not sure what he was waiting for at this point. Then a zombie materialized in the center of P1's crosshairs and P1 promptly shot him.

Researcher: Oh! You wanted to get a good shot on him as they come in!

P1: Yep!

Another zombie materialized again in the same spot and P1's quickly shot him.

P1: Ha ha, you like that? Another headshot.

3.3.1.4 Focus

This principle called for the researcher to allow the users to talk about things that they found were important. At the same time, since this was an exploratory study on the subject of 3D navigation, the researcher tried to place emphasis on collaboration and navigational matters. This was done by first interpreting correctly what the users were saying and what they, the masters, found important. If the matter was related to the mentioned area of interest, the researcher would then steer the conversation in that direction. The following is an example dialog in which this happened. P1 and P2 were both expert gamers, and were talking to the researcher about the PlayStation 3 controls in COD Black Ops. The researcher noticed something interesting in the dialog and steered the conversation in that direction.

P2: This is nice [holding up the game pad], because you have all your buttons at your fingertips.

P1 looked intently at his game pad

P1: How do you strafe? Because you aren't using these. [P1 holds the gamepad up so P2 can see and wiggles the analog joysticks]

P2: You use these. [P2 wiggle the left analog joystick on his game pad]

P1: So you can't look around on this?

P2: Wait, sorry, hold on…

At this point, the researcher suspected that the players didn't consciously remember specific controls.

Researcher: Wait, can you guys tell me what the controls are?

P1 and P2 struggled and finally pieced it together. They then went on to explain that they didn't explicitly remember what the controls were. It was just "instinct".

3.3.2 Findings

CI models were constructed using the data gathered during the observation phase. The physical model shows the settings under which the observations were carried out. The artifact model shows the virtual environment in each of the screen sharing modes. The flow model shows the communication between the participants in each group.

3.3.2.1 Physical Model

Figure 9 shows the physical model with overlays identifying the physical components in the environment over the photographs taken of the environment in which the gaming sessions occurred. On the right, the figure showed the native environment in which one of the participants regularly played video games. In this case, it was the participant's living room. Items of interest were highlighted. The white dashed lines provide depth references of each item in the photos so that the relative position to each other was clearer.

To begin, located at the back of the room was the large screen display on which video and audio of the game were emitted. In front of it was the PlayStation 3 (blue square), the video game console on which both games were played. Moving towards the front, between the console and the user was the wireless controller (four-way directional icons) via which the user provide inputs to the game. The controller also provided vibrational feedback to the user. The users (circles) were observed to be sitting on coaches (clouds) facing the screen with the controllers in front of them. Two pieces of non-native equipment were introduced into the environment by the researcher. A camera (triangle) was placed in close proximity to the screen facing the users in order to capture the users' physical actions during the play. A laptop equipped with an integrated webcam was placed in front of the game display in order to capture the users' actions in the virtual world.

Due to logistical reasons, two of the groups were invited to play at the researcher's lab space instead. An effort was made to ensure that all of the relevant components observed during previous sessions at participants' native environment were present at the lab. They were also placed in very similar sequential order starting with the screen, then the console, the controllers, and finally some form of comfortable seating.

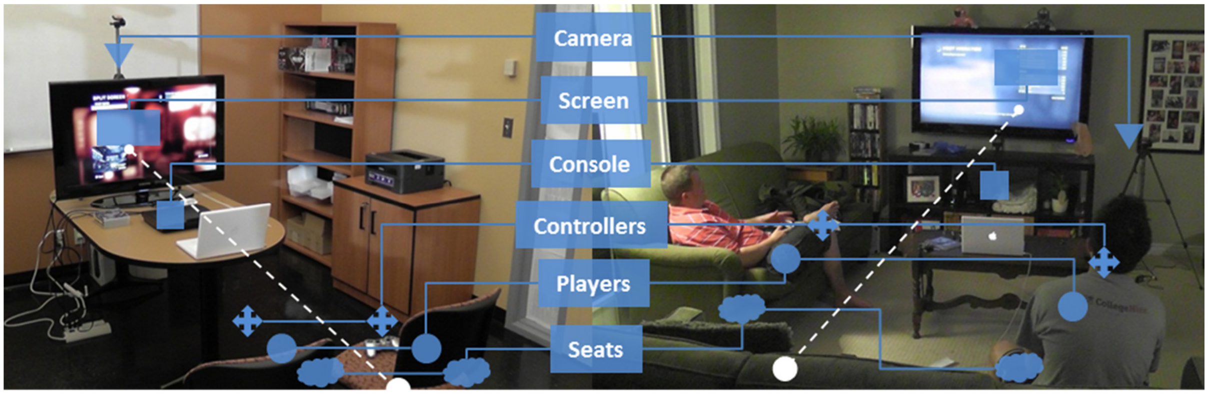

Figure 9: The video gaming environment at the participants' living room (right) and the setup at the lab (left)

While not explicitly told to do so, the participants were not kept from moving the items (or to move about on their own) in either their native environment or the lab. Participants were not observed to be very mobile over the course of the gaming session. Movements were observed to be limited to their heads and upper bodies.

During the follow-on interview, the participants were asked if they had any comments regarding the physical environment. There appeared to be a preference for large displays and comfortable seating. No apparent breakdowns were observed in this model.

3.3.2.2 Flow model

CI's flow model was used in the following to represent the information flow between players. Normally, each flow model diagram is constructed in the perspective of a single user. Due to the collaborative nature of the task being studied, it was modified to include the player and his partner (labeled Player 1 and Player 2). Two flow models were constructed, one for split screen mode, and one for single screen mode.

3.3.2.2.1 Split Screen Mode (Call of Duty Black Ops)

The model in Figure 10 uses colour codes to show a composite view of the interaction between players of varying expertise. The colour red represented the information flow between expert players, amber represented that between intermediate players, and green represents that between the intermediate and the beginner player. Immediately apparent from the start was that, between expert players, the communication was almost always two-way. The amount of information flow was very symmetrical. This was represented by the dual arrow connectors between expert players. For each piece of information or comment voiced by one player, the other expert player would either respond with another piece of related information or with an acknowledgement. The following was an example exchange between two expert players during Call of Duty Black Ops.

P1 was at the top of a long stair case, fighting off enemies that were coming up.

P1: There is three coming up the stairs…

P2: That’s not good, P1!

P2 started running up a different staircase to get to P1.

P1: Are you upstairs yet? Ah… I can’t get headshots on them.

P2: I am coming up on the other side. HEY, guys RIGHT behind you!

P1: Ack! Running.

P2 started shooting the enemies behind P1 as P1 started running for cover.

P2: Got your back! I got them. I got them. You got a guy chasing you still.

Also being illustrated in the diagram was the flow of information between Intermediate-intermediate (amber) and intermediate-beginner (green) players. They tended to employ more one-way communication and the flow was not necessarily symmetrical. This was illustrated by the single direction arrows. Figure 10 further highlighted specific types of information that flowed between players. The model divided the specific types of exchanges into those that occurred pre-game and in-game.

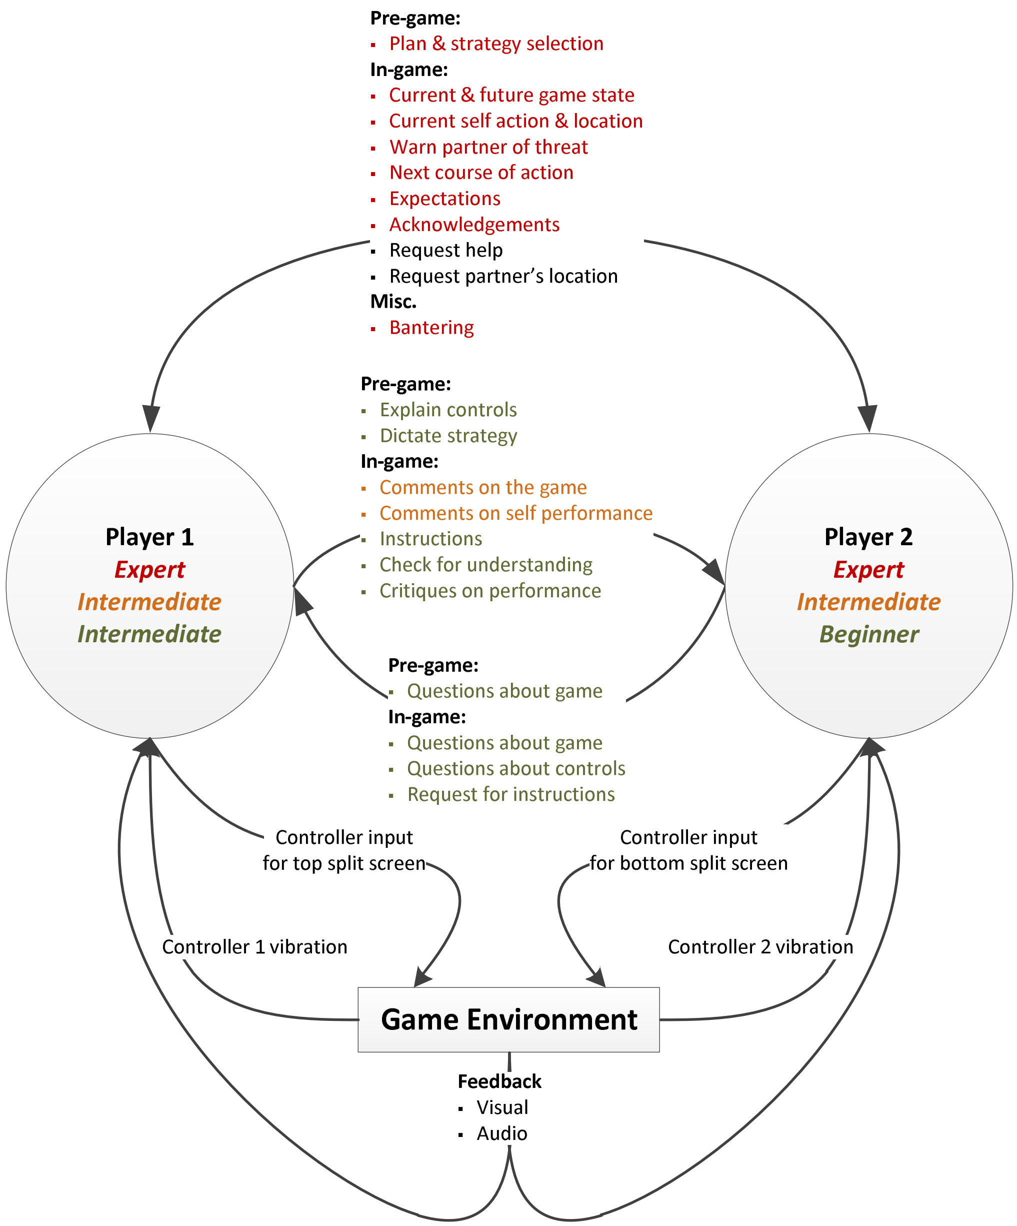

Figure 10: Flow model showing the interaction between players with different skills (represented by different colours) and the game environment in COD Black Ops.

3.3.2.2.1.1 Pre-game

The difference in information exchange began early. Expert players were observed to have begun their planning before entering the game environment (while the game was loading). They would bring up a set of possible strategies, discuss them, and then choose one. Such planning or pre-game exchange did not occur between the intermediate-intermediate players. The Intermediate-beginner pair did have a pre-game exchange. The intermediate player began by reminding the beginner player of the buttons on the controller. The intermediate player also picked a plan and told the beginner player that "I will let you know what to do."

3.3.2.2.1.2 In-game

Similar to pre-game planning, there was also a significant difference in the communication flow between the pairs of players of various expertise levels.

During the actual game play, expert players were constantly exchanging with each other the current game state, where they were, and what they were doing. These were listed as the top two types of communication between expert players in the flow diagram (Figure 10). The information was always given voluntarily, without being prompted by the partner. For instance, one of the players would say, "I am upstairs." His partner would immediately reply, "I am downstairs."

The exchanges also exhibited a degree of least collaborative effort (Clark & Brennan, 1991). It was a principle in communication grounding that stated participants in a conversation would do the least work possible to achieve mutual understanding. The following exchange was one of the examples.

P1: You watching the… uh… [*pause*]
P1 became occupied with an enemy who just appeared.

P2: Main door?

P1: Yeah.

P2: I'm out.

P1: Out? [*pause*] Oh you left! Ah, vanity!

P2: Ha yes!

In the exchange above, P1 stopped short of finishing the sentence due to time pressure. He also figured from experience P2 would know what needed to be covered. This assumption was correct judging from P2's response. P2 also exercised least collaborative effort when he simply said he was out and left P1 to figure out what he meant. P1, puzzled at first, quickly figured out that "out" meant outside the current room they were in, where there were better weapons and more enemies outside that would give P2 more points. Hence he made the "vanity" comment, to which P2 laughed. It was likely that their ability to communicate using what would appear to a newcomer as disjointed exchanges was supported by their prior experience and history together as well as visual clues provided by the game.

While the expert players did not always have their partner in view, when they spotted possible threats moving towards their partner's approximate location, they always passed that information along. In between waves of enemies, when it was not busy, the expert players would turn their conversation to planning their next move (i.e, when to unlock the next level and move onto the next location). Silence was rare.

Two-way communication that was common to players of all levels was requesting help and requesting their partner's location (both in black text in the diagram). The latter request was often a response to the former. This was interesting, since players still requested this information verbally when visual aid showing the relative location of their partner was available on the display.

There was significantly less traffic between intermediate players during the game. They made comments about the game such as, "I wish the shotgun comes with more than two shells," and "Oh, that didn't work." These comments did not appear to be directed at anyone in particular. A response was rarely given by the partner. In between rounds, where expert players would talk about the game in great details, breaking down what happened and what went wrong, intermediate players simply had much less to say.

The information flow between the intermediate and beginner player was very much like that of a teacher and a student. The intermediate player mostly gave instructions and told the beginner player if he was performing the right actions. The beginner in turn mostly asked about what he should do next, what controls perform specific action, and general questions about the game. The following were two observed exchanges between P3 (intermediate player) and P4 (beginner player).

P4: How do you change weapons?

P3: Triangle [button]. But you don't want to, because the zombies right now takes six or seven shots to kill them.

Here, the beginner started asking a question. The intermediate player provided the answer, as well as advice that he should keep his current weapon rather than switching to one that's less powerful. The following example dealt with directions.

P3: Okay, let's go outside. Do you see me? Okay, I am coming back in. Do you see me now? [*P3 could see P4 in his view*] Right behind you, right behind you.

P4: Okay… [*P4 turned around to follow P3*]

P3: So it's to your left. See the light *[shining in through the door from the outside*]? Just follow the light… just follow me…

Note that P3 provided redundant instructions in the last comment. He told P4 to follow the light, as well as to follow him. When asked why he did that, P3 commented that he was in front of P4 (in the game) and couldn't tell if P4 was following him or what he was seeing. He decided to give both instructions just in case. This was interesting. In split screen mode, P3 only had to look down to see what P4 was looking at. Despite this, all the players, expert, intermediate and beginners alike, stated that they rarely look at the bottom screen.

3.3.2.2.2 Single Screen Mode (LEGO Star Wars)

The following (Figure 11) shows the flow model for LEGO Star Wars. Since none of the participants recruited for this study had ever played the game, skill level was not a factor in this model. As such, the pair of participants was simply labeled as Player 1 and Player 2. There was also a difference in the players' interaction with the game environment between this and the previous flow model. Referring back to the flow model for Call of Duty Black Ops' split screen mode (Figure 10), even though both players had equal access to the game environment, the controller input from each player affected only the viewing perspective of the split screen that was under his control. Here, since there was only a single viewing perspective, action from either player affected the entire visual display (illustrated by both arrows joining in the controller input). The details of the controls will be discussed later in with the artifact model.

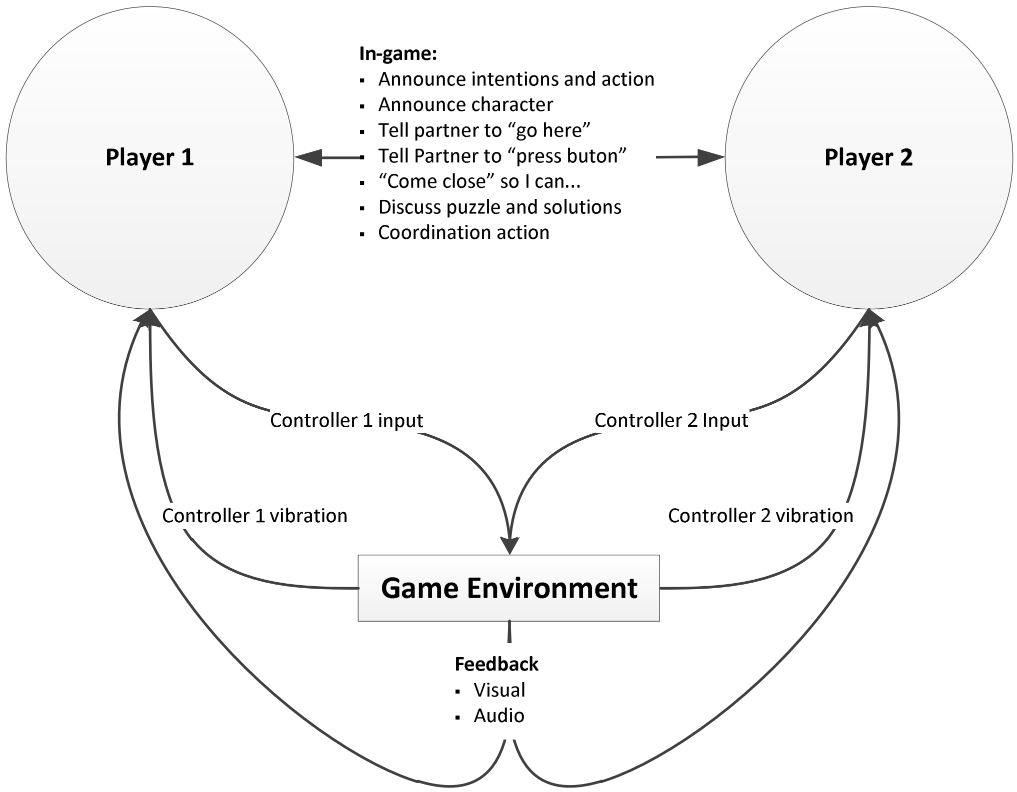

Figure 11: Flow model showing the interaction between players and the game environment (LEGO Star Wars)

The interaction between the participants while playing LEGO Star Wars was different from when they were playing Call of Duty Black Ops. The LEGO game play required much closer cooperation between the players. In many case, in order to advance to the next level of the game, both players must each press a button or toggle a lever at different parts of the map at the same time to unlock a door or a gate. This promoted heavy two-way communication between the two players. Information used to coordinate flowed continuously between players through the game as they coordinated with each other to solve puzzles. Figure 12 showed an example of such collaboration. Here, the two players were trying to get to the control panel at the top via the yellow elevators. These elevators could only be operated by a player's telekinesis ability. The catch here was that a passenger could not operate his own elevator. Here, the players solved this puzzle by each getting on an elevator, and then using his telekinesis ability (shown by the respective blue and green aura) to lift his partner's elevator. Exchanges that facilitate this type of coordination such as those described in the model represented the bulk of the communication traffic.

