

Supporting Spatial Collaboration: An Investigation of Viewpoint Constraint and Awareness Techniques

Wendy Schafer

**Dissertation submitted to the faculty of the Virginia Polytechnic Institute and
State University in partial fulfillment of the requirements for the degree of**

**Doctor of Philosophy
In
Computer Science**

**Doug A. Bowman, Chairman
John M. Carroll
Christopher L. North
Mary Beth Rosson
Robert W. Morrill**

**April 19, 2004
Blacksburg, Virginia**

**Keywords: computer-supported cooperative work, collaborative virtual
environments, spatial discussions, radar views, frames of reference, fisheye
projections**

Copyright 2004, Wendy A. Schafer

Supporting Spatial Collaboration: An Investigation of Viewpoint Constraint and Awareness Techniques

Wendy Schafer

Abstract

Spatial collaboration refers to collaboration activities involving physical space. It occurs every day as people work together to solve spatial problems, such as rearranging furniture or communicating about an environmental issue. In this work, we investigate how to support spatial collaboration when the collaborators are not colocated. We propose using shared, interactive representations of the space to support distributed, spatial collaboration. Our study examines viewpoint constraint techniques, which determine how the collaborators individually view the representation, and awareness techniques, which enable the collaborators to maintain an understanding of each other's work efforts. Our work consists of four phases, in which we explore a design space for interactive representations and examine the effects of different viewpoint constraint and awareness techniques. We consider situations where the collaborators use the same viewpoints, different viewpoints, and have a choice in viewpoint constraint techniques. In phase 1, we examine current technological support for spatial collaboration and designed two early prototypes. Phase 2 compares various two-dimensional map techniques, with the collaborators using identical techniques. Phase 3 focuses on three-dimensional virtual environment techniques, comparing similar and different frames of reference. The final phase reuses the favorable techniques from the previous studies and presents a novel prototype that combines both two-dimensional and three-dimensional representations. Each phase of this research is limited to synchronous communication activities and non-professional users working together on everyday tasks. Our findings highlight the advantages and disadvantages of the different techniques for spatial collaboration solutions. Also, having conducted multiple evaluations of spatial collaboration prototypes, we offer a common set of lessons with respect to distributed, spatial collaboration activities. This research also highlights the need for continued study to improve on the techniques evaluated and to consider additional spatial collaboration activities.

Acknowledgements

This dissertation is a result of many years of hard work. It was only possible with the help and support of many colleagues, friends, and family members. I wish to thank my advisor, Doug Bowman, who helped me grow personally and intellectually. He always asked the hard questions and was patient with my responses. My committee members John (Jack) Carroll, Mary Beth Rosson, Chris North, and Robert Morrill were also extremely influential. Their continual support and direction guided my research and facilitated the completion of my degree.

I also wish to thank my undergraduate advisor, Dr. Angela B. Shiflet. She introduced me to research and pushed me to accomplish more than I ever thought I could.

My family also provided their love and support in every step of the process. From the day I drove to Blacksburg in August of 1998 until graduation day, they have been a constant in my life. My parents may have lived 500 miles away over the course of the years, but they were always close by to hear my joys and troubles, and offer advice. Their house also provided the best kind of get away, very familiar, and yet, stress free. My grandparents, Mary and Tom Deusenbery and June Schafer, were also supportive as they helped me to remember what truly matters in life. Thanks also to my sister, Colleen, for her daily phone conversations, which made life interesting and enjoyable, especially when I had no drama.

The folks with the Center for HCI at Virginia Tech were also invaluable during the past years. Dan Dunlap was a mentor who shared an office space with me at the start of my graduate career. As we worked together on various research projects, Dan helped me to understand graduate life. Philip Isenhour also provided technical support that was crucial to my research. It was Philip's programming skills that allowed me to build my collaborative tools. He was also a terrific help with all of my unconventional Linux explorations.

I also wish to thank all the students and faculty who have worked with the Center in my time there. This group, along with the 3D Interaction research group, made this work possible. They listened to my practice talks, had numerous discussions with me, and offered suggestions and words of encouragement daily. Thanks to Jon Howarth, Con Rodi, Jochen Rode, Chad Wingrave, Kibum Kim, Dheva Raja, John Lucas, Jian Chen, Debbie Reese, Dennis Neale, and Jason Zietz. In particular, Craig Ganoie repeatedly talked with me about my ideas and helped with equipment configurations. He also showed me there was more to Southwest Virginia than just schoolwork.

Lastly, I wish to thank some of my good friends. When I began graduate school, Nathan Whitehead insisted that I could be a graduate student and that I could learn to program in Java. He was definitely right. My friends Rob and Miranda Capra and Tanya and Nathan Conklin were also very supportive. They reviewed my presentations, helped with the data analysis, as well as pulled me away to have some fun. Being with each of them is always a treat, as you can discuss HCI issues or just enjoy a good game of WizWar or Settlers. Most recently, Jonathan Laughnan allowed me to take time off, play, and enjoy life. He also critically read my writing, provided an alternative viewpoint, and offered suggestions.

To
Mom, Dad, and Colleen

Table of Contents

1	Introduction.....	1
1.1	Motivation.....	1
1.2	Collaboration Technology	2
1.3	Problem Statement and Hypotheses.....	4
1.4	Scope.....	5
1.5	Approach.....	8
1.6	Significance.....	8
1.7	Overview.....	9
2	Review of the Literature.....	11
2.1	Spatial Collaboration in Geographic Information Systems	11
2.2	Two-Dimensional Collaborative Solutions.....	16
2.2.1	Radar Views.....	16
2.2.2	Other Two-Dimensional Collaborative, Spatial Representations	17
2.3	Three-Dimensional Collaborative Solutions.....	19
2.4	Summary.....	23
3	Preliminary Work.....	24
3.1	Classification of Interactive, Map Software	24
3.2	Map-Based Navigation in MOOsburg.....	28
3.3	Collaborative, Spatial Discussion Prototype.....	31
4	Two-Dimensional Map Study	35
4.1	Study Details.....	35
4.1.1	Variables of Interest.....	35
4.1.2	Task.....	41
4.1.3	Interface	42
4.1.4	Experimental Design.....	44
4.2	Experimental Results	45
4.2.1	Final Questionnaire Results	45
4.2.2	Strategies of Use	47
4.2.3	Task Rating Sheets.....	49
4.2.4	Task Completion Times.....	51
4.2.5	Agreement Metric	53
4.3	Discussion of Results.....	53
4.3.1	Usability issues with discrete navigation.....	53
4.3.2	Approaches for collaborating with discrete navigation	54
4.3.3	Advantages of continuous navigation for collaboration	56
4.3.4	Other Navigation Possibilities	57
4.3.5	Negative effects of the fisheye, discrete combination	58
4.3.6	Performance and usability of the fisheye design	58
4.4	Implications.....	60
4.4.1	Issues in designing for spatial collaboration.....	60
4.4.2	Issues in designing distributed, spatial collaboration solutions.....	61
4.5	Conclusions.....	62
5	Three-Dimensional Virtual Environment Study	64
5.1	Spatial Collaboration Experiment.....	64
5.1.1	Variables of Interest.....	65

5.1.2	Task.....	68
5.1.3	User Interface.....	70
5.1.4	Software Design.....	71
5.1.5	Experimental Design.....	71
5.2	Experimental Results	72
5.2.1	Questionnaire Results	72
5.2.2	Awareness Question Results.....	76
5.2.3	Timing Results	76
5.3	Discussion.....	79
5.3.1	Actor and Director Frames of Reference	79
5.3.2	Frame of Reference Combinations	81
5.3.3	Awareness Information.....	82
5.4	Implications.....	83
5.5	Conclusions and Future Work	85
6	Real-World Application and Validation.....	87
6.1	Introduction.....	87
6.2	Preliminary Studies	87
6.3	Prototype	88
6.3.1	Interaction Techniques.....	92
6.3.2	Software Design.....	93
6.4	Qualitative Study	93
6.5	Observations	94
6.5.1	Validating the Previous Studies	94
6.5.2	General Collaboration Approaches.....	96
6.5.3	Strategies Employed	97
6.5.4	Uses of the Two-Dimensional Map	98
6.5.5	Uses of the Three-Dimensional Virtual Environment	99
6.5.6	Coordinated Activities Using Both Interfaces	100
6.5.7	Gesturing Issues	101
6.6	Implications.....	102
6.6.1	Realistic Setting	102
6.6.2	Multiple Representations	103
6.7	Conclusions.....	106
7	Conclusion	108
7.1	Contributions.....	108
7.1.1	Summary of Findings.....	109
7.1.2	Lessons Learned.....	110
7.2	Advances in Understanding	114
7.3	Future Work	116
8	References.....	119
Appendix A	Map software applications surveyed.....	125
Appendix B	Map-based Navigation Prototype Evaluation	131
B.1	Background Questionnaire.....	131
B.2	Evaluation Data Sheet.....	132
Appendix C	Spatial Discussion Prototype Evaluation.....	138
C.1	Evaluation Task	138