Figure 12: Two players collaborating in LEGO Star Wars

3.3.2.3 Artifact Model

According to Bayer & Holtzblatt (1997), artifacts are things that people create, use, and modify during the course of their task. CI's artifact model captured the details of these artifacts and used them to tell a story on how people worked. While the users did not create anything during their gaming session, they were actively using the feedback provided by the display and modifying the 3D virtual environment through the PlayStation 3 controller. The following artifact models explored the two different display modes and the controller.

3.3.2.3.1 Split Screen Mode

In Call of Duty Black Ops, the users were presented with a split screen that provided the display for Player 1 on the top and that for Player 2 on the bottom. The artifact model is illustrated in Figure 13. In Part 1 of the model (Figure 13), the display showed the virtual environment through an eye-level first person perspective of the player's virtual presence in the environment. A hand holding a weapon could be seen extended from the bottom of the player's screen to show the direction of fire and to provide the user with a reference for depth perception.

The display also sported a virtual heads-up display (HUD) that provided support information related to the virtual world by overlaying it onto the screen in real time. For example, player information such as weapon ammunition and accumulated points were shown to the player on the bottom right. In the center of the player's view (not labeled) was a small crosshair that showed where a player's shot would land.

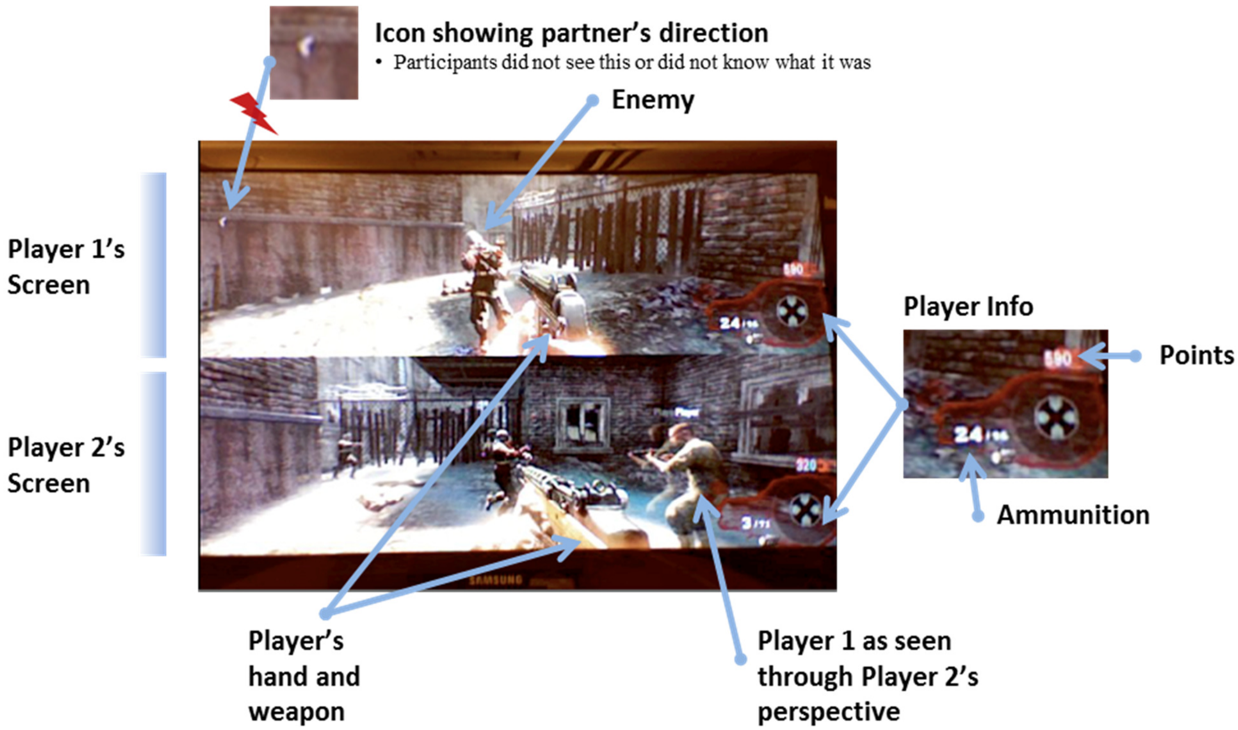

Figure 13: Artifact model of Call of Duty Black Ops' split screen display mode (Part 1)

To support situation awareness, the HUD employed arrow icons (top left in Figure 13) to indicate the relative direction of the player's partner. In Figure 13, looking at Player 2's screen (bottom split screen), it could be seen that Player 1's avatar (with a white label above its head) was standing in Player 2's view. This meant that from a top down point of view, Player 2 was standing to Player 1's left, and a little in front. To represent this spatial relationship, Player 1's HUD displayed an arrow icon that was pointed at the screen's 10 o'clock position. In this convention, the 12 o'clock position was directly forward. As an example, if the partner were directly behind the player, the arrow would be pointed at the six o'clock position (drawn at the bottom where the "hand and gun" was, and pointed towards the ground). Figure 14 showed where the icon would appear in a player's split screen depending on the partner's relative position (gray circles) to a player facing 12 o'clock (blue figure in the middle).

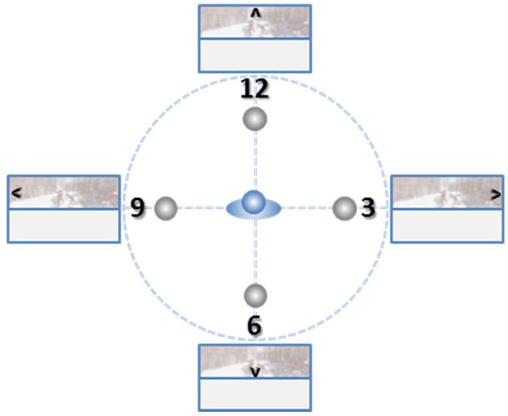

Figure 14: Icons indicating the various relative directions of Player 1's partner

A similar technique was used by the game to show the direction of enemy attacks. Figure 15 showed how the HUD used a red circular arc to indicate that the player was being attacked from a certain direction. In the top split screen in Figure 15, the lower left arc indicated that Player 1 was being attacked from the 8 o'clock direction. Player 2, whose screen was showing the top arc, was being attacked directly from the front. The enlarged insert on the right was from a separate instance in the game where the player was being attacked from both 7 and 8 o'clock direction.

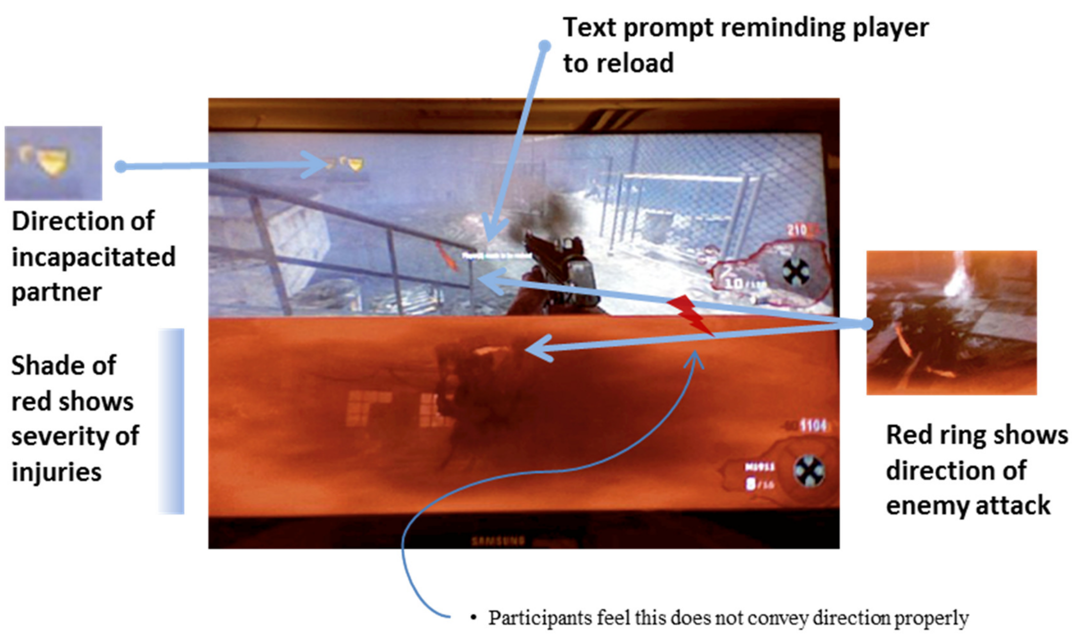

Figure 15: Artifact model of Call of Duty Black Ops' split screen display mode (Part 2)

Both of these techniques share a common issue. There was a mismatch between the mapping of the direction shown by the arrow and the player's first-person front-facing perspective. For instance, because the icon/arc system used a top down view of the world as shown in Figure 14, a forward direction into the screen (at 12 o'clock) would be shown as an arrow pointed upwards. However, from the player's viewing perspective, this upward direction indicated by the icon appeared to point towards the sky, which was above the player's field of view instead of the intended forward direction. The directions were being superimposed onto the wrong plane and failed to convey depth. The same was true for the down direction, which could easily be mistaken as below the user as opposed to behind.

3.3.2.3.2 Single Screen Mode

LEGO Star Wars provided its users with a more exocentric perspective of the virtual environment. The two players shared a single view of the environment. The game automatically and continually adjusted the viewing angle to keep both players' avatar in view as the players saw themselves moved around in a third-person perspective. Figure 16 is an artifact model showing the components of the screen. Similar to Call of Duty Black Ops, game information such as each player's health, points, and the player's current character was overlaid on the screen (at the top). If the objective of the level were to defeat an enemy boss, the boss' health would also be displayed.

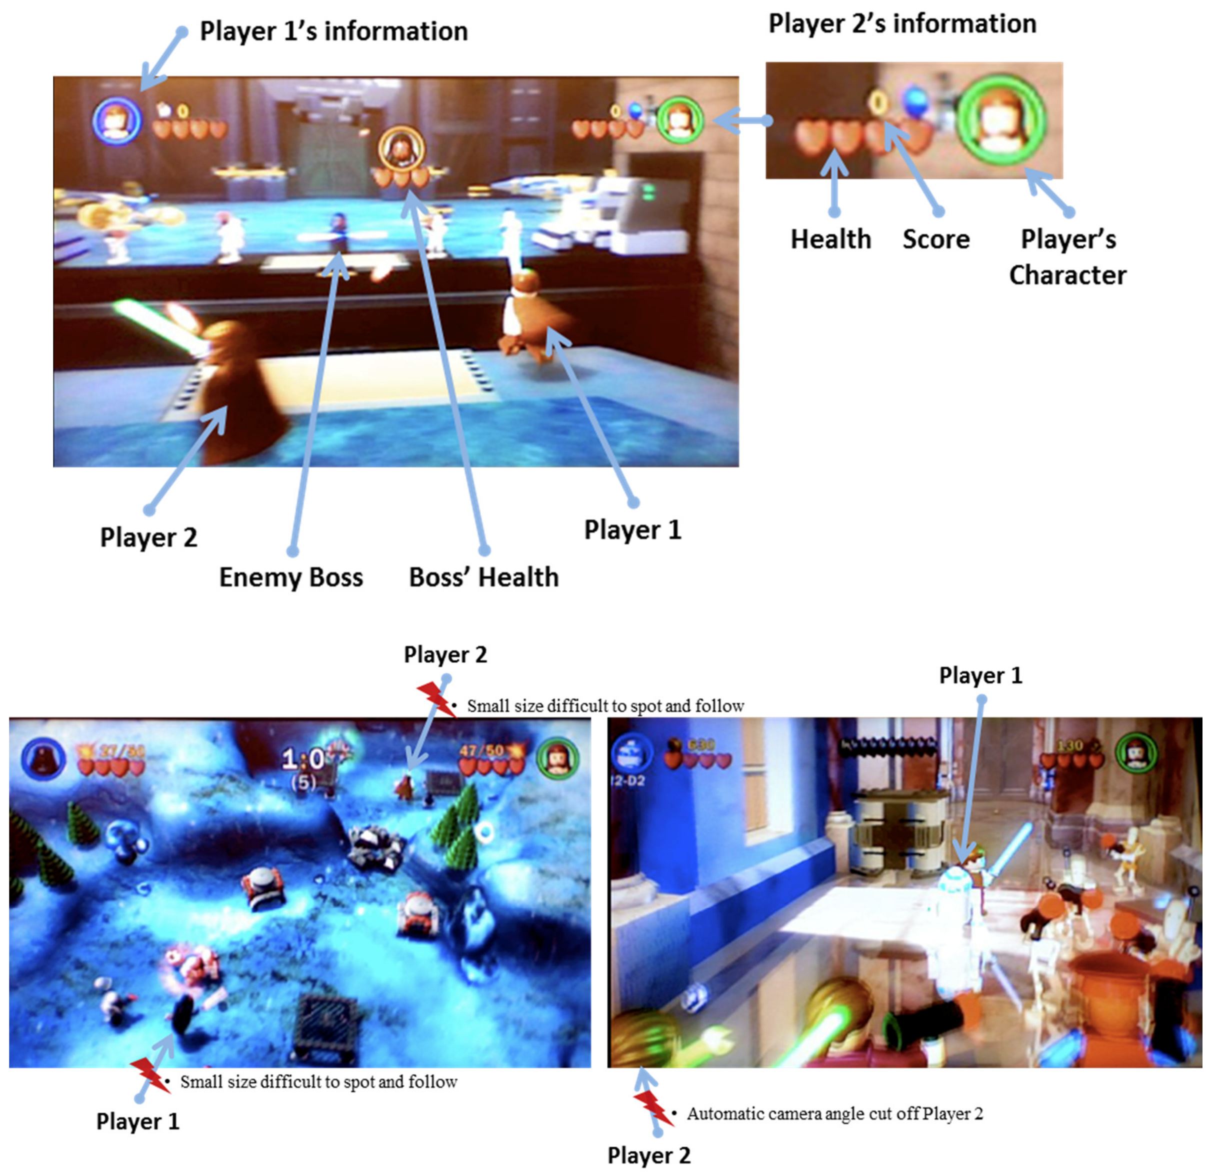

Figure 16: Artifact model of LEGO Star Wars. The top screenshot shows the visual overlays in the game. The bottom left screenshot shows players' avatars became very small as the automatic camera perspective adjustment mechanism attempts to keep them in view. The bottom right screenshot shows Player 2's avatar almost fell completely out of view due to the automatic camera perspective.

As mentioned earlier. The camera perspective was managed automatically by the game. The camera would pan, rotate, and zoom on its own in order to keep the players' avatars in view. The viewing perspective could also be affected by the virtual terrain as well as storyline elements. While the players had little or no control over panning and rotation, the players indirectly affected the camera's zoom by varying their proximity to each other. It appeared that the camera's zoom level was largely dependent on the distance between the players. The bottom left screenshot in Figure 16 showed how as the players moved apart, a larger portion of the world could be seen, although at a lower resolution, as the camera zoomed out. Conversely, as the players converged, the camera zoomed in to provide a more detailed view of the environment up close. This appeared to follow the reasoning that players often moved apart to explore the world and therefore needed to see the larger picture. When they came together, it was often to examine a location or item of interest together and therefore needed the additional details afforded by the closer zoom.

The aforementioned perspective management technique was far from perfect, however. A number of participants complained about the automatic camera adjustment during their interview after the game session. There were multiple instances where participants remarked that they could not see their avatar on the screen when the camera is zoomed out to the level similar to that shown in the bottom left screenshot in Figure 16. At this zoom level, the perspective often included many more distractions, such as enemy characters, landscape features or particle effects from explosions. Participants reported that they tried reacquiring their avatars by swinging a light saber or jumping to make their avatars more salient. Three out of four groups indicated that they had problems navigating when the camera perspective constantly changed. They complained about not being able to get to certain parts of the map or walking into things when the game's perspective management placed the camera at an awkward angle. An example of an awkward camera angle can be seen in the bottom right screenshot in Figure 16. Player 2 was almost completely cut off from view. Issues with the automatic camera perspective adjustment will be further explored in the critical incident analysis in Section 3.5.

3.4 Situation Awareness

The use of situation awareness (SA) analysis was brought on by the need to objectively determine the strengths and shortcomings of the display sharing techniques in the study. Just like navigating in the real world, a user navigating in a virtual 3D space while playing video games also built a mental model of the virtual world. The ability to determine the richness of that mental model and how close that model resembled the actual virtual environment could help in the aforementioned objective evaluation. The goal here was to find out what and how much players from each display mode knew of their virtual environment during game play. Endsley's (1995) SA model and the Situation Awareness Global Assessment Technique (Endsley, 1988), also known as SAGAT for short, were the tools for this evaluation.

3.4.1 Method

As mentioned, this study employed Endsley's SA model. Endsley defined situation awareness (SA) as a person's state of knowledge of the environment (Endsley, 1995). The relevant parts of this model are illustrated in Figure 17: Endsley's (1995) situation awareness model is modified to illustrate the increasing collaboration as SA levels increase. The two separate pipelines inside the group situation awareness box represented the progress of each player in the group. This will be explained in more details in the next section. There were three levels of SA in the model. The first level was perception. This occurred when the state of the environment was detected by the person's senses. Relating this to video games, first level SA was exhibited by the user's ability to perceive the video from the display, the sounds from the speakers, and the vibration provided by the game controller. The second SA level was comprehension. This occurred when the person derived meaning from the sensory stimuli from the perception stage and made sense of the environment. This study looked for second level SA in the user's knowledge of the floor plan, terrain and environmental features, the location of himself in relation to the overall floor plan and that of his partner. The third SA level was projection. Most often seen in experts, this was when the person used the current comprehension and made a prediction of the future state of the environment. This study looked for Level 3 SA in the user's ability to predict the spawning points of incoming enemies (COD Black Ops), and the movements of the stage boss (LEGO Star Wars).

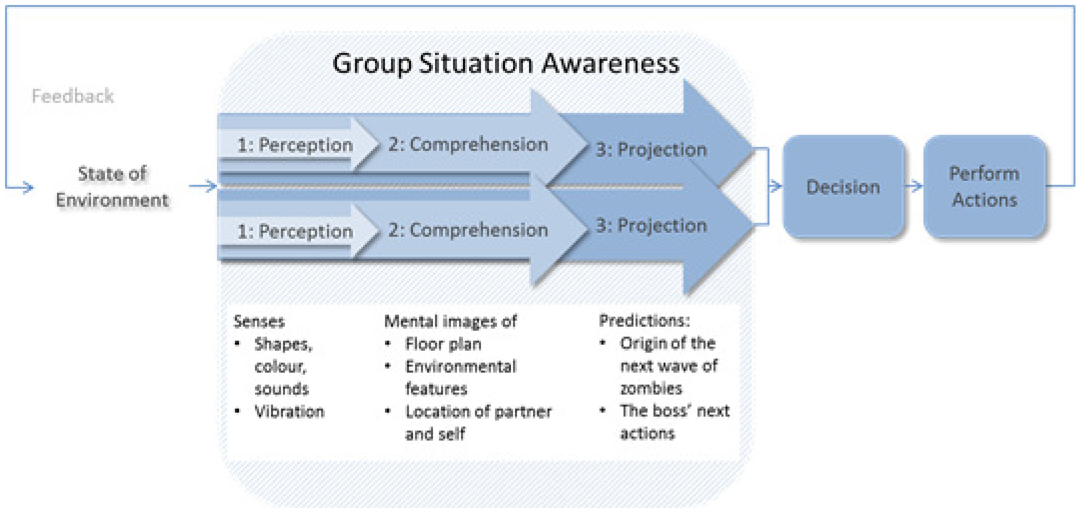

Figure 17: Endsley's (1995) situation awareness model is modified to illustrate the increasing collaboration as SA levels increase. The parallel tracks illustrate the awareness level of each player. The overlapping arrows shows the increasing collaboration at each level.

To collect the data mentioned above, there needed to be a way to capture a user's knowledge of the environment while the task was in progress. In this case, this study needed to capture the user's transient mental snapshot of the virtual environment during game play.

Endsley's SAGAT (1988) was adapted to do this. Participants were subjected to do a blank test. At a certain point during the game play, the users were told to pause the game. The display screen was then immediately turned off. The participants were then given a piece of paper. For COD Black Ops, the game was stopped approximately two minutes prior to the end of the game session. For LEGO Star Wars, the game was stopped once the players had gotten to a particular part of the map. These stopping points were chosen to ensure that players had time to explore and interact with the specific map in each game. Level 1 SA was not explicitly tested. For Level 2 SA, the participants were asked to draw the floor plan of the virtual environment in which they were playing. They were then asked to indicate their own location on the floor plan, the location of their partners, and finally any additional details that they remembered of the place. To test their prediction capabilities (SA Level 3), they were asked to indicate on the diagram from where the next wave of enemies would be entering (COD Black Ops) or what the boss was going to do next (LEGO Star Wars).

3.4.2 Findings

Aside from the SAGAT test, there were also some observations made regarding SA during the game play that was worth mentioning. Figure 17 showed two separate pipelines inside the group situation awareness box. This was to represent the distinct progression of the two individual players as they ran through the SA levels. In Level 1 SA, players perceive stimuli individually. In Level 2 SA, players were observed to communicate with each other to comprehend what it was that they were seeing. In Level 3 SA, even more collaboration was observed as they reasoned together to predict the future state of the environment. The following was a conversation from a group during game play in COD Black Ops.

P5: What the? [*P5 entered SA Level 1 when he noticed balls of blue flashes around him on the screen.*]

P6: What was that? [*P6 also noticed the flashes and entered SA Level 1*]

SA Level 2 began here when they tried to figure out what it was.

P6: Was a boss? [*Something jumped out of the ball and ran at him*]

P5: Dude, it's a dog!

P6: Yo, what are these things? They are everywhere.

Soon, both players were overwhelmed. They exhibited Level 3 SA by predicting their demise together.

P5: Out of ammo?! Hey uh, we ain't gonna to last.

P6: Yeah, we are dead. There is no way. It's ridiculous.

This increasing collaboration through the stages is illustrated by the incrementally overlapping arrows from left to right in Figure 17.

Participants produced a total of 16 drawings from the SAGAT test. Four groups of two users with each user producing two drawings, one for the third-person single screen mode of LEGO Star Wars, and one for the first-person split screen mode of COD Black Ops. For both display modes, novice and expert players alike, participants were able to produce accurate floor plans, terrain features, and environment elements of the virtual space.

All users displayed Level 2 SA in their drawings. While this study did not establish specific criteria in identifying levels of details, expert users produced richer drawings. For example, location of weapons, specific environment features like fences and window openings were only found in expert drawings. An expert also drew areas that he was not seen to have entered during the observed play session.

Similarly, a degree of Level 3 SA was displayed by all players. They were able to identify from where the enemies were coming. While the drawings of enemy spawning points were largely the same between players of different skill level on paper, expert players were observed to be better at predicting the enemy movement during the game play and were able to achieve much higher scores in the game.

More related to the difference between display modes was participants' ability to pinpoint where they were and where their partner was. In COD, the participants were generally confident about their own location in the map they have drawn. This was not the case when they were asked the location of their partner. Responses were mostly approximations. Figure 18 showed an example of this. In the figure were photos of the drawings produced by the participants. The drawing on the right was made by a participant playing the first person split screen mode of COD Black Ops. It could be seen inside the area marked by the circle that the user pinpointed his location with an "x" in one room while drawing an approximate area of where his partner could be. When asked why he did that, the particular user indicated that his partner was not in his field of view at the time. He could only make an educated guess. This turned out to be correct, as corroborated by his partner's drawing (not shown).

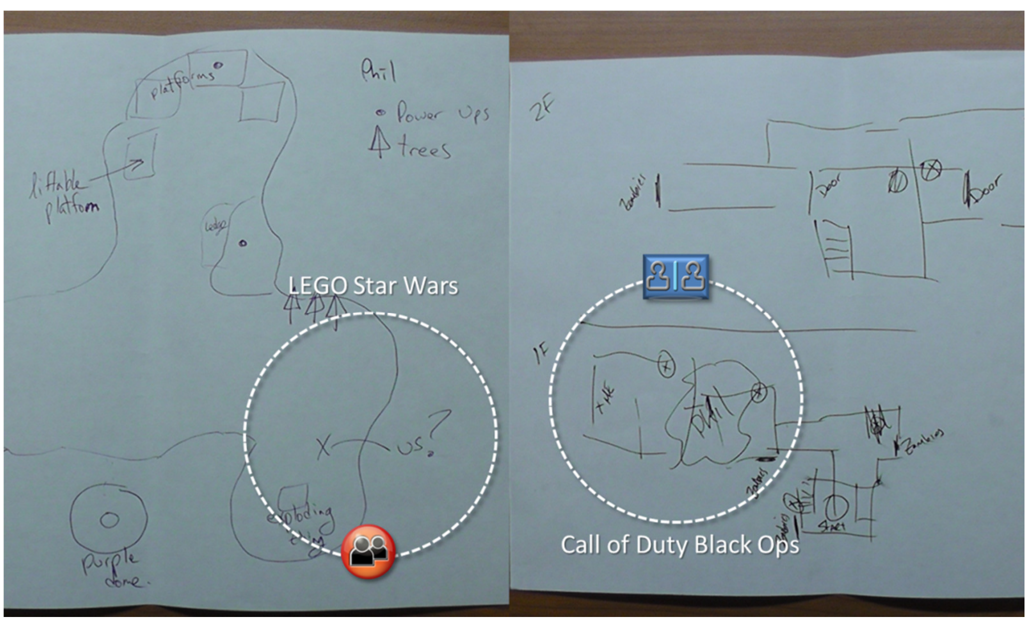

Figure 18: SAGAT results showing the difference in SA between display models

For LEGO Star Wars, the participants did not appear to have a difference in certainty when identifying their own location and that of their partner. While all of the participants were correct in identifying their position in the virtual environment, many of the remarks were, “we were mostly here.” When asked to identify individual positions, the markings tended to be close together. In one extreme case, as shown in the left drawing in Figure 18, the participant simply put a single “x--us?” When asked why it was he put the “?” the user simply stated that the camera was zoomed out pretty far. Although he knew where they both generally were in the map overall, he could not be exact in pointing out their individual locations.

3.4.3 Discussion

Findings from COD were interesting in that the split screen display provided an equal amount of information for both users. Player 1 had as much available information as Player 2 about Player 2’s location. Nonetheless, players reported that they typically did not look at the other player’s portion of the screen. In over 20 minutes of game play, players reported that they might have glanced over between 0 to 2 times. It is possible that experts have the cognitive ability to glance at the other player’s portion of the screen. However, the self-reported experts in this study stated that they do not do this often.

While differing SA Levels were observed between player skill levels, the same cannot be said about that between the different display modes. This was because the environments between the two display modes were different and a comparison between the results simply would not be meaningful. All that could be taken from this was that participants were accurate in reproducing the floor plan and the items they were asked to include. This was a limitation in the study. An alternative study design, namely Yang & Olson (2002), would be to use two navigation techniques conducted in the same virtual environment.

While the user’s knowledge of the environment between the display modes was inconclusive, the differing certainty in which users using the different display modes identify their location and that of their partners deserved attention. It suggested a degree of tradeoff between the two display modes. The first-person split screen view provided rich information in the limited eye level perspective but not anything that was outside of the first person perspective. Note, as discussed earlier, COD Black Ops tried remedying this by using arrow icons to indicate the relative position of the player’s partner. This apparently was not enough to overcome the player’s uncertainty over his partner’s location.

On the other hand, the third-person single screen view in LEGO Star Wars provided a good global view. It attempted to keep both the player and his partner in view at all times. As shown in the findings, both players became much more aware of where they both were. However, the view was now further away and resolution was reduced to the point where the players could only tell approximately where they were. This also restricted each player's movements, as later sections will discuss.

3.5 Critical Incident Technique

Much of the work in this study was to find out from the user's point of view the experience of using these different display modes during game play. Important insights could be gained from allowing the user to describe what they considered to be memorable parts of the experience and retell it in the form of a story. Flanagan's Critical Incident Technique (CIT) is well suited for this. CIT is a flexible set of principles that helps to discover the factors that contribute to the success or the failure of a mission (Flanagan, 1954). The user was guided through to narrate what was considered to be a notable event that had either positive or negative impact to his or her performance. Aside from being a useful technique, CIT's story telling component also conveniently lends itself to be integrated as a part of the CI interviews.

3.5.1 Method

CIT was employed after the participants were finished with the video game and towards the end of CI interview. At this stage, the researcher posted an open-ended question regarding what was a memorable experience during the game play. To guide the participants along in their narrations, open questions similar to the questionnaire used by Serenko (2006) were used:

- Was the incident positive?
- Provide details
- What was the outcome?
- Why was this critical?
- What were your feelings and perception?
- What action did you take?
- How did you change after?

3.5.2 Findings

While there were two display modes in this study, participants mainly told stories of their experience about LEGO Star Wars. A number of the stories were related to the camera view. Of the four groups, two groups provided stories of a particular part of the game where players had to navigate a set of overhanging platforms extended from a cylindrical structure. Figure 19 showed a simplified top down view of the players' path. The players had to travel counter clockwise from right to left around the dark blue structure. The players did this by jumping from one platform to next along the way. All except for two of the platforms were stationary. The two moving platforms (shown in orange in Figure 19) were retracted inside the dark blue structure by default. To get across, one player would have to extend the platform using telekinesis and keep it extended while the other player used it to jump across.

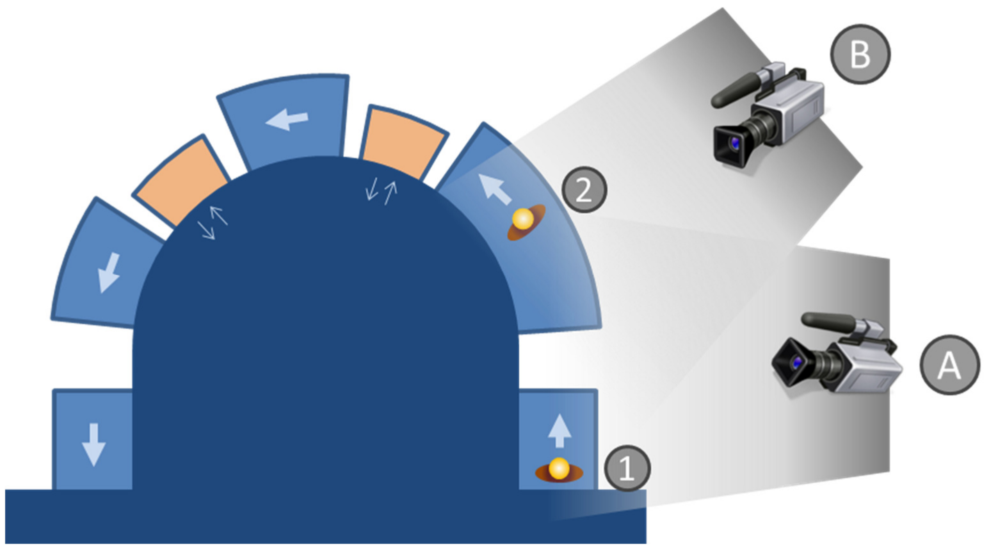

Figure 19: A top down view of the most frustrating stage in LEGO Star Wars

All of the participants who told the story indicated that this was not a positive experience. The stories all started with the two players entering the stage and beginning to tackle the platforms independently. At this point, one or both players would fail in the first jump and would fall off to a platform below the level. Getting back was quick. The players only had to return to the starting point, jump back up to the first platform and retry the first jump.

In cases where only one player fell off, the other player would proceed onward. The camera perspective would zoom out to accommodate the increasing distance between the two players. Eventually, the camera would reach its zoom threshold and could not zoom out any further. This was where the players realized that there was an issue. As seen in Figure 20, Player 1, the player who didn't fall off (to the right of the screen) was unable to proceed while Player 2, the player on the left, returned to the beginning to jump back on to the first platform. Player 1 was also trapped because he was unable to jump back on top of the platform he came from due to the change in the camera angle. Note that the platform on top (one from which he came down) was now directly on top of him. He was unable to manipulate his controls to get back on top. He was also unable to back up, since there was a ledge behind him. At the same time, Player 2 (on the left) did not have enough slack to jump back on to the first platform. The players were at an impasse.

Figure 20: At the camera perspective threshold players going in opposite directions were unable to go any further

The groups solved the issue in different ways. One group recalled that Player 2 (on the bottom) moved to the right to give Player 1 on top room to jump back on to the top platform. Then Player 1 backtracked towards the left just enough to let the Player 2 on the bottom jump back onto the first platform. Another group's story was less cooperative. Player 2 on the bottom refused to relent and eventually dragged Player 1 off the ledge. This caused Player 1 to fall off the stage and die. Player 2 reasoned that there was no cost to Player 1's death anyway, since Player 1 would respawn after a small delay at the closest point to where he died (at the edge of the screen). While Player 1 was gone, Player 2 was free to jump back onto the first platform. When asked how he felt about this, Player 1 laughed and stated that it was slightly unfair that he had to die (and lose his points) even though Player 2 was the one who didn't make the jump.

One of the group stated that it was a critical moment for them as they realized the limited camera perspective meant that they must work closer together. Failure of one player affected the other. Instead of exploring independently and going at separate pace, they learned to maneuver their characters together so that they would be on the same part of the screen. They also started to communicate more to ensure their actions were synchronized.