C.2	Questionnaire	138
C.3	Session Log.....	139
Appendix D	Two-Dimensional Map Study Materials.....	142
D.1	Background Questionnaire.....	142
D.2	Instructions.....	144
D.3	Task Rating Sheet	147
D.4	Final Questionnaire.....	150
D.5	Evaluation Data Sheet.....	152
Appendix E	Three-Dimensional Virtual Environment Study Materials	156
E.1	Background Questionnaire.....	156
E.2	Instructions.....	158
E.3	Task Survey	159
Appendix F	Real-World Application and Validation Materials	162
F.1	Evaluation Data Sheet.....	162
Appendix G	Vita	165

List of Figures

Figure 1.1	Time space taxonomy.....	3
Figure 2.1	Spatial Group Choice	12
Figure 2.2	Argumentation Maps	13
Figure 2.3	Geovisualizations	14
Figure 2.4	Interactive, Web-based GIS.....	15
Figure 2.5	GroupARC.....	16
Figure 2.6	Radar view.....	17
Figure 2.7	Ramonamap.....	18
Figure 2.8	ActiveMap.....	19
Figure 2.9	Augmented Reality examples.....	20
Figure 2.10	Collaborative Architectural Layout Via Immersive Navigation (CALVIN)	21
Figure 2.11	Collaborative navigation in a CVE.....	22
Figure 3.1	Interactive, map software classification.	27
Figure 3.2	MOOsburg.....	29
Figure 3.3	Zooming with the map-based navigation tool.....	30
Figure 3.4	Clicking with the initial map-based navigation prototype	30
Figure 3.5	Spatial discussion prototype	33
Figure 4.1	Discrete navigation technique	37
Figure 4.2	Continuous navigation technique	38
Figure 4.3	Fisheye projections.....	39
Figure 4.4	Traditional radar view	40
Figure 4.5	Fisheye radar view.....	40
Figure 4.6	Multiple user fisheye radar view	41
Figure 4.7	Two-dimensional map study prototype.....	44
Figure 4.8	Perceived effort interaction	50
Figure 4.9	Navigation support interaction	51
Figure 4.10	Time to position a light marker	52
Figure 4.11	Time to position a sign marker.....	53
Figure 4.12	Discrete navigation modification	57
Figure 5.1	Exocentric viewpoints	67
Figure 5.2	Egocentric, actor condition.....	69
Figure 5.3	Exocentric, actor condition.....	70
Figure 6.1	Collaboration with the virtual environment representation.....	89
Figure 6.2	Manipulation with the two-dimensional map representation	90
Figure 6.3	Observing manipulations with the virtual environment representation	91
Figure 6.4	Lab design product	98
Figure 6.5	Organizing objects with the two-dimensional map.....	99
Figure 6.6	Exploring object positions with the virtual environment	100

List of Tables

Table 4.1	Navigation and radar view ranking results	46
Table 4.2	Subjective navigation results	46
Table 4.3	Subjective radar view results	47
Table 4.4	Participant effort differences	47
Table 4.5	Radar view and marker strategy results	48
Table 4.6	Navigation and marker strategy results	48
Table 4.7	Complex task and pointing strategy differences	49
Table 4.8	Confirmation action results	49
Table 4.9	Task completion time results	52
Table 5.1	Frame of reference, role conditions	72
Table 5.2	Collaboration rating statements	72
Table 5.3	Mean ratings for collaboration statement 2.	73
Table 5.4	Perceived effort statements	73
Table 5.5	Mean ratings for the perceived effort statements.....	73
Table 5.6	Mean ratings for perceived effort statement 1.....	74
Table 5.7	Mean ratings for perceived effort statement 3.....	74
Table 5.8	Awareness questions.....	75
Table 5.9	Mean ratings for awareness question 1	75
Table 5.10	Mean ratings for awareness question 2.....	76
Table 5.11	Mean ratings for awareness question 3.....	76
Table 5.12	Mean times for selection.....	77
Table 5.13	Mean times for selection and coarse positioning.....	77
Table 5.14	Mean times for coarse positioning.....	78
Table 5.15	Mean times for fine positioning.....	78
Table 5.16	Mean times for positioning	79
Table 5.17	Mean times for task completion.....	79

Chapter 1

1 Introduction

1.1 Motivation

Spatial collaboration is a specialized form of collaboration where the discussion relates to a physical place (Schafer & Bowman, 2003). Every day, people engage in complex activities and projects that require them to think spatially. Such thinking is a consistent element in any given individual's daily decision-making process, from deciding which route to take to work, and engaging in hobbies such as hiking and traveling. In many cases, people collaborate in groups to solve problems related to spaces. The task might be a question of getting from point A to point B in the middle of rush hour traffic, arranging the existing furniture in a new office area so that each coworker is content, or ensuring proper maintenance of a set of hiking trails by the members of an outdoor club. Each of these activities involves individuals thinking about spatial relations and communicating with others to achieve a desired goal.

Spatial collaboration describes the combined work of multiple participants to solve problems involving physical space. This physical place can vary in size, but it is typically a large space people occupy and interact with. For example, the place might be a building, such as a house or a multiple-story office space. In these cases the collaboration activity might be focused on discussing floor plan layouts and the use of various rooms. Spatial collaboration could also relate to an outdoor area, such as a nature park or an intersection of roads. For example, environmentalists and highway designers may need to work together to achieve their goals for an area of land. Conditions for spatial collaboration vary can according to the goals and objectives of each activity, the number of people collaborating, each person's role in the effort, each person's level of involvement and time commitment, and each person's physical location. The common factor in all of these situations is that the people involved are focused on a spatial problem concerning a large, physical area that can be physically inhabited.

Spatial collaboration is a complex activity. Thorndyke and Hayes-Roth (1982) discuss the different types of spatial knowledge people use. These include knowledge of landmarks, routes, and map-like information. People also have different mental models of the same physical space. Important landmarks for one person may be insignificant to another, and people may differ in their knowledge about various areas within space. These individual differences make it difficult to exchange ideas about a static space, not to mention a dynamic one. In fact, collaboration problems due to individual differences are common. For example, people get lost trying to follow someone else's directions, they disagree about navigation when reading a map together, and commonly have discussions with different outcomes about where things are located.

Spatial collaboration is also difficult because many of the ideas relate to a particular spatial environment. These ideas need to be understood with respect to this environment, including its structure and scale. For example, an idea might involve location information, such as where a park is located within a town, or where to position a piece of furniture within a room. This location information makes sense within the context of the environment. The information is

directly related to a particular place and it must be understood with respect to the place. As such, in the example about a park's location, or the position of a piece of furniture, the collaborators can discuss one another's ideas because they have a similar understanding of the town's structure or the room's layout.

Ideas exchanged in a spatial collaborative activity also use notions of distance and direction. When discussing a spatial environment it is often important to convey distance. Determining the shortest route between two locations is one example where distance comes into play. One can consider the straight-line distance between two places, or a complex path from one location to another. In both cases, an understanding of the scale of the environment allows us to comprehend the distance measurements. Likewise, the ideas exchanged in spatial collaboration also use directional concepts. For instance, consider a road that primarily runs north and south, or a baseball field that is to the left of the soccer field, or a couch that is centered along a wall. All of these statements provide directional information and use a specific reference frame in their description. In the baseball field example, "to the left" is understood with respect to certain position and orientation. Directional information also varies in scale and dimensionality. Objects can be positioned with scales varying from miles to millimeters. The information can also use multiple dimensions such as a three-dimensional location and orientation. This range of possibilities complicates the process of spatial collaboration.