Every group in the study came to this realization at some point in this stage as no one was able to make it through without falling at least once. Those who fell off earlier actually had an advantage as there was less distance to backtrack. The groups that fell near the last platforms had to spend much longer backtracking and redoing all the actions. It was a pain point for a lot of the players. Having to repeat these maneuvers over and over again also exposed another issue. Players found this stage disorienting. After reaching the other end, one player actually questioned whether it was the place where they started.

3.5.3 Discussion

The stories told by the groups in the CIT sessions suggest that LEGO Star Wars' automatic camera adjustment, a mechanism that was designed to keep the avatars of both players in view in the shared-screen environment, hindered individual navigation by limiting the full range of movement. Because the view perspective was tied, movement of one player interferes with that of the other player. This points to a broader shortcoming that is also applicable to tabletop display in that the tied perspective hinders individual navigation.

Another issue was the method by which the camera perspective was managed. The constantly changing perspective was a result of averaging the two players' movements. In effect, players cannot rely on a predictable view through which they could navigate as each player had no control over the other player's movements, which contributed to perspective change. This contrasts with COD BlackOps whose players did not mention issues of interference and unpredictable view controls.

3.6 Summary of Findings

This section presents a summary of the findings from the three methods. The findings are categorized into four groups: general, split-screen mode, shared screen mode, and controller. Each finding is preceded by the method from which it was drawn. Related findings are also grouped together.

3.6.1 Difference in Situation Awareness

According to the blank test results, both split-view and shared-view players were able to produce the floor plan of the game environment accurately. What was different was their awareness of their own location and that of their partners. This is evident in Figure 18, the diagram drawn by a Call of Duty:

Black Ops player (split-view mode). It can be seen that the player was precise in pinpointing his own location. As for the location of his partner, he drew an area indicating where his partner could be. When asked why this was during the interview, the player indicated that he only had an approximate idea of where his partner was. Similar results were observed in other groups.

In contrast, the pictures drawn by LEGO Star Wars players told a different story. In Figure 18 (left), the participant drew an “x?” to indicate the approximate location of himself and that of his partner. When asked to explain the uncertainty, he indicated that their individual locations were hard to pinpoint because the camera perspective was zoomed out too far.

The difference found in the blank test results between the two screen modes suggests a variation in the type of situation awareness afforded by the two screen modes. The split-view mode appears to facilitate situation awareness at a more personal level, whereas the shared-view mode provides the players with a more global awareness of the group’s surroundings.

3.6.2 Interference between Collaborators

In split-view mode, players are free to navigate the virtual space on their own. While changes in their first person perspectives are tied to their navigation movements, the changes are localized to their own halves of the screen. A player’s own navigation does not cause interfere with his/her partner’s viewing perspective.

In contrast, share-view mode has only one viewing perspective (one that encompasses the avatar of both players). Navigation movements of either player can cause a change in the entire screen when the viewing perspective is automatically adjusted to keep both players in view. This was found to be an issue. For example, as the players move apart in the virtual space, the perspective zooms out to keep both players’ avatars in view. Eventually, the perspective reaches its zoom limit and the players are prevented from moving any farther apart. In another example, a player was attempting to jump over a gap while his partner was moving away in the opposite direction. The resulting perspective change caused the first player’s landing to fall out of view. His avatar fell to his death as a result.

While the shared-view mode places constraints on the players’ movements, these constraints were not completely counter-productive. Some participants indicated that they become more coordinated as a team because they have to be mindful of their partners’ movements. In comparison, players in split-view mode operate much more independently.

3.6.3 Difficulties of Automatic Perspective Adjustment

During the open interviews, a particular issue was consistently brought up by every group. In the shared-view mode, players felt disoriented in a specific part of the game where the pair had to navigate their avatars through a curved path while the camera perspective was continuously changing to keep both avatars in view. Participants indicated that the perspective changes felt unexpected at times, and navigation errors (in the form of the players' avatars falling off a cliff) occurred as a result.

This was due to the fact that player controls were tied to the perspective. When the perspective changes, the player must adjust his/her control input in order to maintain the original heading. Because this perspective was influence by both players' movements, an individual player was not able to completely predict the upcoming perspective change and adjust the control accordingly.

When asked how they felt this issue could be addressed, a participant suggested a default perspective that the players can recall and hold. Another suggestion was to give players the ability to temporarily freeze the perspective when precise navigation was required.

3.7 Exploratory Study Summary and Inspiration

The three findings discussed above each yield interesting implications to the sharing of screen space on the table top. The difference in the players' SA points to a trade-off between the two screen modes in terms of the type of situation awareness they provide. Split-view offers better individual SA while shared-view offer better group SA. Interestingly, the players of the most successful group (survived the longest) in the split-view mode were observed to verbally broadcast their individual locations and actions to each other continuously through the game. In doing so, they received the benefits of the individual SA offered by their split-view as well as the group SA offered by their verbal communication. This speaks to the importance of a collaborative system on the table top to facilitate both types of SA.

Interference between users occurs when the perspective between collaborators conflict. Binding their perspectives together, as the share-view play sessions have shown, creates restrictions on individual player movements. These restrictions also had the side effects of fostering closer co-ordination between players. Building on the need for both individual and group SA in tabletop collaborative systems, perhaps the side effect of individual perspective change could be used as a device for communicating SA between team members. Perhaps there is a way that individual perspective changes and its resulting interference can be made useful.

The shared-view's automatic perspective adjustment was disorientating for user. Players found it difficult to anticipate the upcoming changes and to adjust their controls accordingly. A hypothesis could be made that it was the instantaneous perspective change that left little time for the players to react. Building on the idea that individual perspective changes can be made useful, what if the perspective changes instigated by individual players did not cause an immediate screen wide change but a gradual distortion that propagates outwards from the origin of the change? Such propagation can also capture the temporal nature of the change. Changes that occurred first would radiate from outward from their origins followed by waves of later changes from behind. As the waves travel across the screen, players on the other end will have time to see the incoming perspective changes and adjust accordingly. The user receiving the propagation could, perhaps, have the option to maintain his/her current view perspective, as one of the participants suggested, or allow the incoming wave to wash over his/her own workspace.

Chapter 4
Design of Interaction Technique

The results of the exploratory study have provided the inspiration for the development of Waves, a new multiuser collaborative interaction technique for the navigation of virtual spaces. Specifically, this technique is designed for navigation of 2D information, such as a map from an overhead viewpoint. This technique is facilitated by Lattice, a flexible multi-user screen sharing framework developed in this thesis. Together, Lattice and Waves aim to support the following objectives:

- Support multiple personal workspaces
- Provide workspace awareness
- Mediate interference between workspaces

This chapter shows the design and implementation of Lattice as well as how this framework supports Waves and traditional collaborative navigation techniques such as split screen and single shared screen as explored in Chapter 3.

4.1 Design Guidelines for Collaborative Navigation Techniques

The design of the Waves technique draws from related work as well as findings from the exploratory study presented in Chapter 3. In this chapter, I focus on the following three guidelines in the design of a new collaborative navigation technique for the tabletop.

4.1.1 Support Multiple Personal Workspaces

Scott et al. (2004) observed that collaborators around a tabletop often switch between individual and collaborative work, and thus designers of collaborative tabletop systems should support both, as well as the transition between them. The exploratory study also showed that single shared displays and split screens each have their advantages in terms of awareness and navigation. For example, the single shared screen gameplay offered players more awareness of their partner's location, whereas split-screen gameplay offered players more freedom in individual movements.

A new collaborative navigation technique should aim to support both individual navigation and group interaction, such as the creation of a personal workspace from the main collaborative workspace and the subsequent integration of workspaces after the individual work is completed.

4.1.2 Provide Group Workspace Awareness

Since the collaborators are navigating within a spatially connected virtual space (i.e., a map) and doing related work, an awareness of the overall state of the workspace and of one another's view location and actions would allow them to work more effectively. The importance of this awareness was observed in the exploratory study where expert players verbally communicated with each other regarding their location and actions in detail throughout the game play. Aside from direct verbal communication, Gutwin & Greenberg (1996) also pointed out that consequential communication, feedthrough, and environmental feedback are also mechanisms through which workspace awareness can be maintained.

A collaborative navigation technique should support workspace awareness by having actions on one side of the table be visible on the other, without introducing sudden changes.

4.1.3 Mediate interference between workspaces

Gutwin & Greenberg's mechanics of collaboration (2000) also included protection of one's individual workspace as a key activity in collaboration using shared workspaces. This action is important, as individual work is often needed to complete collaborative work. An individual contributor needs a stable personal workspace that is free from interference from other collaborators in order to get this work done. On a tabletop, when multiple users are working on spatially connected data, such as a map, individuals should be able to navigate in their own protected workspace. Since the data is spatially connected, the individually changing views may interact and cause interference with each other. In the exploratory study, LEGO Star Wars protected these spaces using an automatically adjusted view that changes the shared view based on the "averaging" of both players' movements. Players had trouble with this technique, as the view often changed as a result of their partner's action. The players find this unpredictable and the resulting disorientation was detrimental to the players' navigation. A number of participants expressed a desire for greater control over their individual view.

Collaborative navigation techniques should account for the interaction between each collaborator's different view, and strike a balance between the communication of context and the management of interference between individual workspaces.

4.2 Collaborative Navigation Techniques

In this section, I discuss three alternative collaborative navigation designs: a single shared screen, a split screen, and Waves, the new proposed collaborative navigation technique. The first two techniques correspond to those studied in the exploratory study, but are considered in the context of a tabletop display, and the third represents an alternative that makes use of a flexible lattice to represent the navigation surface, and leverages the benefits of both. The implementation details of this navigation surface will be discussed in Section 4.3, and I here focus on a high-level description of the techniques and how they address the three design guidelines from a tabletop standpoint.

4.2.1 Single Shared Screen Navigation

This particular mode mimics the default collaborative navigation mode in which all users share a single screen with a single view. Panning done by any one user around the table causes a view change across the entire screen without any delay (Figure 21). This particular mode is suitable for single-user applications, as well as presentations where one user has complete control of the entire work space.

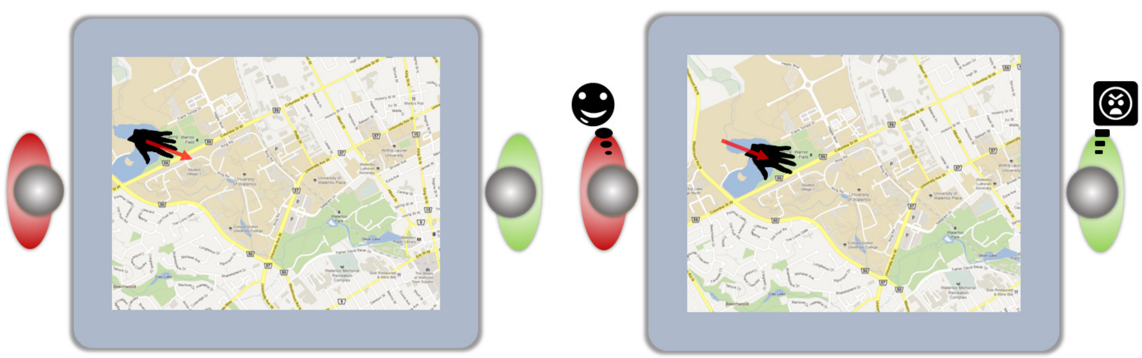

Figure 21: Single shared screen on the tabletop

Supporting Multiple Workspaces. This particular method does not support multiple workspaces as the display is limited to only a single view. For example, in order to explore two different parts of the map in detail, users must manage the single view using social conventions. This entails either positioning the view so that both areas of focus are visible or explore one after the other.

Provide Group Workspace Awareness. This method provides excellent workspace awareness as all collaborators share a single view. Collaborators can freely use any part of the viewable map as a reference for communication as there is no ambiguity of what the other person is looking at. There is also no overhead from needing to integrate a difference in view before establishing context between sections of the map.

Mediate Interference between Workspaces. Since any navigation has a global and instantaneous effect, any input by any collaborator will affect everyone else's view around the table. This technique does not provide any support for mediation of interference between collaborators. As there is no mediation, users will have to rely on social convention to mediate interference and take turns to perform individual work.

4.2.2 Split Screen Navigation

The split screen mode mimics Call of Duty: BlackOps' two-player game mode in which the surface of the tabletop is divided in half. The person on each side can navigate freely without interfering with the other person's view.

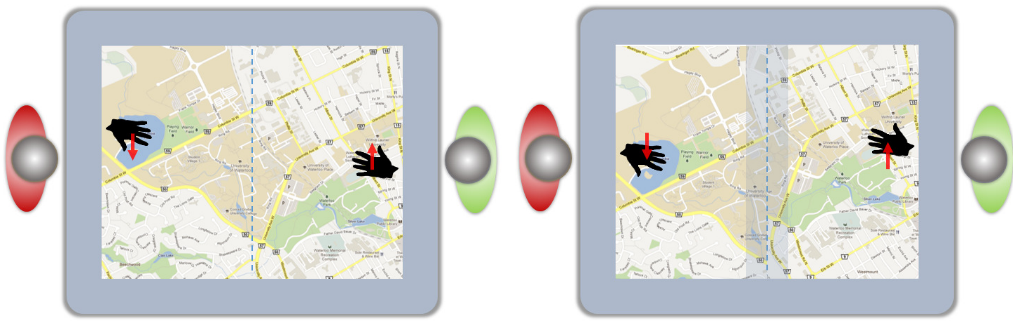

Figure 22: Split screen on the tabletop

Supporting Multiple Workspaces. Split screen methods support multiple workspaces. The two-independent-view split screen method provides complete independence between the views of the two users. Two different people can explore the same area of the map each in their own workspace without interfering with each other. However, it is not flexible in terms of the number of users it supports. The boundaries of the work area on the tabletop is predefined at startup and cannot be changed at runtime. The number of users must be known beforehand in order for the proper number of workspaces to be created.

Provide Group Workspace Awareness. This particular mode shared the same workspace awareness issues found in the console counterpart. It provides independent workspaces for the users and does not offer any cues that indicate the relative position between the two viewpoints. When comparing to Call of Duty: BlackOps, which offers an overlay arrow pointing to the direction of the other player (Figure 15), this mode provides even fewer cues. Two users can navigate independently in their own half of the tabletop and will have to rely on their knowledge of the features on the map to interpolate their relative positions. As such, this mode offers little support in the way of providing group workspace awareness.

Mediate Interference between Workspaces. The split screen mode provides two independent views of the map. Each user has complete control over his/her own view. Any view change as a result of individual navigation stays within the boundary of that workspace. As there is not interaction between individual workspaces, interference is non-existent, mediation is therefore unnecessary.

4.2.3 Waves

The above two collaborative navigation techniques each has its strengths and weaknesses. The single shared screen navigation mode does not support multiple workspaces, but offers group workspace awareness, as all users share the same view. The split-screen mode on the other hand offers an individual workspace for each user. However, it offers poor group workspace awareness as users do not have any visual reference indicating each other's relative view position. Neither technique mediates interference between collaborators. The table below summarizes the comparison.

Table 1: Comparison of how each navigation technique satisfy design guidelines

Design Guidelines	Screen sharing technique		
	Single	Split	Waves
Support multiple workspaces	No	Yes	Yes
Provide group workspace awareness	Yes	No	Yes
Mediate interference between collaborators	No	No	Yes

Waves is the proposed navigation technique that aims to fill the gap between the two different navigation modes. As seen in the table above, it aims to support all three design objectives. It does so by drawing inspiration from Information Cloth's (Mikulecky et al., 2011) malleable and stretchy surface. Waves allows multiple users to simultaneously navigate a map on a shared touch surface. Waves mimics the behaviour of a pool of water. Individual user navigation in the form of pan, zoom and rotate can be thought of as the user causing a disturbance to the pond's surface. The navigation action changes the view under the user's touch immediately. However, this action does not change the view of the entire surface right away. Instead, the view of a given area is updated based on its proximity to the disturbance and an adjustable time delay. Figure 23 illustrates this concept.

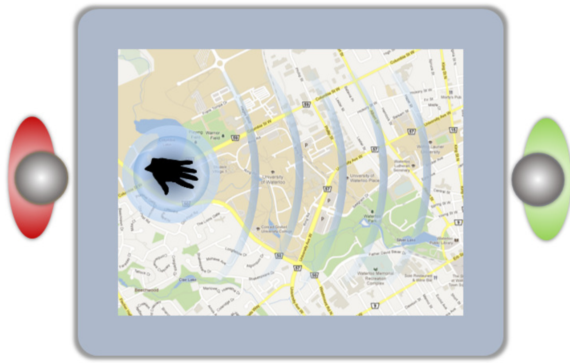

Figure 23: Concept of Waves on Tabletops

The Waves technique uses the physical property familiar from pools of water that movements occurring first will arrive before movements occurring later, as opposed to occurring instantaneously as in many computer interfaces (e.g., single shared screen movement). The state of change also persists for the duration of the propagation, allowing neighbouring users a larger time window to become aware and react, if so desired. Other users around the table who are closer to the movement origins will feel the view change sooner than those who are further away. The behaviour supports the idea that users who are near the origin of the view change are likely to be working on related tasks and that they should be informed of the change first. Those who are further away will also see the oncoming wave and have the opportunity to prepare his/her workspace.

Supporting Multiple Workspaces. Waves allows the creation of an individual personal workspace in two ways: on-demand and as-needed.

On-demand workspaces are created when the user touches the yellow border at the edge of the tabletop as seen in Figure 24.

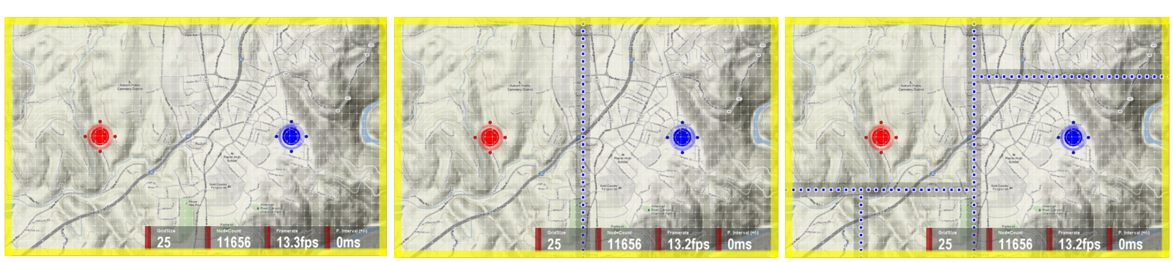

Figure 24: On-demand workspaces separated by explicit borders

This feature allows users to explicitly create a workspace in advance, before a navigation occurs, resulting in static borders. The blue on-demand borders extend from the user's touch point perpendicularly across to the other end or until it reaches another on-demand border. A newly created workspace can be further subdivided.

Another way a personal workspace appears is by the user simply navigating and changing their own view. A workspace is automatically created when this conflicts with another user's navigation (i.e., another view change). When users from each end pan in opposite directions (Figure 25), the view change from each navigation gesture propagates outward from the touch points.

Figure 25: Simultaneous navigation by two different users

When the edges of the propagation reach each other, a barrier is created at the point of the collision. This edge becomes the border between the workspaces. When this occurs, additional propagation from the same touch will not ripple beyond the border, thus creating a personal workspace. This workspace is maintained until the next navigation occurs. This technique can scale to the number of users that can comfortably share the screen. Figure 26 shows the dynamic creation of three workspaces by three different people interacting with the screen simultaneously.

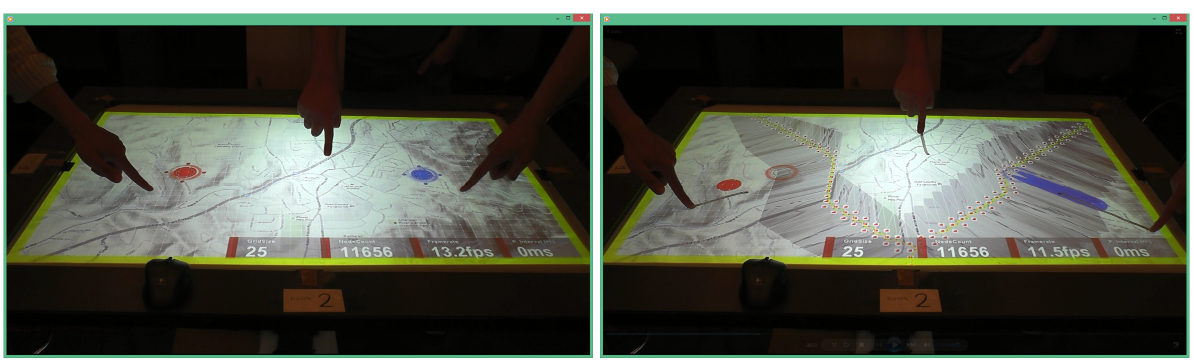

Figure 26: Ad-hoc workspaces separated by dynamic borders

Provide Workspace Awareness. Waves provides workspace awareness in two ways: communicating the view changes of each collaborator to others via *waves*, and representing the discontinuity between workspaces as a result of divergent navigation using *distortion*.

In the default navigation mode, where the entire group shares a single view, feedthrough (Baker et al., 2002) communicates workspace awareness. As one person navigates, the resulting global change in view of the entire tabletop is communicated to everyone around the table. In split screen mode, this mechanism is not available, as navigations are limited to each player's own half of the screen. In the exploratory study, expert players compensated by using continuous verbal communication to keep each other appraised of the overall state of the game. Waves serves the same function as the verbal communication in the split screen game in the exploratory study and provides workspace awareness by broadcasting the view change gradually and visually between collaborators. The changes in navigation over time are captured by the radiating waves. As waves from other collaborators reach a user, he/she becomes aware of how the views are changing.

The second way Waves provide workspace awareness is through the use of distortion to represent the discontinuity between workspaces with different views on the same surface. As each user navigates individually on the same surface, they begin creating and modifying their own individual view of the map. Between these views are areas that conflict. Figure 27 shows a surface with two users navigating.

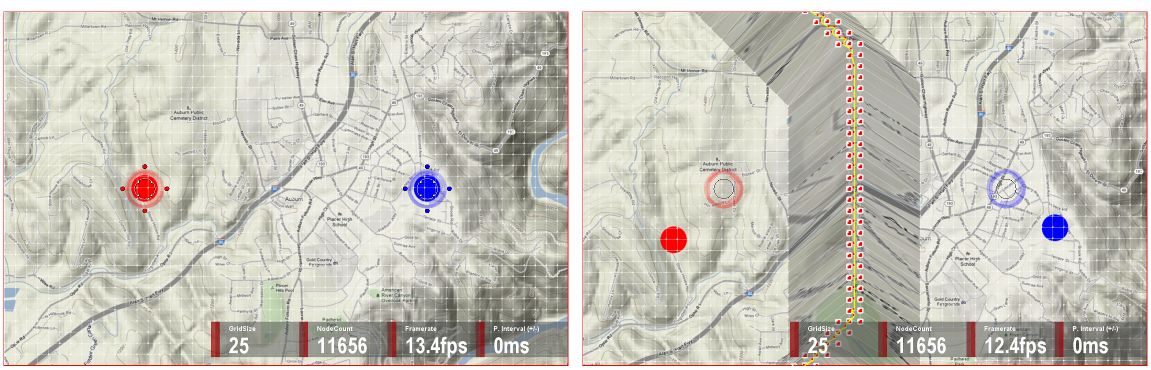

Figure 27: Distortion between workspaces provides spatial awareness

To represent the connection between spaces, these two areas are distorted; for example, a road running between the two workspaces can still be seen. Users can tell how their view is connected spatially in relation to the views of those around the table. This is especially useful when the screen is partitioned into multiple workspaces. Figure 28 below shows a road that runs through each of the partitioned workspaces and provides a spatial reference for every user around the table.

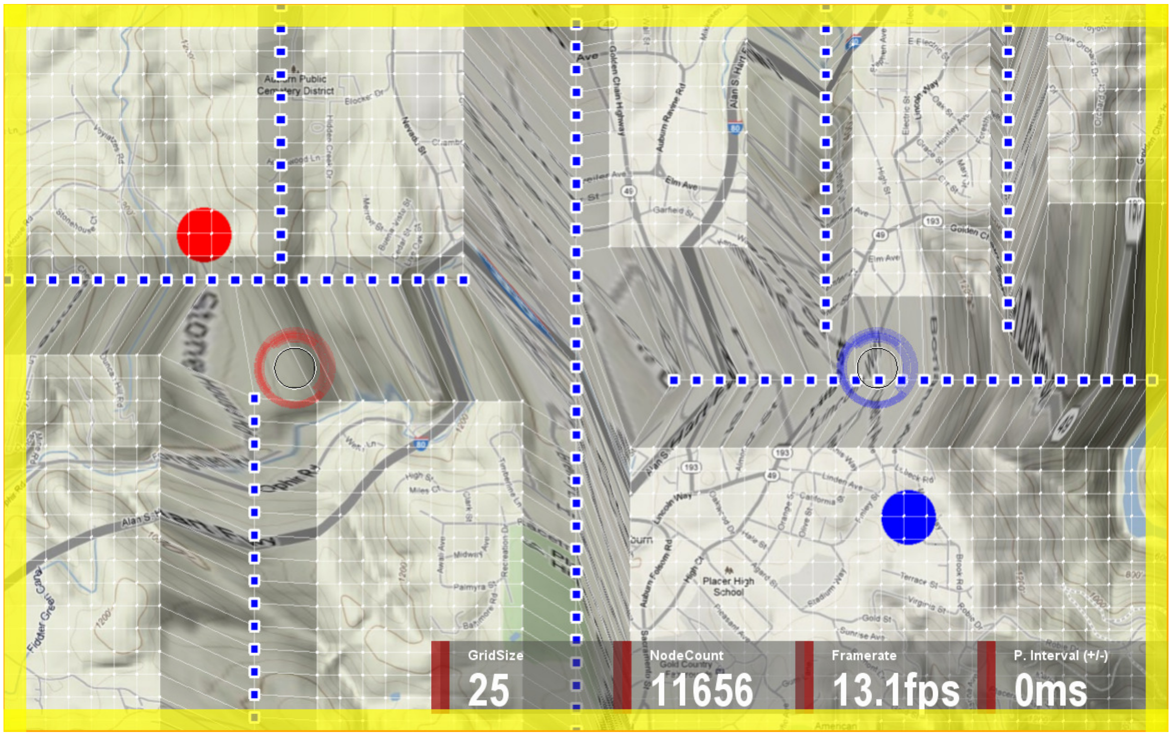

Figure 28: Distortion shows continuity between workspaces

Mediate Interference between Workspaces. Interference between workspace occurs when two or more users navigate simultaneously and their view collides. Waves mediates the interference between users' workspaces by dynamically managing the border of each user's influence. Figure 29 demonstrates the behaviour of the mediation mechanism. Navigation movements by the two users are represented by blue and red. Touch points are marked by "x". When a user navigates from the area marked by dashed to solid outline Figure 29 left pane), the view changes propagate outwards as marked by the darker color wave (Figure 29 middle pane). The border of the workspace expands until the waves propagate the entire space or when it collides with waves from a different workspace (Figure 29 right pane). The collision begins the formation of the dynamic border, as marked in yellow.

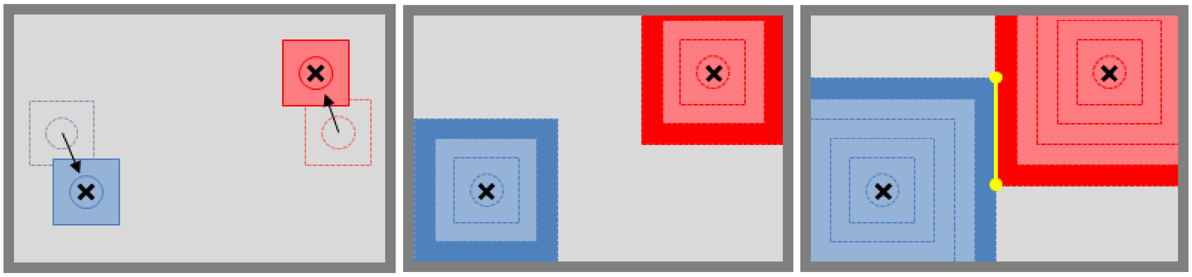

Figure 29: Wave's workspace mediation behaviour shows expanding waves (blue and red) and the resulting border formation (yellow)

Figure 30 shows the additional collisions that occur as waves from each user's movement propagate outward, and the dynamic border continues to take shape. The propagations continue until they cover the entire space.

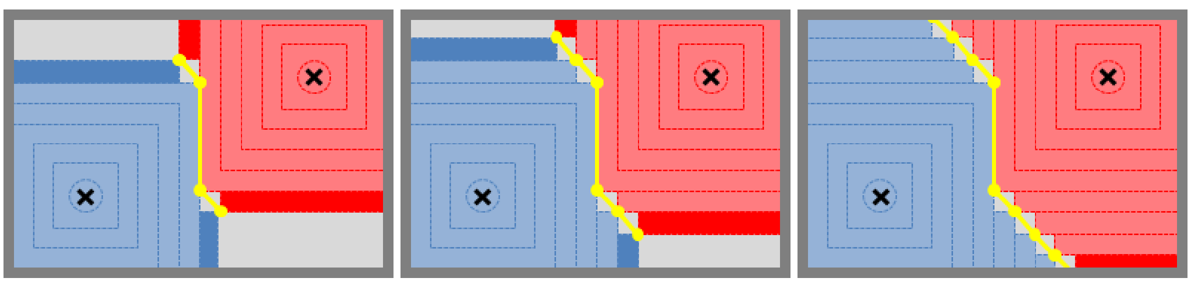

Figure 30: The forming of a dynamic border (yellow) in Waves between two expanding workspaces (blue and red)

Once the dynamic border is created, it stays in place for as long as the user continues to hold down the touch point that created the border. Waves given off by the same touch points by either user will not alter the border. This allows the users to continue to navigate. The assumption here is that the user does not wish for his/her space to be altered by others within the timeframe of a single touch transaction. Maintained action is required to retain the workspace.

Users can overwrite existing borders by navigating using new touch points. Subsequent waves from a different touch point reset and flatten the space. Figure 31 illustrates this.

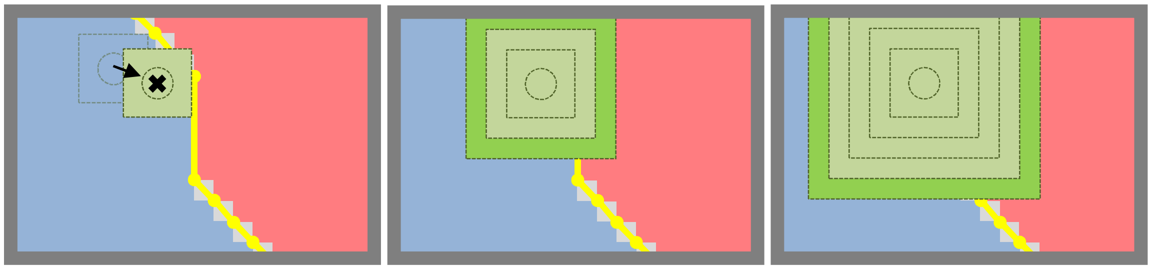

Figure 31: A new workspace (green) expands to overwrite an existing dynamic border (yellow) that divided two existing workspaces (blue and red)

Here the green area represents the user's view propagation. The user moved from the dashed area to the solid green area (left). Propagation then begins and expands to the entire screen. The expanding view propagation overwrites any existing borders.

If a user around the table sees this incoming wave propagating towards his/her workspace, he/she can choose to defend by initiating their own navigation, even if very small. Figure 32 shows a user (orange area) maintaining his/her view by interacting with his/her part of the screen (the orange area), so that the wave that gets produced propagates outward and collides with the incoming propagation to create a dynamic border (centre). The size of the resulting orange area depends on how early the user reacts to the incoming wave. The sooner the user reacts to block the wave, the large the area the user will be able to maintain.

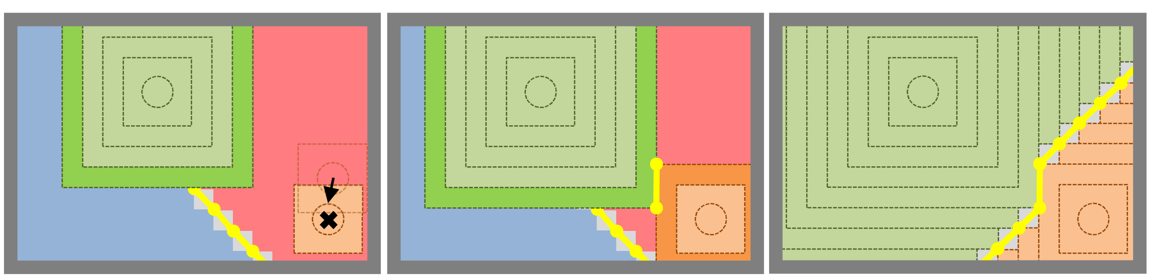

Figure 32: In Waves, one can prevent an existing workspace from being overwritten by creating another wave (orange) to block an incoming wave (green).

4.3 Waves Implementation

The described techniques were all implemented using a framework involving the mapping of screen graphics onto a vector map. They were created by varying the framework parameters that govern pixel movements in response to user navigation inputs. In doing so, varying degrees of communication or isolation between user workspaces can be adjusted. The major components of the framework are as follows.

The Wave interaction proposed in this thesis describes the propagation of the view change outwards from the origin of the touch. The delay as a result of the propagation allows multiple users to navigate simultaneously. The ability to provide context for collaborative navigation between users and individual navigation spaces for each user are provided through the framework's communication and isolation mechanisms.

4.3.1 Equipment

The prototype was developed using the Processing Java library. Multitouch functionality is provided by the Simple MultiTouch Toolkit[2], an open source toolkit that is based on the TUIO framework. The prototype runs at a variety of resolutions. Initially it ran at 1024 × 768 pixels on a SMART Table, a multitouch tabletop developed by SMART Technologies. Figure 33 shows the SMART table running the prototype at 8.1 frames per second with no movement or interaction occurring.

Figure 33: Waves running on the SMART Table

Subsequently, the development of the prototype was moved to a custom-made rear-projected table with a display resolution of 1280x800 and touches detected using diffused infrared lasers across the surface of the table, with a PlayStation 3 camera used to detect reflections of this infrared light.

[2] http://vialab.science.uoit.ca/SMT/

Common CV was used to process camera data and provide TUIO input to the same software prototype. The laser table was powered by a desktop PC with Intel i7 960 @ 3.20GHz and 12GB of RAM. In the following image, it can be seen that the prototype was running on the laser table at 9.9 frames per second with interaction from two touches. Note that in both cases, the frame rate depends on both the resolution of the display and the number of nodes used in the lattice implementation (see next section).