To summarize the issues of spatial collaboration, consider an actual example of a local community group interested in preserving and restoring its local watershed. The group consists of three leaders from three different organizations that coordinate unique watershed-related activities. One leader is an assistant director for the local natural history museum, another a director of the water research center at the local university, and the last is the director of the local children's exploratory camps. Each leader is familiar with different areas of the watershed and examines different aspects of the environment. The museum organizes biological sampling efforts to investigate invertebrate counts at designated sites. The water research center coordinates research projects to classify and catalogue watershed-related issues along the primary stream's path. The children's camps director schedules in learning activities and trash pickup efforts at select sites.

In order to work together, the leaders share information and communicate their concerns for the watershed. The leaders have a common understanding of the stream, but they often negotiate individual differences in spatial knowledge. They reference locations familiar to everyone in the group and they explain their ideas by describing locations along the primary stream. In their conversations, they also discuss the distance along the stream between locations, and reflect on the direction and shade sources for the water flow. The stream is a complex spatial entity that winds its way through different terrains and takes on different forms. As a result, the communication of the group involves the many spatial aspects of the watershed.

1.2 Collaboration Technology

Technology can be used to support and enhance spatial collaboration activities. It can aid collaborators in establishing a common understanding, allow them to be more specific in their references and measurements, and even help them in making decisions. The recent use of

collaborative software is one example of such technology. It allows people to go beyond using a computer individually, providing a means to communicate and share ideas with others. The research field of computer-supported cooperative work (CSCW) investigates this emerging field. It examines how cooperative work is conducted, explores the design and development of collaborative systems, and studies the effects of technology on collaborative activities. Ellis, Gibbs, and Rein offer an application-level taxonomy of collaboration software (1991). Their categorization is based on where and when people use the software and uses a time space taxonomy (Figure 1.1). Collaboration can occur in the same location through co-located interaction, or from different locations through distributed interaction. It also occurs either at the same time synchronously, or across different times asynchronously.

	Same Time	Different Time
Same Place		
Different Place		

Figure 1.1 Time space taxonomy

This taxonomy is one way to classify collaboration software.

Particularly interesting are the situations where people are located in different places. In these cases, the collaborators are typically not located within the same room, requiring them to rely on technology to communicate. Referred to as computer-mediated communication, the participants must use the technology to explain their ideas and learn of the other participants' ideas. This exchange of ideas very often lacks the nonverbal cues of face-to-face communication. Unless a video-sharing technique is configured, the participants will not be able to see each other. Not visually seeing someone drastically changes the style of communication, and even with video support, the shared images cannot reproduce face-to-face contact (Egido, 1990; Fussell, Kraut, & Siegel, 2000). Movements such as head nodding and hand gesturing will be much less useful and subtle cues with eye positioning and body posture will often go undetected. Supporting effective communication without these nonverbal signals is a challenge taken on by distributed technology.

Distributed collaboration is a common occurrence as people often work in different locations. Technology, such as email, telephone, instant messaging, and video conferencing, is frequently used to communicate between multiple people in multiple locations. Numerous software tools have also been written to support distributed collaboration and many use a graphical user interface to enable the collaborators to share information. Often referred to as groupware, these systems enable multiple people in different locations to interact with an interface as they work toward a common goal (Ellis, Gibbs, & Rein, 1991). For instance, collaborative editors support multiple authors composing a single document, and synchronized web browsing allows multiple collaborators to navigate online together (e.g. Greenberg & Roseman, 1996; Leland, Fish, & Kraut, 1988). Other collaborative applications allow multiple collaborators to share a calendar (e.g. Palen, 1999), write web pages together (e.g. Guzdial, Rick, & Kerimbaev, 2000), and explore a three-dimensional virtual environment together (e.g. Hindmarsh, Fraser, Heath, Benford, & Greenhalgh, 2000). Some of these tools support synchronous work, while others are designed for asynchronous interaction. Each one tries to address the challenges of distributed collaborators.

1.3 Problem Statement and Hypotheses

The field of CSCW explores the issues and challenges of collaboration. It provides many technology-based solutions for collaboration activities. However, it does not directly address the issues of spatial collaboration and, in particular, the difficulties of distributed, spatial collaboration.

When the participants are at different locations, the issues of spatial collaboration are further complicated. To express ideas, people frequently use physical cues such as pointing and gesturing to convey location and direction information. This nonverbal communication can be instrumental in establishing a common understanding, especially as people differ in spatial knowledge and abilities. In a distributed setting, resolving individual differences becomes more problematic without these natural movements. The collaborators have to find another way to convey their thoughts so that others clearly understand them. Typically, this involves a negotiation with spatial references to create a common ground.

Similarly, relating ideas to a physical place becomes more complex. If two people are physically located at the place of interest, the environment around them establishes a spatial context for their work. For example, in designing a new building, a visit to the intended location will allow people to walk around the site, point out features, and discuss the location. Such positioning within the spatial context enables the ideas to be put in perspective with the environment. On the other hand, if the individuals are located in different places with the same task of designing a new building, the spatial context must be established through a different means. This might be achieved through an intricate discussion of memories of a site, or having each collaborator look at a similar representation of the space, such as a map.

Current collaborative software does not address these challenges of distributed, spatial collaboration. To begin with, some collaboration technology does not support the spatial aspects of the task. This software requires the distributed participants to establish the spatial context through speech or text. Audio conferencing, instant messaging, and discussion lists are prime examples of technology that require establishing such context. Other collaboration solutions offer limited support for exchanging ideas about space. They provide representations, or visualizations, of the physical space that provide context to activities, but lack a way for collaborators to converse and express thoughts spatially. Web pages with graphical images, for example, typically only present an idea. Most Web sites allow an author to post a picture, but they do not support an exchange of ideas among multiple people. Emails with image attachments also have problems. One collaborator can send a message to others with an image, but numerous messages and attachments are required to engage in a spatial discussion. Email can also be slow, and communication is typically between two individuals, rather than a group of people.

Looking at current groupware solutions, this software includes graphical representations in its user interfaces. Groupware usually supports interacting with the representation, very often allowing the distributed users to share the changes they make. For instance, when drawing on shared whiteboard, this action is mimicked on every other participants' whiteboard. However, current groupware lacks the features to support many spatial collaboration activities. With shared whiteboards in particular, most of the interaction techniques are tailored toward drawing.

These tools often do not support discussions of existing spaces, investigations of multiple types of spatial data, or explorations of distances. Spatial collaboration encompasses these more complex activities, and it requires distributed solutions that enable participation in a variety of spatial problems.

Given this state of collaborative software, the problem statement for this work can be summarized as follows:

Spatial collaboration is a common and everyday task. Current collaboration research does not address the issues inherent in distributed, spatial collaboration. Its use of representations does not facilitate the diverse exchange of ideas within a spatial context that is characteristic of spatial collaboration activities.

Hence, there is a need to explore new collaborative solutions to address the issues of distributed, spatial collaboration. One way is to expand on the design of single-user, spatial software. These technologies typically offer an interactive representation of the space, enabling the user to reference features and not rely on his/her individual, spatial knowledge. For example, in providing directions, a person might use current map software to look up the precise location of a road, as opposed to determining the location from his/her spatial, mental map. Such software also provides a range of features and tools supporting spatial tasks that could be useful to spatial collaboration problems.

Distributed, spatial collaboration solutions can also extend the ideas of groupware. Interface designs can include a representation that provides a spatial context. This visualization enables the collaborators to explain their ideas through referencing. They can point out elements in the representation to one another, and they can discuss spatial layout issues together (Churchill & Snowden, 1998). Integrating both the interactive features of current spatial software and the collaborative features of groupware will create a shared, interactive representation. Such a representation can support discussions involving spatial concepts of distance, orientation, and location information, all of which are difficult to convey using only a verbal form of communication such as text or speech.