Figure 34: Waves running on a custom made infrared laser table

4.3.2 VectorNodes

The backbone of this technique is the lattice, an array of movable VectorNodes that forms the vertices of a mesh on which the image from the graphics buffer is mapped as a texture. The image is a large geographic map in PNG or JPG format. Figure 35 shows how each node on the lattice is laid out. To the left is the undisturbed lattice in its original state. Each node has two sets of coordinates, the node's location in texture coordinates *(u,v)*, which is bound to the map's texture and remains constant, and the node's display coordinates *(x,y)*, which represents the location of the node in the display space as seen by the user. On the left, the Figure 35 shows the initial state before user touches disturb the lattice. When a touch action is initiated, the node group surrounding that touch point is identified. The node group circled in black has matching *(u,v)* and *(x,y)* coordinates, with the top-left node at (25, 25); that is, there is no deformation in the display space. The image on the right shows the state of the lattice in the next frame of a touch move event. The touch movement was displaced by $(\Delta x, \Delta y)$ units, shifting the texture coordinates in the display space. The mismatch between *(u, v)* and *(x, y)* creates the visual distortion used in the Waves navigation technique.

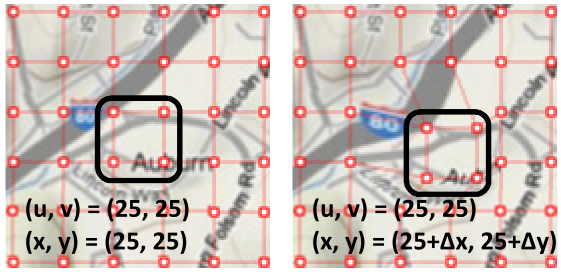

Figure 35: VectorNodes bound to map image

4.3.1 Propagation

To facilitate communication between user workspaces, the VectorNode Framework uses the metaphor of water waves in a pond. Movements made by a user in one area of the screen become waves that radiate out to other parts of the screen. In each frame, a node's transformation is determined by a combination of any touch actions occurring on that node, as well as any transformation propagating from its neighbouring nodes in the lattice. That is, the transformation *T* for node *i* at time *t* is:

$$T_{t,i} = \prod_{j \in N_i} T_{t-1,j} \cdot A_{t,i}$$

where N_i contains the set of node indices, *j*, for which n_j is in n_i's proximity (node *i*'s neighbourhood). *At,i* is the touch action for node *i* at time *t*. It is the identity matrix if there is no user touch near it at time *t*. If there is, then it is the matrix of the user touch transformation. This process is illustrated in Figure 36, where the three diagrams, from left, middle, and right, represent the state of the system at *t*=0, *t*=1, and *t*=2 respectively. On the left, a move vector is added to a node at *t*=0. In the middle, at *t*=1, the original node is moved and the same move vector is applied to adjacent nodes. At *t*=2, the nodes are moved, and the move vector is applied to nodes surrounding the first propagation.

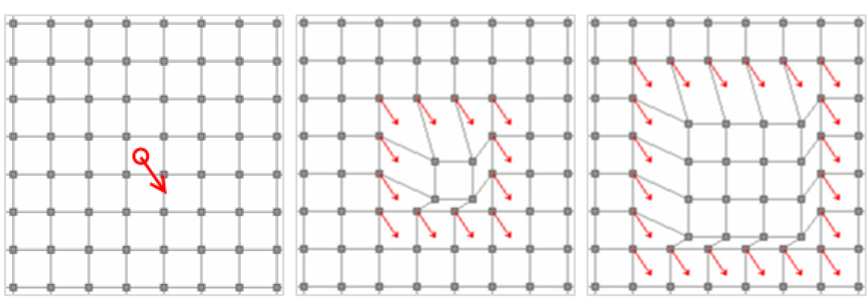

Figure 36: Propagation in the VectorNode Framework

This changes the node's display location on the screen and deforms the displayed image. This change is then propagated to adjacent nodes and eventually across the entire screen. The movement and update process is further described in the Propagation section below.

Different screen sharing modes can be created by adjusting the rate at which the move vectors propagate. For example, Figure 36 shows a propagation "radius" of one adjacent node per time interval. This produces a ripple effect as changes can be seen propagating outwards from its origin over time. Alternatively, the propagation radius can be changed to include the entire screen. This in effect would reproduce the traditional single shared screen navigation mode.

This framework supports the basic navigation movements such as pan, rotate and zoom, as shown in Figure 37, Figure 38, and Figure 39 respectively.

Figure 37: Translation in the VectorNode Framework

Figure 38: Rotation in the VectorNode Framework

Figure 39: Zooming in the VectorNode Framework

4.3.2 Wave

Propagations of touch actions are organized in waves. That is, a touch action (translation, rotation, or zoom) can occur over several frames, and each frame initiates the next "wave" of transformation. Each wave is started by the transformation as interpreted by the SMT touch package over a small period of time. This wave is implemented as a data object that contains a list of nodes on the lattice that is currently touched by the wave. Figure 40 below illustrates this:

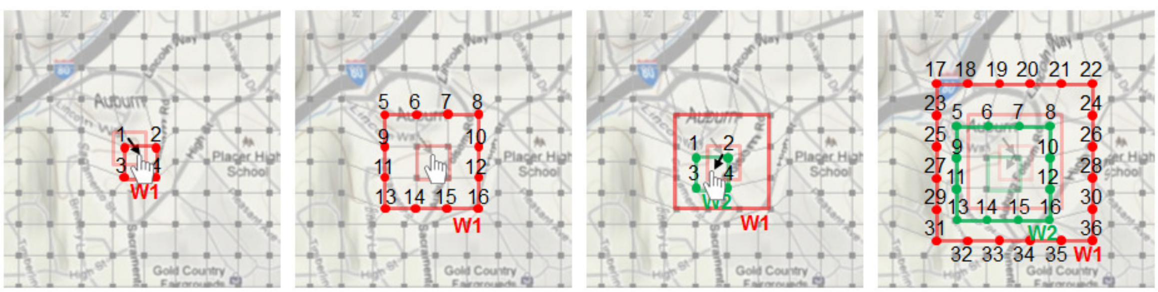

Figure 40: Wave formation and propagation

From the left, user touch movement (direction as indicated by the black arrow) creates the first wave (W1), which contains nodes [1–4]. The second pane from the left shows W1 has propagated outwards and now includes nodes [5–16]. This is done by building a list of nodes that neighbour the origin of the wave. The third pane from the left shows that the next frame of movement from the current touch event starts the second wave (W2), with nodes [1–4]. The forth pane shows the next propagation where both W1 and W2 move outwards. W1 propagates outwards to nodes [17–36], and W2 follows behind with nodes [5–16] moving from their previous positions, as changed by wave W1 (light red), to the current dark green position as numbered. As the user touch moves, additional waves are created at the touch point and the existing waves continue to propagate radially outward until it has covered the entire map or collided with another wave that was created by a different touch point.

It is worth noting that the interval between each propagation can be adjusted. By varying this parameter, it is possible to vary the speed at which view changes travel across the tabletop. This changes the amount of time each user has to react to the interference coming in from other users.

4.3.3 Collisions and Borders

The resolution between competing views is handled by the creation of borders at the locations where two waves created by different touch points collide. Once collided the border is frozen in place until a wave from a different touch point that was not involved in the initial collision reaches the border location. This collision behaviour is implemented by the movement characteristics built into each node on the lattice. Figure 41 shows the logic behind the node's movements. Logic items are designated from A to J for ease of reference. Before each propagation cycle, the movement of each node in the lattice needs to be determined individually according to the movement logic. The process begins at Step A for Node j. Movements of nodes come from their contact with waves. In the wave's data structure is a list of nodes it currently touches. Step B looks for Node j in each of the active waves in the workspace. Depending on the number of waves that "touch" this node, four different possibilities will occur from Step C, which I will next describe in detail: *no movement, movement by a new wave*, *a new collision, no movement due to a previous collision*, and *movement clearing previous collisions*.

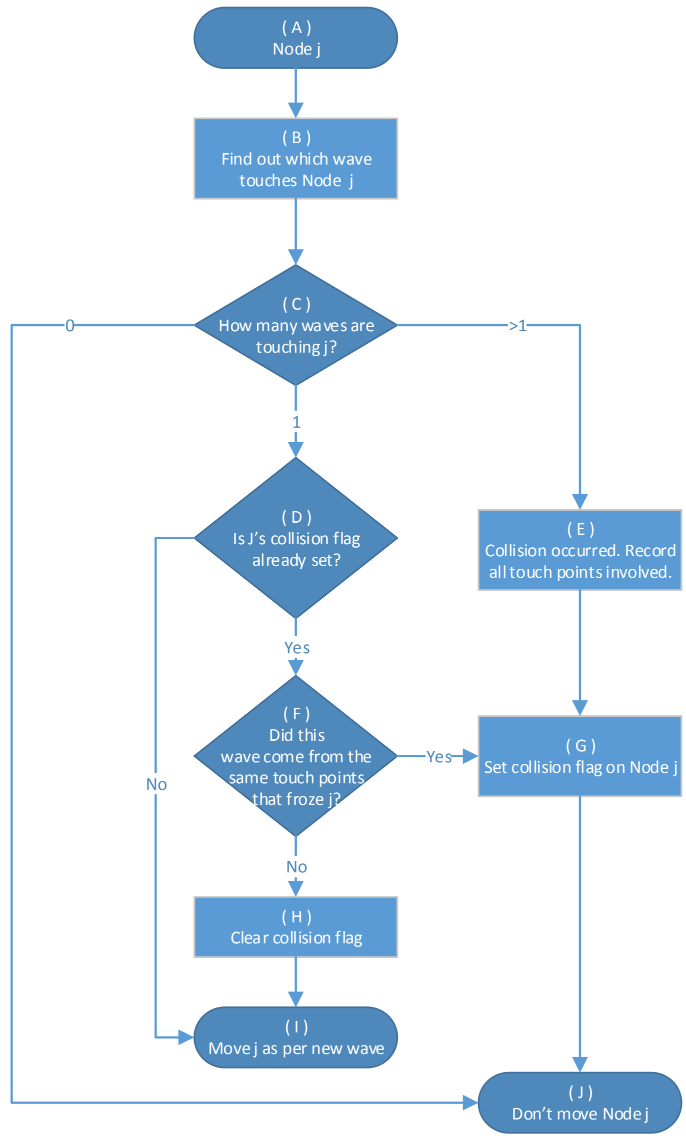

Figure 41: Logic diagram of node movement and collision

No movement has step sequence {C, J}, where Node j did not show up on any wave's node list. This means there are zero waves touching Node j and it is not involved in any movement this propagation cycle. The logic flow then goes to Step J, where Node j does not move.

Movement by a new wave occurs with the sequence {C, D, I}, when Node j is touched by a single wave and was not previously frozen as part of a collision (i.e., its collision flag is not set, Step D). This is the case for standard movement and the logic concludes at Step I where Node j is moved as per the wave's movement.

A new collision has the sequence {C, E, G, J}, and occurs when multiple waves are touching Node j. When this occurs, Step E records the touch IDs from which the colliding waves originated. This information is used to identify the touch points that created the collided node later on, as this mechanism is used to contain future waves created by the same touch event. Step G sets the collision flag in Node j's properties to indicate that this node is now part of a collision. Since Node j is collided, it does not move (Step J).

No movement due to a previous collision has the sequence {C, D, F, G, J}. This occurs when a single wave hits up against a collided (frozen) node that is part of a border that was previously created to contain the waves from the same touch event. A single wave is touching Node J (Step C), the Node is frozen in place, recorded by its collision flag (Step D), and this single wave originated from the touch points whose wave this border is trying to contain (Step F). Thus, the collision flag on Node j remains set (Step G), the Node j does not move (Step J), and the border remains in place.

Movement clearing previous collisions has the sequence {C, D, F, H, I}. This occurs when a single wave hits up against a collided (frozen) node and the wave did not originate from the same touch points that created the border (which the frozen Node j is a part of). This sequence is common to the no movement due to a previous collision sequence up to Step F. Because the wave did not come from the same touch points that created this border, of which Node j is a part, the process moves to Step H, which clears the collision flag in Node j's property. As Node j is no longer collided, Step I moves it as per the movements of the new wave.

4.3.4 Dynamic and Static Borders

The Wave prototype allows for the separation of workspaces using visible borders. There are two types of borders: implicit and explicit.

Dynamic borders appear as a by-product of users' touch navigation. They are created and destroyed automatically by the system according to users' navigation actions. The automatic nature of this is similar to the automatic camera view adjustment from LEGO Star Wars. The difference here is that the Wave technique maintains each user's view and dynamically resolves the interference between competing views when necessary. Figure 42 shows the sequence of the actual interaction of the prototype. Overlays are added in the diagram to highlight the interactions. At t = 0, two users pan the map in opposite directions creating two waves, W1 and W2 (illustrated by nodes with white outlines). At t = 1, W2 propagates outwards (illustrated by nodes in outlined in yellow) and collides with W1. The node where the collision occurs, the node is marked as collided (its colour turns red), and is frozen in place. The darkened area around the collided node is distorted due to the displacement of nodes caused by the wave. At t = 2, W1 propagates outwards (illustrated by nodes with yellow outline) and collides with more nodes from the previously expanded W2. This causes additional collisions. At t = 3, W2 propagates again, and more nodes are collided.

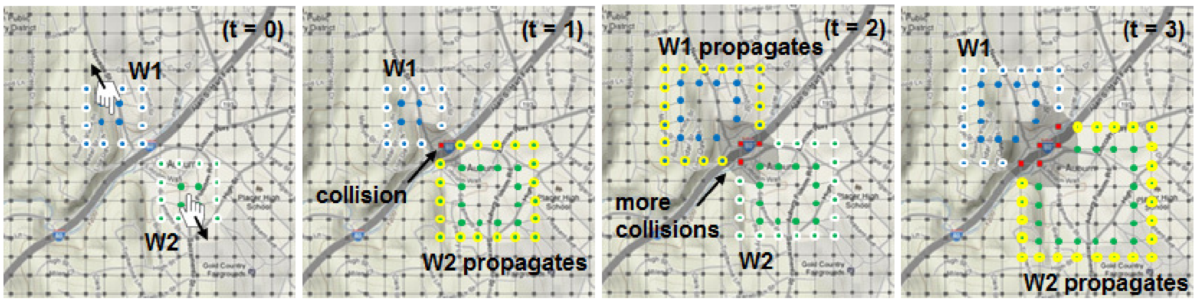

Figure 42: Sequence of dynamic border creation

As collisions happen at the wave fronts, they create a chain of collided nodes that forms the border of each user's workspace. Figure 43 continues at t = 4 and shows the sequence of dynamic border creation with the yellow overlays showing W1's propagation. The collision continues to grow as W1 and W2 expand from t = 6 through 10. The prototype connects adjacently collided nodes with Bézier curves to further highlight the border between the workspaces.

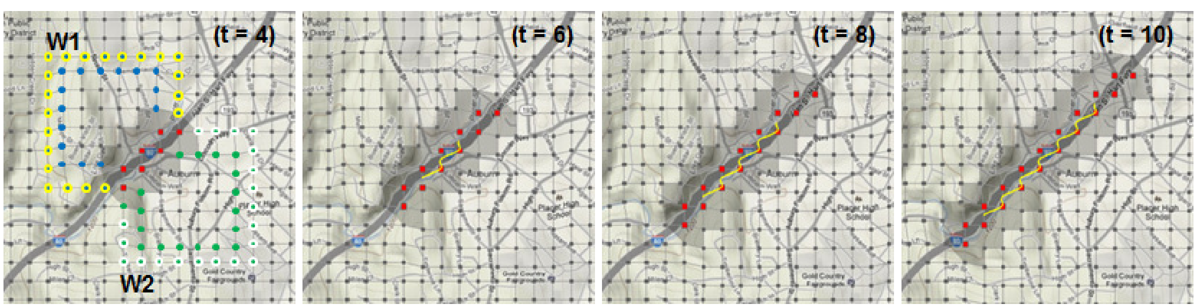

Figure 43: Dynamic borders are drawn based on positions of frozen lattice nodes

Once a dynamic border is created, the nodes along the border remain frozen in place and do not allow propagation of any other incoming waves generated by the same touch event that created it in the first place. As long as the users continue to navigate using the same touch event (i.e., a continuous drag), the borders will remain intact and their personal workspaces are maintained.

To reintegrate workspaces separated by dynamic border, one of the users at the table simply has to lift off from the current touch points that created the border and initiate a new touch event. This is shown in Figure 44. At t = 0, the area is divided by a dynamic border that was created by earlier user interaction. When the user pans with a new touch point and creates W1, since W1 did not originate from the same touch point, as per the collision logic from Figure 41, at t = 3, W1 is allowed to propagate through the dynamic border. As it does so, W1 clears the border by realigning the frozen nodes on the border to the lattice within its expanding perimeter. This process continues through t = 9 and beyond until W1 propagates to the entire surface.

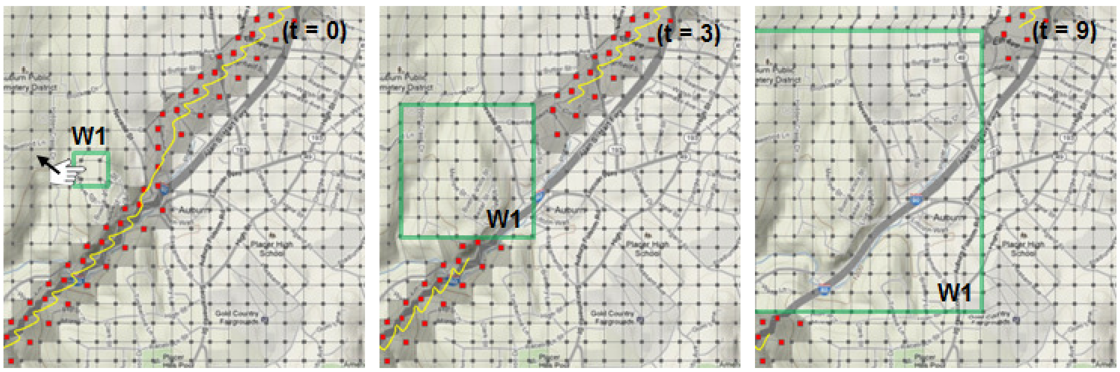

Figure 44: Dynamic borders are overwritten

In the situation where a user would like to maintain her workspace and protect her existing view from an incoming wave from a different user, she can "defend" her border, using a different touch event to create waves of her own to collide with the incoming wave. This is shown in Figure 45. Assuming that at t = 5, when W1 in Figure 44 is clearing the dynamic border, a second user pans on her side and creates W2. At t = 6, W1 and W2 collides to create another dynamic border. At t = 9, W1 continues to clear the old border. At the same time, however, the collisions between W1 and W2 create a new border.