This work incorporates three general hypotheses relating to this problem statement:

- Distributed, collaborative solutions can be designed to support spatial collaboration activities.
- Shared, interactive, representations, similar to those found in other groupware systems, can enable distributed, spatial collaboration.
- An assessment of representation and interaction techniques will help guide the design of spatial collaboration solutions.

1.4 Scope

The purpose of this work is to initiate an assessment of representation and interaction techniques for distributed, spatial collaboration. It does not provide a complete analysis, but reveals the advantages and disadvantages of different features with respect to specific spatial collaboration tasks. This study sets the stage for future spatial collaboration work. It encourages the analysis

of other features as well as inspires the design and evaluation of new representation and interaction techniques tailored for spatial collaboration.

This work investigates techniques that support distributed collaborators sharing a software representation. More precisely, the research examines interactive representations that support synchronous, distributed work. There are multiple ways a spatial entity can be represented in a software solution; some representations can be more understandable and usable than others. One way representations can differ is in terms of dimensionality. They can be two-dimensional (2D), such as a map, or three-dimensional (3D), such as in virtual reality. They can also implement different ranges of possible viewpoints and vary in how much control the user has in switching between these viewpoints. Also, representations can include different awareness features that allow multiple participants to maintain an understanding of each other's work effort.

In distributed collaboration, one can also consider how the individual collaborator's displays differ. One technique is to have each person use the same type of representation. For example, two people might view a similar 2D representation, or similar 3D representation, where the display supports the same range of viewpoints. Alternatively, the spatial representations could vary with the collaborators. The displays could use different viewpoint techniques, different dimensionality, or different awareness techniques. For instance, instead of using the same 3D viewpoints, two collaborators might navigate the space from different perspectives. Another example occurs when one participant interacts with a 2D representation, while another has a 3D representation. A better understanding of the advantages and disadvantages of these different techniques for representing a spatial entity will better inform the design of spatial collaboration tools.

Within the area of spatial representations, this work focuses on investigating viewpoint constraint techniques and awareness techniques. Viewpoint techniques affect how the users view the representation. Constraints on these techniques determine the available viewpoints of a representation. In a spatial collaboration solution, viewpoint constraint techniques need to support the collaborators in their exploration of the space. The techniques employed should enable them to make sense of the representation and understand it with respect to the real-world place. They should also provide the spatial views that will be most helpful in solving the spatial problem at hand. The issues with viewpoints are further complicated when one considers multiple users engaging in collaborative navigation. A useful and usable viewpoint technique will recognize individual differences and encourage the group to establish and maintain a common understanding throughout the activity.

In addition to viewpoint constraint techniques, this work investigates the familiar collaboration issue of social awareness. In distributed collaboration, participants need to be aware of each other's activities and work effort. Gutwin and Greenberg describe synchronous, groupware awareness in terms of knowing about the people participating in the collaboration, including who they are, what they are working on and the specifics of their locations and actions (Gutwin & Greenberg, 2002). This information is widely available in face-to-face settings, but is less available in distributed situations. Awareness techniques provide this crucial information in distributed software applications (Dourish & Bellotti, 1992; Gutwin & Greenberg, 1999; Gutwin, Roseman, & Greenberg, 1996). In spatial collaboration activities, these techniques can provide

information about how collaborators have navigated independently, which spatial features each participant is focused on, and any actions the users have performed to manipulate the representation. A useful and usable awareness approach will enable the collaborators to explore different areas of the space, yet maintain group focus and coherence. It will also enable the group members to understand each other's ideas as they closely work together in different areas of the space.

Combining different awareness and viewpoint constraint techniques also warrants consideration. The awareness technique should complement the viewpoint style, so that it is clear where each collaborator is working. A complex viewpoint technique should be combined with appropriate awareness techniques, so that the awareness information is clear and easy to comprehend. If the awareness technique is less useful, this creates more collaboration issues than it resolves. Users will struggle to understand the awareness information, causing them to question their collaborators work effort, possibly distracting them from the real collaborative task.

With this focus on viewpoint constraint and awareness techniques, this work is interested in the individual and combined effects on spatial collaboration. Specific techniques can cause collaborators to use different strategies and take different approaches with spatial problems. This research examines how various techniques affect the participants' problem-solving process. Similarly, it evaluates the product of the collaborative effort. It questions whether certain techniques lead to greater efficiency, greater precision, or a more complete solution. Lastly, this work considers the collaborators' experiences. It asks how satisfied each collaborator and the group as a whole is with the collaboration process and their solution to the problem.

This research limits the field of spatial collaboration by focusing on synchronous work. Spatial activities can occur asynchronously, but this work concentrates on the real-time situations that parallel traditional, face-to-face communication.

This work also focuses on common collaboration tasks completed by knowledgeable, but not necessarily technical users. The approach is to explore the use of spatial representations to enhance informal, everyday spatial collaboration tasks. These tasks involve non-professionals working together at a high level of detail. Friends who discuss directions to a place in town, or colleagues who are rearranging office furniture are prime examples. This focus is in comparison with more technical users such as geographers or scientists, who would require more sophisticated representations and functionality.

Lastly, in its investigation of viewpoint constraint and awareness techniques, this research is focused on supporting both individual and shared efforts (Dourish & Bellotti, 1992; Gutwin & Greenberg, 1998). The thought process of individual participants often involves expressing ideas as well as understanding their partners' ideas. To support this mixed-focus collaboration, each user navigates the representations individually and can perform actions not visible to the others, while awareness techniques are used to inform the collaborators about the group effort. For example, a group of people designing a town park from distributed locations might want each person to have an individualized map of the park, where everyone sees the same set of participant-created annotations. This enables each collaborator to explore the space individually, while the shared annotations act as an awareness technique. They allow each participant to

understand where others are working and how the discussion is progressing. The spatial representations in this work use this approach of allowing users to explore the space independently, which is termed *relaxed what you see is what I see* or relaxed WYSIWIS (Stefik, Bobrow, Foster, Lanning, & Tatar, 1987).

1.5 Approach

This research specifically addresses two questions. First, how do different viewpoint and awareness technique combinations affect spatial collaboration activities when the collaborators use the same techniques? And secondly, what are the effects when the collaborators use similar and different viewpoint constraint techniques? The questions examine two aspects of spatial representation variability. The first focuses on differences between representations. By providing the collaborators with the same representation, it investigates the direct effects of specific viewpoint constraint techniques and specific awareness techniques. The second question, on the other hand, focuses on how the individual collaborator's displays differ. It explores the advantages and disadvantages of having the same and different viewpoints into a representation.

These research questions are answered through a four-stage approach:

1. This first phase explores a design space for interactive representation techniques. Surveying interactions supported in Web-based and commercial map programs, it reveals techniques that can be applied to spatial collaboration.
2. This phase uses two-dimensional spatial representations to address the first research question. It compares both common viewpoint techniques and typical and novel awareness techniques using a secondary display. It also explores the interactions between the viewpoint and awareness techniques.
3. This phase looks at a different type of spatial representation with three-dimensional, collaborative virtual environments. It compares viewpoint constraint combinations and explores the awareness effects of the combinations. Collaborators use similar and different viewpoint constraint techniques, and awareness information is integrated into the spatial representation.
4. This phase uses a more informal approach, exploring distributed spatial collaboration in a realistic setting. It examines the favorable techniques from the previous studies to validate the findings. It also incorporates a novel solution that adds new insight into spatial collaboration design.

1.6 Significance

This work furthers the understanding of distributed, spatial collaboration and how to support it through technology. It distinguishes between different viewpoint constraint and awareness techniques, providing guidelines for their use and offering direct comparisons between techniques. These findings are not specific to this work and are applicable to other distributed, spatial collaboration problems. This work also raises additional considerations in designing for distributed, spatial collaboration. Through the design and development of novel solutions and in observations of small groups performing distributed, spatial collaboration tasks, it informs the design of future tools. This approach provides both an analysis of the novel techniques and a better understanding of spatial collaboration processes.

This work also highlights the need for a continued study of designing spatial collaboration solutions. It focuses on a small, yet important, aspect of using spatial representations. This prompts a further investigation of representation and interaction techniques. Possibly, there are more useful ways for representing the space, or more creative interactions that enable collaborators to share ideas. Such an investigation can lead to the design and evaluation of novel approaches for spatial collaboration that address the issues unresolved in current solutions.

This research directly contributes to the fields of Computer-Supported Cooperative Work. It provides a rich description of distributed, spatial collaboration, the issues inherent in designing solutions, and some initial approaches to the problem. It also offers new insight for the Collaborative Virtual Environments community. This work furthers an understanding of the use of viewpoint constraint techniques in a virtual environment, and presents a unique design for supporting a three-dimensional, collaborative task.

1.7 Overview

This introductory chapter has described the topic of spatial collaboration and its complexity. Reflections on the current research field of Computer Supported Cooperative Work emphasize current distributed, collaboration solutions and their failures to address the needs and issues of spatial collaboration. This research explores this problem by investigating viewpoint constraint and awareness techniques that support shared representations. Most of the remaining chapters of this work correspond to the approach, which includes four stages.

In Chapter 2, continues with a review of CSCW research and presents related work. This discussion includes previous literature that relates to spatial collaboration, as well as prior two-dimensional and three-dimensional collaborative solutions that incorporate spatial representations, although not necessarily for spatial collaboration.

Chapter 3 corresponds to the first stage of this research. It describes preliminary studies conducted that relate to the topic of spatial collaboration. A review of a design space for spatial representations and an initial prototype to support distributed, spatial discussion offer motivation for second phase of this work.

Chapter 4 presents the second phase. Through a detailed description, it provides a formal comparison of two viewpoint constraint techniques and two awareness techniques. The reasoning for the comparisons, the experimental design, and the results are presented, including a discussion of the implications of the findings.

Chapter 5 follows with the third phase of the research. It describes a formal study of similar and different viewpoint constraint techniques and their effects on a three-dimensional, spatial collaboration task. The chapter presents the results from the study and discusses the implications for supporting distributed collaboration.

The fourth phase, covered in Chapter 6, builds on the findings from the experiments. It validates the results in a real-world application setting and explores the combined use of two-dimensional and three-dimensional representations.

The final chapter considers the previous chapters and draws conclusions. It summarizes the main findings and contributions, offering possibilities for future research.

Chapter 2

2 Review of the Literature

Spatial collaboration is a relatively new area of research. Related prior work exists in the fields of Geography, Computer Support Cooperative Work, and Collaborative Virtual Environments. The next few sections discuss the relevance of each of the fields to distributed spatial collaboration, in general. Additional literature is also presented with the individual studies to describe the closely related research (Chapters 4, 5, and 6).

2.1 Spatial Collaboration in Geographic Information Systems

Most of the past literature dealing with spatial collaboration relates to the field of Geographic Information Systems (GIS). These systems are typically used to analyze spatial data sets in solving domain specific problems. For example, a recent water main break might be explained by simultaneously displaying underground pipelines and various water sources. These systems focus on details, precision, and static visualizations of different data sets. The approach of this work, in comparison, is to explore the use of interactive spatial representations to enhance informal, everyday spatial collaboration tasks. The tasks involve knowledgeable, but not necessarily technical people working together on a spatial problem.

GIS literature has investigated both same place and distributed spatial collaboration with a focus on three types of tasks. These include collaborative spatial decision-making, collaborative spatial planning, and collaborative data analysis. With respect to same place spatial collaboration, Piotr Jankowski and Timothy L. Nyerges have explored applying decision support solutions to collaborative spatial decision-making (Jankowski, Nyerges, Smith, Moore, & Horvath, 1997). Their product, Spatial Group Choice (commercially sold as GeoChoicePerspectives), allows multiple people to make decisions about a fixed set of sites. For example, one of their scenarios has a variety of stakeholders prioritize habitat sites for restoration. Decisions are made using a range of numerical procedures including weighting criteria and ranking different options. Each person works with a separate computer to input his or her numerical contributions through a decision support module and explore representations through a separate, single user module. This meeting support approach is based on their task model that describes how groups can use GIS-based tools for selecting a restoration location. The model expresses a relationship among a specific selection problem, various tool capabilities and features, and a decision-making process. It incorporates a task analysis of other restoration selection activities and a theory of how GIS-collaboration tools will alter a small group's interactions and decision-making strategies (Nyerges & Jankowski, 1997).

This GIS work is an interesting example of how decision-making can be enhanced with same-place collaboration features. In particular, it demonstrates one way to support spatial collaboration through models of decision-making and a corresponding tool that uses ranking techniques and a single user representation. The representation enables each user to explore the initial available data about the problem and spatially view the results during various stages of the decision-making process. In this way, the representation acts as a decision aid, and it is not a vital element of the activity. The work reported in this thesis focuses on a more direct

involvement of the representation in spatial collaboration. It also investigates less technical tasks in a distributed setting.

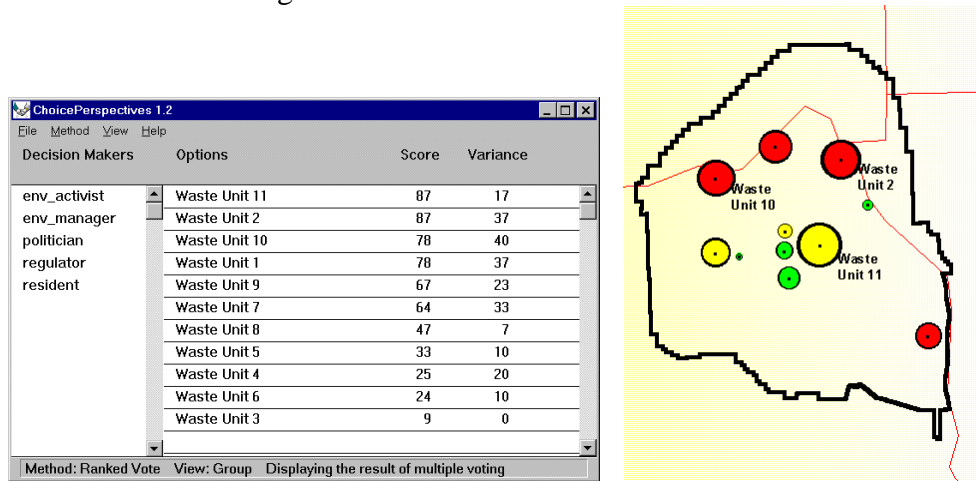


Figure 2.1 Spatial Group Choice

This commercial software facilitates spatial decision-making among colocated users. It focuses on the use of rankings and weighted criteria, allowing users to view the results with spatial representations.

A number of different GIS research projects have also investigated distributed spatial collaboration. In Claus Rinner’s dissertation, “Argumentation Maps – GIS-based discussion support for online planning,” he realizes the need to support group discussions through linking individual contributions to map elements (Rinner, 1999). Working with a land-use planning application, Rinner analyzes how GIS and Argumentation Theory can be combined to support asynchronous, distributed debates. He offers a simple object-oriented model for a geographically referenced discussion, as well as proposes a set of use cases for the interactions and discusses the theoretical implications for spatial planning. An initial prototype presents the idea of displaying arguments using a hierarchy and a spatial representation. This technique encourages the exploration and analysis of planning ideas spatially. Each argument is associated with a spatial location. The spatial representation indicates the distribution of multiple arguments. The representation supports zooming and panning operations, but no awareness features if two people happen to be using it at the same time.

Rinner’s work is another example of a spatial collaboration solution that uses GIS concepts. His combination of Argumentation Theory and geographical representations presents a novel approach for asynchronous, spatial planning. The work reported in this thesis differs in that it studies synchronous communication in a different task domain. Also, Rinner’s dissertation does not include an evaluation of the prototype, while this thesis evaluates multiple representation techniques.