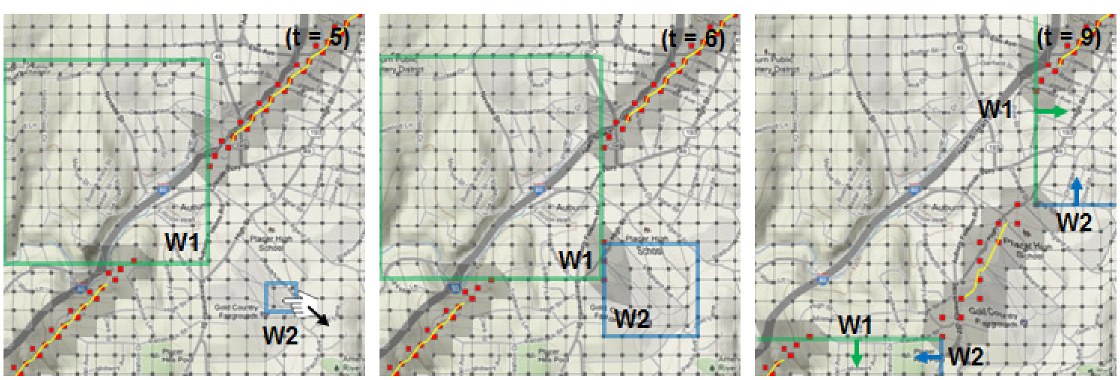

Figure 45: User initiates a wave to defend her workspace from an incoming wave

Static borders can be used in cases where users would like to create an on-demand workspaces on the tabletop that are free from interference from others. They can be created on the workspace by tapping on the yellow border that runs along the edge of the tabletop workspace. This extends a border perpendicular to the yellow edge out towards the other side of the table. Figure 46 shows the location where the user taps to create an explicit border. The red arrows show the direction of the border extension.

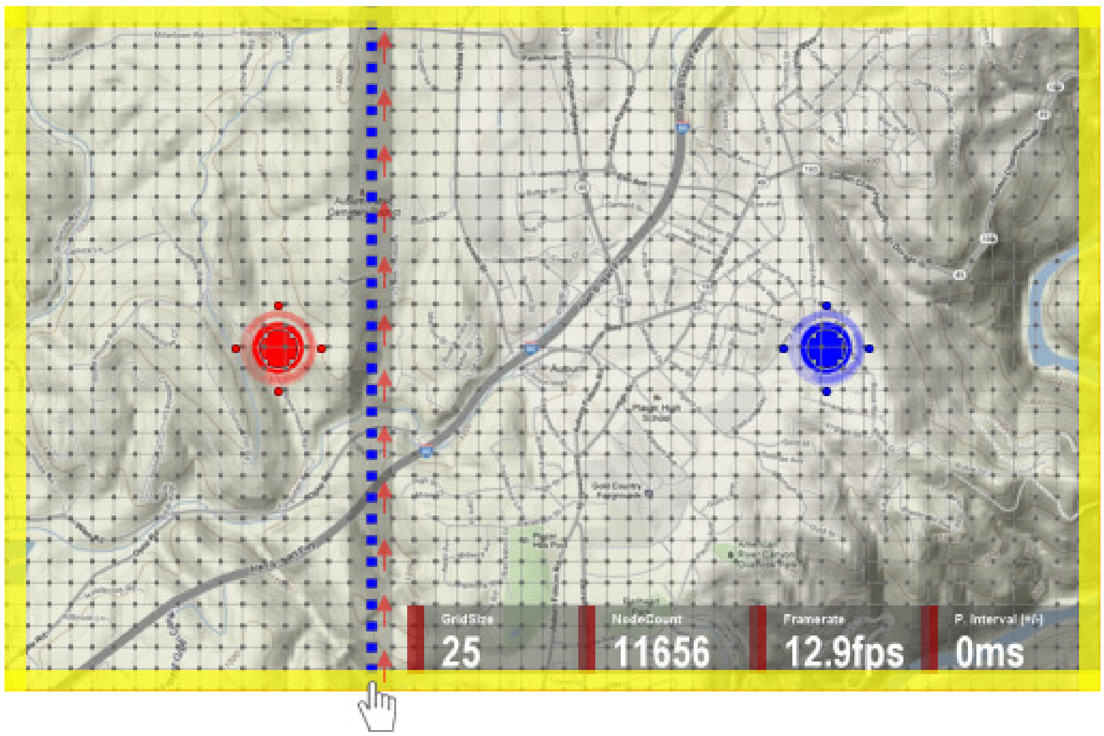

Figure 46: Manual creation of an on-demand workspace using explicit border

The mechanism used here is similar to that of dynamic borders. The implementation consists of finding the node on the lattice that is the closest to the user's touch point and freezing it in place. Once the first node is frozen, depending on which yellow edge was touched, the explicit border creation function traverses all the nodes on the lattice in the direction perpendicular to the yellow edge. Those nodes appear in blue in Figure 46.

To allow the creation of multiple on-demand workspaces, explicit borders is implemented such that subsequent borders do not cross existing borders. That is, the explicit border creation function stops traversing when it reaches a node on the lattice that is already frozen by an existing explicit border. This allows the remaining workspace to be additionally subdivided as seen in Figure 47.

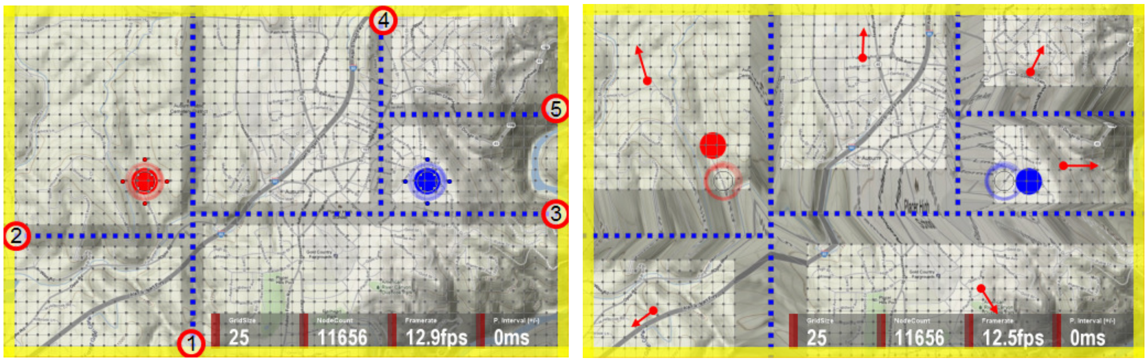

Figure 47: Creation of multiple on-demand workspaces with explicit borders

The left pane Figure 47 shows the sequence in which multiple explicit borders are created to subdivide the workspace. The right pane shows users being able to navigate the area within each workspace independently.

4.4 Waves Limitations

While the Waves technique allows for the dynamic creation of workspaces, and aims to reduce interference between navigation events from multiple users, there remain several limitations.

4.4.1 View competition

The Waves technique relies on the users to actively maintain their views by creating their own waves to prevent their views from being overwritten by waves created by other users. A potential issue may arise when there is an imbalance between the interactions of the user with their respective workspaces. For example, the more the user interacts with her workspace, the more waves she creates, and the more influence over the entire workspaces she exerts. Other users around the table must match the amount of interactivity in forms of wave creation to "defend" their workspaces. Static borders were implemented to provide users with an alternative to constantly needing to defend their workspaces.

4.4.2 No Support for Data Duplicity

While Waves combined the strengths of the two popular screen sharing techniques from the exploratory study, the split screen navigation still provides simultaneous access to the same area on the map by more than one person. The malleable nature of the Lattice allows the simultaneous viewing of locations that are physically close by, but individual workspaces cannot show data from the same location in the map. Users are expected to take turns and share this data through social convention.

4.4.3 Waves Requires a Large Table

User navigation generates waves that occupy screen real estate. Even though they are transient, the areas occupied by waves are not ideal for detailed examination as the map details are likely obscured by the wave distortion. In cases where there is a large number of simultaneous interactions around the touch table, a large portion of the table will be covered by waves. Combined with the area taken up by collisions, the area available for useful work may be limited if the table is not sufficiently large.

4.5 Reception

While this technique has not undergone any formal validation, the Wave prototype has been shown in informal demonstrations to members of the author's research group. The following are some of the reactions and verbal feedback, both positive and negative.

4.5.1 Positives

Users stated that they like the prototype's ability to support more than two people browsing at the same time. The visual presentation of the technique, namely the waves generated by the user's movements, appeared to delight the users. They discovered the dynamic border creation without explicit instructions. They were comfortable creating the additional dynamic borders once the first ones were created. They also recognized that the space between their workspaces is a distortion and could be used to identify the relative spatial positions between workspaces. Static border creation was often discovered only by accident. However, it was said that it was a good feature to include.

4.5.2 Negatives

Users did not like their spaces being overwritten by an incoming wave that they didn't notice until it was too late. They commented that they constantly needed to monitor the area surrounding their own workspace and that being ready to defend their workspace could be tiring. Safeguarding an area also means that the user has to move the view, which may not always be convenient. People also commented that they were distracted when navigating by the numerous distortions generated by others' movement. They also commented that borders also often "sneak up" on them. This occured when a small incoming wave that was not noticed collided with the user's touch location very soon after the touch down event, leading to confusion.

4.6 Summary

Informed by related work discussed in Chapter 2 and the exploratory study in Chapter 3, the Waves technique is designed and implemented to satisfy three design guidelines: (1) support multiple personal workspaces, (2) provide group workspace awareness, and (3) mediate interference between workspaces. (1) was addressed by the use of dynamic borders, a by-product of user navigation, and static borders, an on-demand border that can sub divide the work area, to create separation between workspaces. (2) was addressed via the use of distortion around dynamic and static borders to represent the continuity between the views of workspaces disjoined by individual user navigation. (3) was addressed by the adaptation of the metaphor of the waves on a malleable surface to communicate what would otherwise be instantaneous view changes created by individual user navigation around the table. The wave collision mechanism creates dynamic borders that mediate interference between workspaces. The technique received both positive and negative reception from users.

Chapter 5
Conclusion and Future Work

This thesis began with the definition of the problem: how can multiple users simultaneously navigate and explore a map on the same tabletop without interfering with each other? It then proceeded to explore the problem in the following steps:

- Exploration of existing related work in literature
- Exploratory study of existing screen sharing solution in video games
- Combining findings from literature and study of games to form design guidelines
- Development and implementation of Waves, a navigation that simultaneously supports multiple personal workspaces, provides group workspace awareness, and mediates interference between workspaces

In this chapter, I summarize each step above, revisit the major contributions of this thesis, and discuss future work.

5.1 Summary of Literature Exploration

The literature exploration focused on three areas: interactive tabletops, collaboration with the tabletop, and information visualization techniques. It began with an introduction of the interactive tabletop, the underlying platform on which this thesis is focused. It explored multi-touch techniques, technologies and applications of tabletops. The second area of the literature exploration looked at research in collaboration using the tabletop. Previous work such as Scott et al.'s (2003) system guidelines for tabletop displays and Gutwin & Greenberg's (2000) mechanics of collaboration formed the basis of my design guidelines for a navigation technique that allows multiple users to share tabletop workspace and mediate interference between them. This was followed by the third area of the literature exploration, which explored information visualization techniques that have been proposed in the past to support the exploration of geospatial data, including the use of distortion techniques, such as the Mélange, to illustrate multiple areas of focus on a single surface. The distortion of the areas connecting workspaces on a tabletop provide individual user context on how each workspace is spatially related to the others. Other important previous research includes Mikulecky et al.'s (2011) information cloth, where data, such as a geographical map, is rendered as the texture of a malleable cloth that can be stretched, pinched, and folded. This served as the basis of the Wave navigation technique central to this thesis.

5.2 Summary of Exploratory Study on Games

To gain a better understanding of collaborative navigation techniques currently in use, this thesis looked at how two popular video games solved the problem of having multiple users with a common objective navigate on the same screen. The first game, Call of Duty: Black Ops is a first person shooter where two players navigate the game environment in split-screen mode. Each player is free to navigate freely within their portion of the screen. In the second game, LEGO Star Wars, two players navigate their third-person avatars in the game environment in the same view without partitions. Three different user research methods were used to gain insights into the advantages and shortcomings of each screen-sharing technique.

The study resulted in three major findings: differences in situation awareness, interference between collaborators, and the difficulties of automatic view adjustment. Players have better group situation awareness while navigating using the shared screen mode than split screen mode. This is because in shared-screen mode, players have the benefit of an explicit visual reference to their avatar's relative position. However, players' movements interfered with each other while navigating in shared-screen mode, due to constraints in the shared viewport. For example, players cannot navigate further apart than what the viewport allows. Another issue was LEGO Star Wars' technique of keeping each player in view. It continuously adjusts the viewport based on the average of each player's movement. Players found this disorienting, especially when the view changes due to another player's movements in the mist of their own navigation.

The results of the study identified a tradeoff between the level of interference form view changes between players due to the connectedness of their viewports and the level of group situation awareness that accompanies the interference. This finding lead to the hypothesis that a temporal component could be added to that global view change such that the interference between players could be lessened while at the same time group situation awareness is still communicated. This served as another inspiration for the creation of Waves.

5.3 Summary of Design and Implementation of Waves

The design of the Waves navigation technique draws from related work as well as findings from the exploratory study presented in Chapter 3. The process began with the formalization of design guidelines for collaborative navigation techniques for tabletops: to support multiple personal workspaces, to provide group workspace awareness, and to mediate interference between workspaces. Using these guidelines, an evaluation of the two collaborative navigation techniques in the exploratory study (shared and split screen) in the context of digital tables was carried out. It was determined that the split-screen technique does well at providing support for multiple personal workspaces and does not do well at providing group workspace awareness. In contrast, the shared-screen technique does not do well at providing support for multiple personal workspaces, but does well at providing group workspace awareness. Neither technique mediates interference between collaborators.

The proposed Wave technique satisfied the first two guidelines by combining the split screen technique's ability to provide a personal workspace with the shared-screen technique's ability to provide group workspace awareness. This hybrid method then mediates the interference between the collaborators' competing view views by making the workspace a malleable surface as inspired by previous literature on using a malleable surface to display information.

The Wave technique was prototyped on an infrared laser touch table. It was implemented by mapping the top down view of the workspace (a geographical map) as a texture onto a moveable lattice. As users interact with the workspace (pan, rotate, zoom), the area directly under the user's touch would be sticky and the view local to that area changed immediately. This view change would then propagate outwards to the rest of the global workspace, much like waves in a pond. Collaborators around the table can navigate the map local to their area and the different views are resolved when the waves from competing views collide. The collision generates a border that consists of visually distorted space that shows how the two views are connected.

5.4 Contributions

The contribution of this thesis include:

- An exploratory study providing evidence for differences between two popular collaborative navigation techniques used in video games, split screens and single shared screen, in situational awareness, interference between collaborators, and difficulties with automatic view adjustment (Chapter 3).
- The design and implementation of Waves, a collaborative navigation technique for the tabletop based on guidelines formulated from the results of the exploratory study; Waves

simultaneously supports multiple personal workspaces, provides group workspace awareness, and mediates interference between workspaces (Chapter 4).

5.5 Future Work

While the Temporal Waves technique was designed with the established guidelines in mind and has demonstrated in informal testing that it can sustain navigation by three different users, there is potential for future work. This section presents some of these possibilities.

5.5.1 Explore the transient nature of waves and borders

There are several parameters that were built into the prototype that have potential for further exploration. One of which is the speed at which the view change propagates. It would be interesting to see whether the rate of propagation affects collaboration. Another possibility is to vary the speed of the propagation according to the distance from the origin of the wave. Making the propagation close to the origin faster may make the system appear more responsive to the user creating the waves. To the first user, the faster the propagation, the more responsive the system feels. However, this same propagation is interference to other users and, to them, slower may be better. An additional possibility is to experiment with the attenuation of the size of the initial wave according to the touch movement. For example, small movement would generate a small initial wave while a large movement would generate a larger initial wave that instantaneously affects a larger region. A user study to find the attenuating function of these two variables, namely wave speed and initial wave size, would help find the optimal balance.

5.5.2 Explore the use of alternative gesture to address view competition

The current design uses the side-effect of navigation to create dynamic borders between workspaces. One possible issue with this technique is that a user who frequently interacts with the workspace will easily disrupt others. A wave will always overwrite a dynamic border given that it was not created by the same touch point, regardless of a user's intentions. Currently, the design provides a separate mechanism (in the form of tapping a yellow border) to create an explicit, static border. An interesting alternative that is worth exploring to is to incorporate the user's intention as part of the gesture. For example, the user could use multiple fingers instead of the minimum required to create an individual workspace that is impenetrable. Such an alternative technique can reduce the likely misclassification of users' actions and transform side-effect operations into explicit first-class operations. This method may also eliminate the need for a separate and completely different method of creating static borders, on which the current design relies.