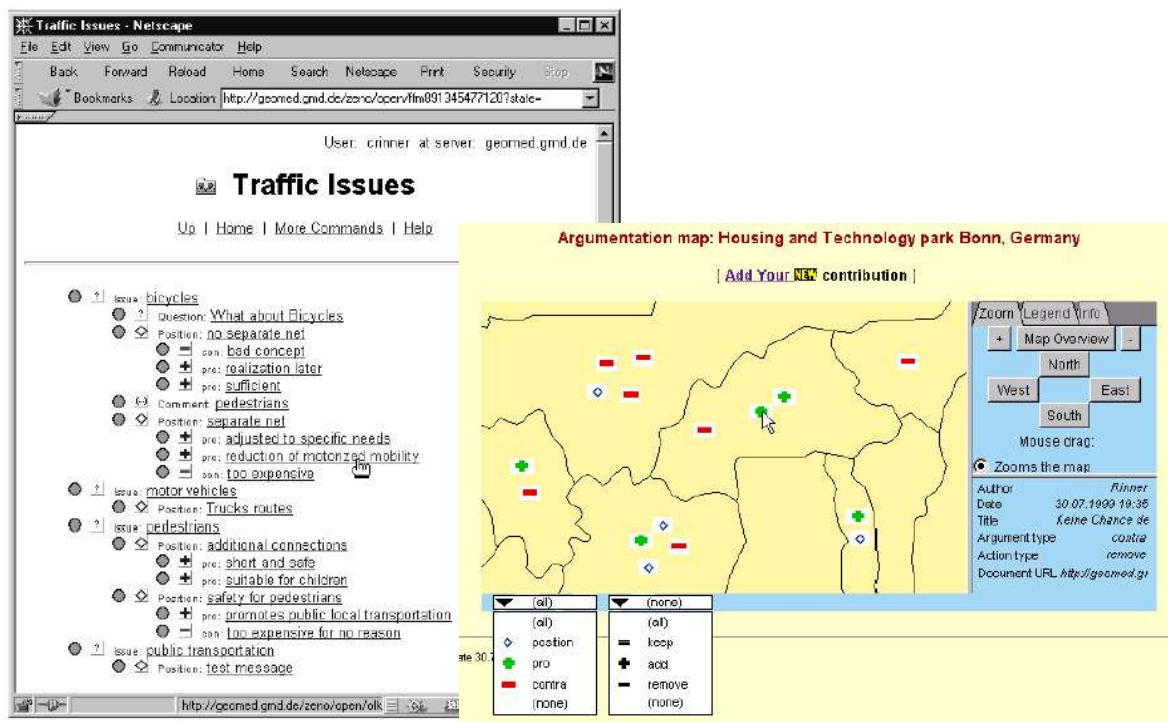


Figure 2.2 Argumentation Maps

Claus Rinner's prototypes support distributed, asynchronous debates by linking arguments with spatial data objects.

Another example of distributed, spatial collaboration is the research of the GeoVISTA Center at Penn State University. They are exploring the use of geovisualizations in distributed work (MacEachren, Brewer & Steiner, 2001). Geovisualizations, or Geographic Visualizations, use visual representations to make spatial contexts and problems visible to scientists. Displaying these representations across multiple locations enables experts to share ideas about geographic spaces. For example, an interactive, 3D view of the Susquehanna River Basin in North America will allow distributed research teams to understand climate and drainage patterns. An initial prototype has been developed to discuss important features with potential users, revealing the need for a variety of awareness and communication techniques (Brewer, MacEachren, Abdo, Gundrum, & Otto, 2000). One issue they hope to address in a second prototype involves communicating who is currently controlling the data displayed and what views each collaborator has available on their desktop. This prototype, called a watcher window, uses a table to indicate each user's view of each representation tool. For example, a geovisualization system may include a map and a spatial scatter plot, each with a different view of the same space. This prototype is still under development and has not been evaluated.

The focus of the GeoVISTA Center differs from this research with respect to the users and the tasks. Its interest lies in using geographic visualizations for science applications. This work, on the other hand, is focused on common spatial tasks performed by non-professionals. The GeoVISTA Center's encounter with awareness problems is also important to note. The watcher window is an interesting technique for providing spatial awareness information. The work

presented in this thesis differs in that it is interested in supporting awareness within a single representation as opposed to multiple representations of the same space.

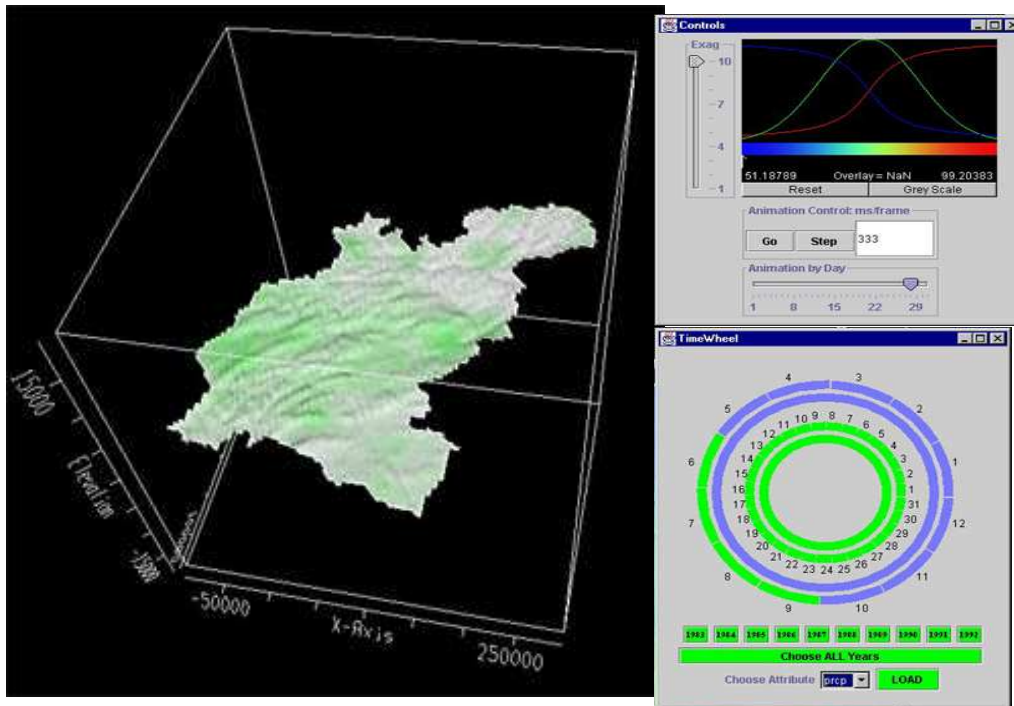


Figure 2.3 Geovisualizations

The GeoVISTA Center envisions distributed research scientists solving geographic problems together using three-dimensional representations.

The Centre for Computational Geography at the University of Leeds in Britain has taken a different approach toward distributed, spatial collaboration. It investigates the idea of increasing public participation in democracy through Web-based GIS (Evans, Kingston & Turton, 1999). Many governments provide information on the Web for the public to view and respond to, but spatial information has been restricted to simply exploration without a response mechanism. As a result, the researchers at the Centre conducted a test-bed project that required local inhabitants contribute town improvement ideas online. Each idea was positioned on a Web-based map using an applet. Comparing this approach to a centrally located, physical, three-dimensional model of the town where people added pushpin flags, revealed that the Web-based system was useful and popular. The Centre's work offers a limited solution with respect to the vision of spatial collaboration discussed in this thesis. They explore a way for the public to provide anonymous, spatial feedback to their government, while the work presented in this thesis is interested in more intricate scenarios of people solving spatial problems together using synchronous communication.

Slaitthwaite Virtual Decision-Making

The following comments were made:

We think that the traffic problem in carrs road ie lorries and other hqvs could be re-routed away from the village center by an alternative road which crosses the river further down towards huddersfield directly on to manchester road. When the canal project is being constructed we envisage terrible congestion due to the heavy traffic which has to travel along carrs road to the industrial estate. Our second point is the appearance of the shop fronts in carrs road. Their aesthetic appearance is not conducive to how we perceive the village once the canal project has been completed. The whole of the shop fronts that line carrs road and line the route of the canal really ought to be refurbished in some consistent style which fits in the canal ambience. Once this has been done perhaps the canal and the shops could be brought together with cobbles.

[Return to previous screen](#)

[Finish](#) [Reset](#)
[View Comments Page](#)

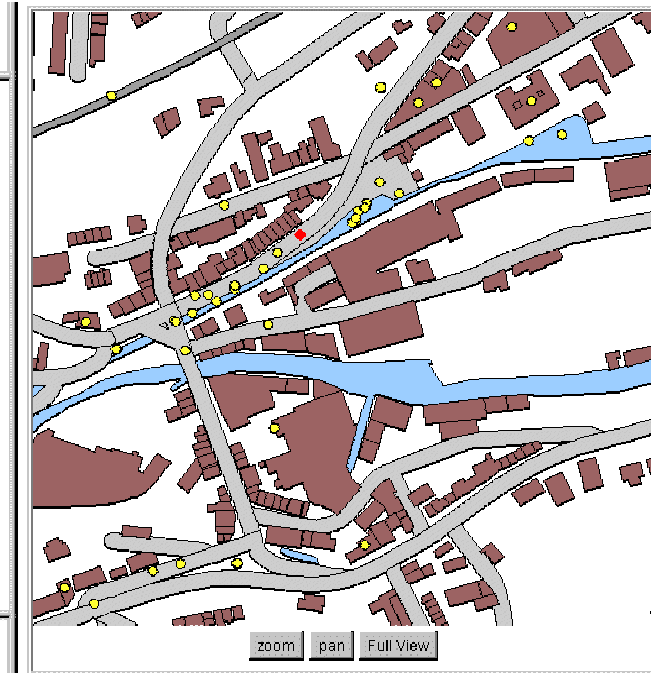


Figure 2.4 Interactive, Web-based GIS

The Centre for Computational Geography investigates two-way, spatial communication between a government and its people, by allowing residents to add comments to a town map.

Lastly, Churcher and Churcher have recognized the need for collaborative GIS applications (Churcher & Churcher, 1996). Through a prototype called GroupArc, they provide a way for physically separated users to concurrently browse and annotate GIS data. The tool demonstrates a union of computer-supported cooperative work and GIS approaches. Using a simplified GIS, it allows users to toggle layers of data on and off, display metadata for spatial features in the representation, and execute queries to filter the spatial data set. Using a groupware toolkit, it also allows users to highlight specific features and make gestures with cursors displayed both locally and on remote users' screens (Churcher & Churcher, 1999). This work offers some features that could also be useful to a spatial collaboration solution. The GIS tools will likely be too complex for informal collaboration tasks, but the use of groupware tools with a spatial representation are desirable. The work presented in this thesis builds upon these ideas by investigating a variety of collaborative tools and evaluating them with users. It looks at their usability, usefulness, the participants' level of social awareness, and their satisfaction with the process and end product of the task.

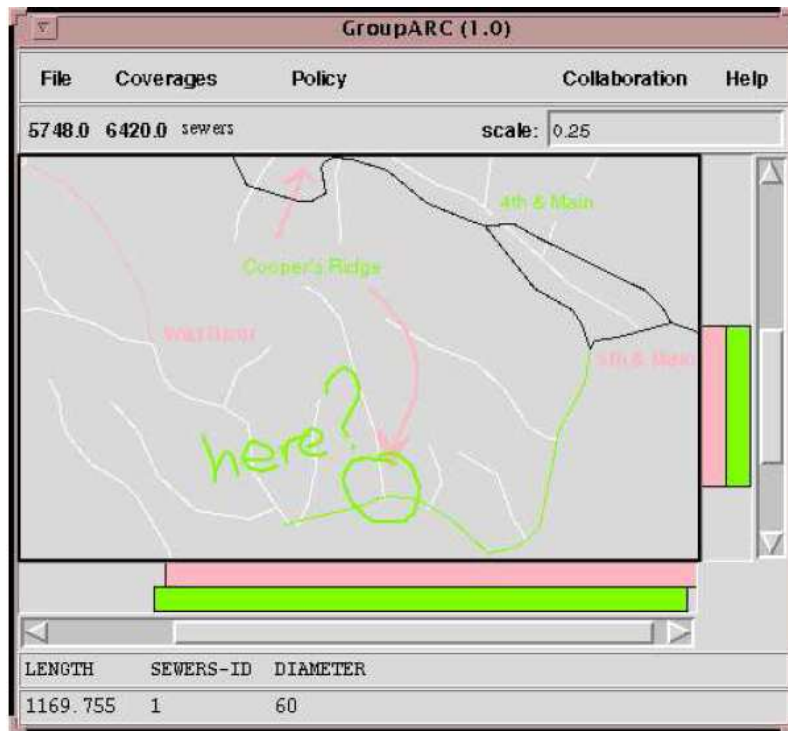


Figure 2.5 GroupARC

Churcher and Churcher explore combining a groupware toolkit with a geographic information system.

2.2 Two-Dimensional Collaborative Solutions


In addition to the GIS literature, there are also prior examples of using spatial representations for collaboration, although not necessarily for spatial collaboration. These examples include 2D, as well as 3D representations. With respect to 2D representations, the GroupLab research group at the University of Calgary has examined using spatial representations to support workspace awareness. Through examining face-to-face collaboration around a shared space, such as a table or whiteboard, they provide a framework for awareness information that is exchanged in a same-place setting (Gutwin & Greenberg, 2002). This framework provides an understanding of how distributed settings differ and encourages the design of collaborative solutions with appropriate awareness features.

2.2.1 Radar Views

The GroupLab research also extensively examines the use of the radar view technique. Initially described as part of the SharedARK system, radar views are one way to provide workspace awareness information (Smith, O'Shea, O'Malley, Scanlon, & Taylor, 1989). They display a miniature of the entire workspace and indicate each collaborator's current viewport into the shared representation (Figure 2.6). This allows each user to visually understand his/her view with respect to the space and with respect to the other users' views. Many different implementations of radar views have been developed and evaluated. They have been used with text documents and simple spaces that are navigated through panning interactions, such as laying out newspaper articles and creating concepts maps (e.g. Begole, 1998; Gutwin & Greenberg, 1998; Gutwin, Greenberg, & Roseman, 1996). They have proved to be a simple awareness feature

corresponding to increased group usability, reduced task completion times, more efficient communication, and increased user satisfaction (Gutwin & Greenberg, 1999)

Figure 2.6 Radar view



The radar view displays a miniature of the shared workspace and indicates where each collaborator is working with a fictional welding task.

However, applying radar views to a spatial collaboration solution is not straightforward. Spatial representations represent a complex space with many intricacies. As a result, the representation can involve numerous levels of detail that are navigated with varying scale factors. For example, zooming in with a two-dimensional map typically displays a smaller area and provides more information. The work reported in this thesis investigates how to portray this style of navigation in the miniature in a useful and usable manner. Radar views have proved useful in the past, and this work examines their utility for spatial collaboration.

2.2.2 Other Two-Dimensional Collaborative, Spatial Representations

Another example of using a spatial representation for collaboration is to integrate it with workplace communication. Ramonamap (Bartlett, 1994) is an interactive map tool that organizes workplace messages by their associated physical locations. For instance, a "printer down" message appears at the printer's location on a shared map, as opposed to being sent in an email notice. In fact, much of the information that flows in a workplace can be related to a physical place, and Ramonamap takes advantage of this. The tool allows users to leave notes for one another, find out conference room schedules, and learn information typically found in a personnel directory.