5.5.3 Optimize prototype

In terms of implementation, the current prototype is not optimized for performance. At the time all features were added, the prototype runs at about 13 frames per second, with the current hardware setup and chosen propagation speed, and gets progressively slower as the number of waves increases. Incorporation of optimization, such as screen buffering and minimizing the number of tasks in each draw loop, could greatly improve the performance of the prototype.

5.5.4 Formal User Testing

While the prototype has been informally tested during development, the next step should be a formal study to test the performance and the degree to which group situation awareness differs between split screen, shared screen, and Wave. The performance can be found by measuring the time it takes for a group of users to navigate to a marked location on a map. Group situation awareness could be measured using similar techniques to what was used in the exploratory study. A second interesting factor to explore would be whether the Waves technique allows users to perform better or worse when the navigating to targets on or off the screen.

Appendix A
Exploratory Study Material

This is a sample Appendix. Insert additional appendices with the "Start New Appendix" command.

A.1 Recruitment Script

Hello, my name is Joseph Shum and I am a graduate student in the Department of Systems Design Engineering. I am currently taking a course (SYDE642 Advanced Ecological Interface Design) with Professor Catherine Burns. I am studying how multiple individual navigate collaboratively through digital 3D space in video games. This research will hopefully lead to a better understanding of collaborative navigation and eventually be used to further research in interface design of groupware systems.

If you volunteer as a participant in this study, you will first be asked about your video gaming habits and your familiarity with two commercially available video games: Call of Duty and LEGO Star Wars. Then, with a partner, you will be asked to play these two games. After the game play you will be asked about your experience playing the games. The session will be video recorded for later analysis. I will also conduct an open interview with you after the play session to discuss your experience. The session should take approximately 1 hour of your time.

I would like to assure you that this study has been reviewed and received ethics clearance through the Office of Research Ethics. However, the final decision about participation is yours.

If you are interested in participating, please reply to this email or contact me in person at CPH3641.

Thank you.

Joseph Shum

A.2 Interview Script

The interview will mainly concentrate on their experience during play. No PID will be recorded.

Previous Habits

- Tell my about your video gaming habits?
    - number of hours played per week, duration in each session
    - type of games played
- Have you ever played these particular games before?

Game play

- Tell me about things you find difficult in the game?
- What are the things you thought you should have been able to do, but couldn't?

Collaboration Specific

- How does having a partner in this game help you?
- How did you keep track of where your partner was?
- How did you communicate with your partner?
- Talk about the kind of strategy you used here [EXAMPLE OF STRATEGY OBSERVED DURING PLAY].

Navigation Specific

- What are the things about the controller that you like?
- What are the things about the controller that you don't like?
- What happened when you and your partner wanted to go to different parts of the map? How did you guys decide?
- What are some of the visual aids you find useful?

- Did you always know where you were in the game?

- Can you contrast the difference between the games modes you experience?

- what are the things you had to keep track of while play?

A.3 Consent Form

Collaborative Navigation Study
University of Waterloo System Design Engineering
Course Instructor: Professor Catherine Burns
Email: catherine.burns@uwaterloo.ca

Participant Information Letter and Consent Form

Overview

You are being asked to volunteer in a study as part of a course project at the University of Waterloo. The intent of the course (SYDE 642 Advanced Ecological Interface Design) is to apply human factor research techniques on tasks performed at work, home, or play. The course instructor is Professor Catherine Burns.

The purpose of this project is to better understand how team members performs a goal oriented task in a shared digital three-dimensional environment. Of particular interest is how the team members interact with each other and the environment as they move through the digital three-dimensional environment. The goal is to find out how, what, and when gets passed between the users themselves, between the user and the digital environment, and between the user and the interface in the physical world.

You will be asked to play two video games with a partner. In these games, you and your partner will work collaboratively to complete the shared objectives in cooperative mode.

Participation Details

If you volunteer as a participant in this study, you will first be asked about your gaming habits and your familiarity with two commercially available video games: Call of Duty and LEGO Star Wars. Then, with a partner, you will be asked to play these two games. After the game play you will be asked about your experience playing the games. The session will be video recorded for later analysis. I will also conduct an open interview with you after the play session to discuss your experience. The session should take approximately 1 hour of your time.

You may decline to answer particular questions, if you wish, and may withdraw participation at any time by advising the student investigator.

Risks

There are no known or anticipated risks to you as a participant in this study other than those associated with the normal use of computers, video game consoles, or mobile computing devices.

Confidentiality

All data collected is considered confidential. Codes, rather than names or other identifying information, will be used in notes and/or recordings. Even though we may present our findings to the class, only the course instructor and our project group will have access to the data collected. Your name or any other personal identifying information will not appear in any publication resulting from this study. However, with your permission, anonymous quotations may be used.

You will be explicitly asked for consent for the use of photo/video data,captured from the video recording, for the purpose of reporting the study's findings. If consent is granted, these data will be used only for the purposes associated with course presentation and you will not be identified by name.

The data collected in this study (i.e., notes and video recordings) will be kept in a secured lab on campus at UW (CPH 3641) over the duration of the study and will be confidentially disposed of. Such material will not be retained. The data will be retained for no more than two years.

Compensation

You will not be compensated for participation in the study.

Questions

If you have any questions about participation in this study, or would like additional information to assist you in reaching a decision about participation, please contact Professor Catherine Burns at (519) 888-4567 ext. 33903 or via email at catherine.burns@uwaterloo.ca.

The contact information of the students conducting this research are:

Joseph Shum	jshum@uwaterloo.ca

This study has been reviewed and received ethics clearance through the Office of Research Ethics at the University of Waterloo. However, the final decision about participation is yours. If you have any comments or concerns resulting from your participation in this study, please contact Dr. Susan Sykes of this office at (519) 888-4567 Ext. 36005 or ssykes@uwaterloo.ca.

Consent Form

I agree to participate in a study being conducted by Joseph Shum for the course SYDE 642 (Interface Design) at the University of Waterloo. The course instructor is Professor Catherine Burns. I have made this decision based on the information I have read in this Information-Consent Letter and have had the opportunity to receive any additional details I wanted about the study. I understand that I may withdraw this consent at any time by telling the student investigator.

I am aware that I have the option of allowing my session to be video recorded to ensure an accurate recording of my responses.

I am aware that excerpts from the interview may be included in the final report, with the understanding that the quotations will be anonymous.

I am aware that I may allow video and/or digital images in which I appear to be used in the course presentation with the understanding that I will not be identified by name.

With full knowledge of all foregoing, I agree, of my own free will, to participate in this study.
___ YES ___ NO

I agree to the session being video recorded.
___ YES ___ NO

I agree to the use of anonymous quotations in any presentation or report that comes of this study.
___ YES ___ NO

I agree to allow pictures and videos of the session to be used in the course presentation and report with the understanding that I will not be identified by name.

___ YES ___ NO

___________________________ (Please print) Name of participant
___________________________ Signature of participant
___________________________ Date
___________________________ (Please print) Name of witness
___________________________ Signature of witness
___________________________ Date

Permissions

Request for permission to use figures from your paper
2 messages

Joseph Shum <strongside@gmail.com> Tue, Jul 22, 2014 at 9:04 AM
To: Mark Hancock <mark.hancock@uwaterloo.ca>

Hi Mark,

Can I please have your permission to use the images in your paper:

Mikulecky, K., Hancock, M., Brosz, J., & Carpendale, S. (2011, November). Exploring physical information cloth on a multitouch table. In Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (pp. 140-149). ACM.

--Joseph

Mark Hancock <mark.hancock@uwaterloo.ca> Tue, Jul 22, 2014 at 9:12 AM
To: Joseph Shum <strongside@gmail.com>

You have my permission.
Mark
[Quoted text hidden]

Request for permission to use diagram
2 messages

Joseph Shum <strongside@gmail.com> Sun, Jul 20, 2014 at 12:57 PM
To: Stacey Scott <s9scott@uwaterloo.ca>
Cc: Mark Hancock <mark.hancock@uwaterloo.ca>

Hi Stacey, I am currently finishing up my thesis... you should be seeing a draft shortly.

In my related work section I talked about command and control scenarios for the table top and referenced your paper

Investigating Tabletop Interfaces to Support Collaborative Decision-Making in Maritime Operations
http://www.eng.uwaterloo.ca/~s9scott/wiki/uploads/Main/scott_iccrts2010.pdf

Do i have your permission to use one of Figure 2 or Figure 3 from that paper? (see below)

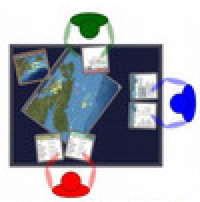

Figure 2. The system provides flexible, adjustable information and data windows to accommodate use from any side of the table.

lows in use.

ort for orienting interface components in order

Figure 3. System-level menus are accessible from any side of the table.

Thanks!

--Joseph

Stacey Scott <stacey.scott@uwaterloo.ca> Sun, Jul 20, 2014 at 1:32 PM
To: Joseph Shum <strongside@gmail.com>
Cc: Mark Hancock <mark.hancock@uwaterloo.ca>

Hi Joe,

Great to hear from you. Yes, you have my permission to use either or both of these figures.

Sincerely,
Stacey

--
Stacey Scott
Assistant Professor, Systems Design Engineering and English Language and Literature
Associate Director, Games Institute
University of Waterloo, Waterloo, ON, Canada

p: 519-888-4567 x32236
e: stacey.scott@uwaterloo.ca
w: www.eng.uwaterloo.ca/~s9scott

References

Antle, A. N., Bevans, A., Tanenbaum, J., Seaborn, K., & Wang, S. (2011, January). Futura: design for collaborative learning and game play on a multi-touch digital tabletop. In *Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction* (pp. 93-100). ACM.

Baker, K., Greenberg, S., & Gutwin, C. (2002, November). Empirical development of a heuristic evaluation methodology for shared workspace groupware. In *Proceedings of the 2002 ACM conference on Computer supported cooperative work* (pp. 96-105). ACM.

Benko, H., Morris, M. R., Brush, A. B., & Wilson, A. D. (2009). Insights on interactive tabletops: A survey of researchers and developers. *research. microsoft. com*.

Bowman, D. A., Kruijff, E., LaViola Jr, J. J., & Poupyrev, I. (2004). *3D user interfaces: theory and practice*. Addison-Wesley.

Beyer, H., & Holtzblatt, K. (1997). *Contextual design: defining customer-centered systems*. Elsevier.

Carpendale, M. S. T., Cowperthwaite, D. J., & Fracchia, F. D. (1995, December). 3-dimensional pliable surfaces: For the effective presentation of visual information. In *Proceedings of the 8th annual ACM symposium on User interface and software technology* (pp. 217-226). ACM.

Carpendale, M. S. T., & Montagnese, C. (2001, November). A framework for unifying presentation space. In *Proceedings of the 14th annual ACM symposium on User interface software and technology* (pp. 61-70). ACM.

Chang, Y. L. B. (2012, November). Improving awareness of automated actions within digital tabletops. In *Proceedings of the 2012 ACM international conference on Interactive tabletops and surfaces* (pp. 283-288). ACM.

Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. *Perspectives on socially shared cognition*, *13*(1991), 127-149.

Dietz, P., & Leigh, D. (2001, November). DiamondTouch: a multi-user touch technology. In *Proceedings of the 14th annual ACM symposium on User interface software and technology* (pp. 219-226). ACM.

Elmqvist, N., Henry, N., Riche, Y., & Fekete, J. D. (2008, April). Melange: space folding for multi-focus interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1333-1342). ACM.

Endsley, M. R. (1988, May). Situation awareness global assessment technique (SAGAT). In *Aerospace and Electronics Conference, 1988. NAECON 1988., Proceedings of the IEEE 1988 National* (pp. 789-795). IEEE.

Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors: The Journal of the Human Factors and Ergonomics Society,* 37(1), 32-64.

Flanagan, J. C. (1954). The critical incident technique. *Psychological bulletin,51(4), 327.*

Furnas, G. W. (1986). *Generalized fisheye views* (Vol. 17, No. 4, pp. 16-23). ACM.

Furnas, G. W., & Bederson, B. B. (1995, May). Space-scale diagrams: Understanding multiscale interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 234-241). ACM Press/Addison-Wesley Publishing Co..

Forlines, C., Wigdor, D., Shen, C., & Balakrishnan, R. (2007, April). Direct-touch vs. mouse input for tabletop displays. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 647-656). ACM.

Gutwin, C., & Greenberg, S. (1996, April). Workspace awareness for groupware. In *Conference Companion on Human Factors in Computing Systems* (pp. 208-209). ACM.

Gutwin, C., & Greenberg, S. (2000). The mechanics of collaboration: Developing low cost usability evaluation methods for shared workspaces. In *Enabling Technologies: Infrastructure for Collaborative Enterprises, 2000.(WET ICE 2000). Proeedings. IEEE 9th International Workshops on* (pp. 98-103). IEEE.

Gutwin, C., & Greenberg, S. (2002). A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work (CSCW),* 11(3-4), 411-446.

Han, J. Y. (2005, October). Low-cost multi-touch sensing through frustrated total internal reflection. In *Proceedings of the 18th annual ACM symposium on User interface software and technology* (pp. 115-118). ACM.

Hancock, M., Carpendale, S., & Cockburn, A. (2007, April). Shallow-depth 3d interaction: design and evaluation of one-, two-and three-touch techniques. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 1147-1156). ACM.

Hancock, M., Ten Cate, T., & Carpendale, S. (2009, November). Sticky tools: full 6DOF force-based interaction for multi-touch tables. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces* (pp. 133-140). ACM.

Hancock, M., Ten Cate, T., Carpendale, S., & Isenberg, T. (2010, April). Supporting sandtray therapy on an interactive tabletop. In *Proceedings of the 28th international conference on Human factors in computing systems* (pp. 2133-2142). ACM.

Holtzblatt, K., & Jones, S. (1993). Contextual inquiry: a participatory technique for system design. *Participatory design: principles and practice. Hillsdale, NJ: Erlbaum.*

Mackinlay, J. D., Robertson, G. G., & Card, S. K. (1991, March). The perspective wall: Detail and context smoothly integrated. In *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 173-176). ACM.

Mikulecky, K., Hancock, M., Brosz, J., & Carpendale, S. (2011, November). Exploring physical information cloth on a multitouch table. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces* (pp. 140-149). ACM.

Ortego, F., Rishe, N., Barreto, A., & Adjouadi, M. (2013). Interaction with 3D Environments Using Multi-Touch Screens. In *Innovations and Advances in Computer, Information, Systems Sciences, and Engineering* (pp. 381-392). Springer New York.

Robertson, G. G., & Mackinlay, J. D. (1993, December). The document lens. In *Proceedings of the 6th annual ACM symposium on User interface software and technology* (pp. 101-108). ACM.

Rodden, T. (1991). A survey of CSCW systems. *interacting with Computers*,*3*(3), 319-353.

Sarkar, M., & Brown, M. H. (1994). Graphical fisheye views. *Communications of the ACM*, 37(12), 73-83.

Schmidt, K., & Bannon, L. (1992). Taking CSCW seriously. *Computer Supported Cooperative Work (CSCW),* 1(1-2), 7-40.

Shneiderman, B. (1996, September). The eyes have it: A task by data type taxonomy for information visualizations. In *Visual Languages, 1996. Proceedings., IEEE Symposium on* (pp. 336-343). IEEE.

Scott, S., Allavena, A., Cerar, K., Franck, G., Hazen, M., Shuter, T., & Colliver, C. (2010). Investigating tabletop interfaces to support collaborative decision-making in maritime operations. In *Proceedings of ICCRTS 2010: International Command and Control Research and Technology Symposium* (pp. 22-24).

Scott, S. D., & Carpendale, S. (2006). Guest editors' introduction: interacting with digital tabletops. *Computer Graphics and Applications, IEEE,* 26(5), 24-27.

Scott, S. D., Sheelagh, M., Carpendale, T., & Inkpen, K. M. (2004, November). Territoriality in collaborative tabletop workspaces. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work* (pp. 294-303). ACM.

Scott, S. D., Grant, K. D., & Mandryk, R. L. (2003, January). System guidelines for co-located, collaborative work on a tabletop display. In *ECSCW 2003* (pp. 159-178). Springer Netherlands.

Serenko, A. (2006). The use of interface agents for email notification in critical incidents. *International journal of human-computer studies*, *64*(11), 1084-1098.

Spence, R., & Apperley, M. (1982). Data base navigation: an office environment for the professional. *Behaviour & Information Technology,* 1(1), 43-54.

Tang, A., Tory, M., Po, B., Neumann, P., & Carpendale, S. (2006, April). Collaborative coupling over tabletop displays. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (pp. 1181-1190). ACM.

Tse, E., Greenberg, S., Shen, C., & Forlines, C. (2006, May). Multimodal multiplayer tabletop gaming. In *3rd International Workshop on Pervasive Gaming Applications, PerGames* (Vol. 6).

Wellner, P. (1991, October). The DigitalDesk calculator: tangible manipulation on a desk top display. In *Proceedings of the 4th annual ACM symposium on User interface software and technology* (pp. 27-33). ACM.

Westerman, W. (1999). *Hand tracking, finger identification, and chordic manipulation on a multi-touch surface* (Doctoral dissertation, University of Delaware).

Wilson, A. D., Izadi, S., Hilliges, O., Garcia-Mendoza, A., & Kirk, D. (2008, October). Bringing physics to the surface. In *Proceedings of the 21st annual ACM symposium on User interface software and technology* (pp. 67-76). ACM.

Yang, H., & Olson, G. M. (2002, September). Exploring collaborative navigation: the effect of perspectives on group performance. In *Proceedings of the 4th international conference on Collaborative virtual environments* (pp. 135-142). ACM.