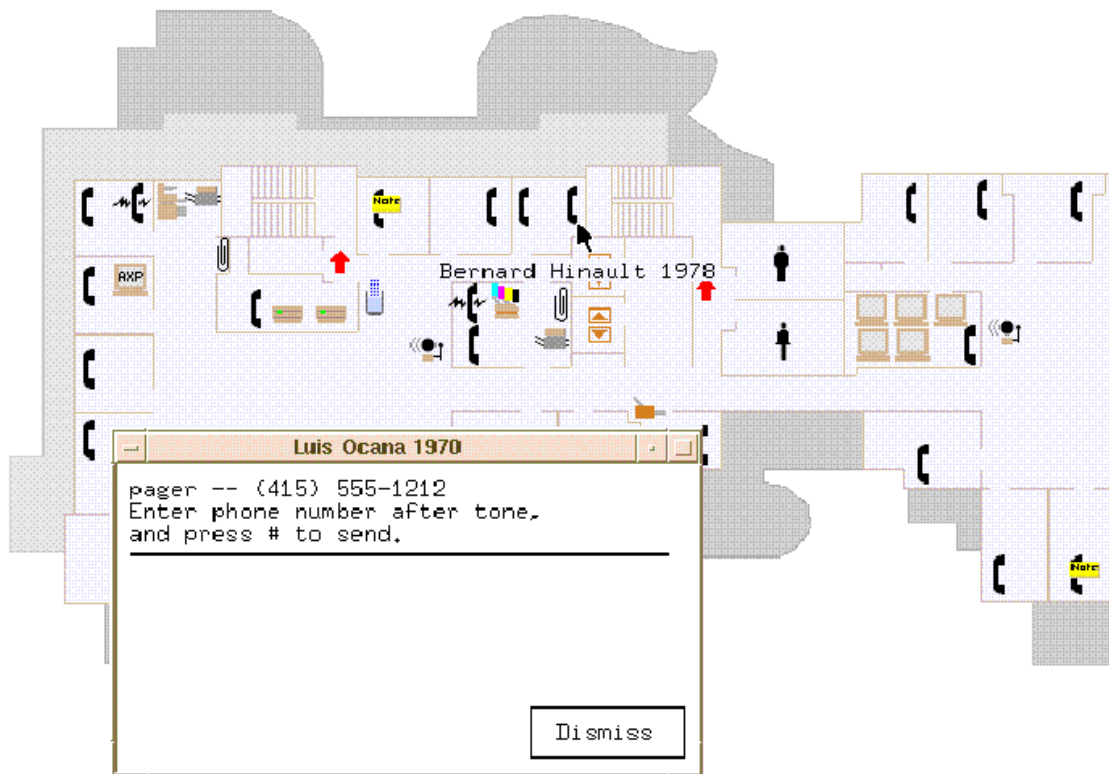


Figure 2.7 Ramonamap

This collaborative software organizes workplace information spatially, providing personnel contact information, printer status information, and conference room schedules.

Similarly, ActiveMap is a visualization tool that provides awareness information about the location of workers (McCarthy & Meidel, 1999). Designed to encourage casual conversations beyond neighboring workspaces, the tool uses a map background and a collection of faces to indicate the most recent location information of individuals. Location data is collected through a system of infrared sensors and badges worn by the workers. Initial user experiences with ActiveMap revealed positive results, but the tool is still under investigation.

The work reported in this thesis extends these projects as it examines other spatial collaboration tasks. It is interested in people using a spatial representation for more than organizing messages and providing awareness information. The idea is that a collaborative, spatial representation can support synchronous, spatial collaboration tasks such as giving directions, brainstorming ideas, and holding a spatial discussion. This work evaluates tasks that require users to explore layout options and make decisions about a space.

It is also important to note that many of the Web-based map programs could be considered as collaborative applications (see Chapter 3). These programs often include a feature that sends an email message announcing a Web page with a particular map display. For example, multiple people might use this if they were planning a trip together. This feature is very limited, however, as many messages must be sent when the task is more involved than simply passing along a map

or a set of directions. Again, the work of this thesis differs in its focus on synchronous communication.




Figure 2.8 ActiveMap

This collaborative software indicates the location of coworkers spatially, using photographs to reflect their current whereabouts.

2.3 Three-Dimensional Collaborative Solutions

In addition to 2D spatial representations, previous work has also explored using 3D representations for collaboration. Collaborative Virtual Environments (CVEs) is a field that explores how virtual reality can be used for collaboration. CVEs provide a way for multiple, distributed users to synchronously share a three-dimensional representation. They are a special case of virtual environments where more emphasis is typically placed on the interpersonal communication. Applications range from teaching online university courses (Lau, Curson, Drew, Drew, & Leigh, 1999) to team-based military training (Mastaglio & Williamson, 1995) to poetry performances (Benford, Reynard, Greenhalgh, Snowdon, & Bullock, 2000). They have even been used in primary education projects (Roussos, Johnson, Moher, Leigh, Vasilakis, & Barnes, 1999).

The work of this thesis differs from most CVE applications. It focuses on CVE applications in which the structure of the environment is meaningful. Instead of simply supporting a group's learning or artistic expression, the environment represents the spatial problem that the collaborators are working to resolve. The CVE represents the physical, spatial entity and the users navigate the three-dimensional space to investigate its structure and details.

One extension of CVE research relates to the field of Augmented Reality (AR). AR technology overlays virtual images onto the real world through specialized head pieces. With multiple

collaborators, AR can be used to augment traditional, face-to-face design meetings. Billinghurst, Kato, Kiyokawa, Belcher, and Poupyrev discuss using tangible augmented reality to support spatial discussions. The use of augmented reality is combined with tangible objects to allow users to explore spatial layouts (2002). Similarly, the Augmented Round Table for Architecture and Urban Planning project uses AR to enhance the architectural design process (Granum, Moeslund, Stoerring, Broll, & Wittkaemper, 2003). It provides placeholder objects in the physical world that correspond to virtual objects in the 3D view. Likewise, the Tiles project allows collaborators to interact with both virtual and physical objects to author interfaces, such as the design of an aircraft instrument panel (Poupyrev, Tan, Billinghurst, Kato, Regenbrecht, Tetsutani, 2001). The work presented in this thesis differs in its focus on typical desktop technology used in distributed settings.



Figure 2.9 Augmented Reality examples

AR examines face-to-face spatial collaboration with physical objects and three-dimensional visualizations.

One example of a CVE used for a distributed, spatial collaboration is a prototype system called CALVIN, or Collaborative Architectural Layout Via Immersive Navigation (Leigh, Johnson, & DeFanti, 1996). The Electronic Visualization Laboratory at the University of Illinois at Chicago recognized that much of the architectural design process involves experimenting with paper models, however virtual environments were only being used for 3D walkthroughs of more complete designs (Leigh, Johnson, Vasilakis, & DeFanti, 1996). To address this, they implemented a multiple perspective approach to virtual reality. The idea is that multiple viewpoints, different information filters, and the ability to experiment with multiple designs will support the more creative aspects of architectural design. The initial CALVIN prototype supported two frames of reference for the same space. The egocentric frame of reference places the user inside the architectural model so that he/she experiences the structure as if it was actually built, while the exocentric frame of reference is one from above looking down at the entire structure. Different participants had different frames of reference, but both could interact with the environment and move objects, such as furniture.

CALVIN has been used in three different architectural design studies, and in each situation the person with the exocentric frame of reference took charge of the large movements, while the user with the egocentric frame of reference performed fine adjustments (Leigh & Johnson, 1996). This corresponds to other researchers' observations that a first-hand view affords fine grain manipulations, while a global view is preferable for gross manipulations (Stoakley, Conway, &

Pausch, 1995). The research of this thesis explores this concept of viewpoints further. It expands upon the CALVIN project, revisiting their exocentric-egocentric combination and investigating additional frame of reference combinations. Using empirical methods, this work analyzes the differences and effects of these combinations. It also differs in that the collaborators interact with the CVE through a desktop display, as opposed to the more immersive technology used in the CALVIN prototypes.



Figure 2.10 Collaborative Architectural Layout Via Immersive Navigation (CALVIN)

Designed for immersive virtual environment technology, CALVIN uses different frames of reference to support interactive architectural design.

Research conducted by Yang and Olson also has explored frame of reference effects within a CVE. Using a fabricated virtual environment created from a scientific data set, they investigated various egocentric-exocentric implementations (2002). The goal was to enable a distributed group to examine trends in a space of three-dimensional data points using a desktop display (Yang, 2002). Conducting an experiment based on their model of collaborative navigation, they found that with different perspectives, groups struggle to perform the spatial, mental transformations required to see the same data objects. However, they also observed that with a pair of participants, an exocentric user could quickly provide directions so an egocentric partner navigated his/her view to include the target object. The work of this thesis is similar in its comparison of multiple frame of reference combinations. It is interested in the effects of an exocentric-egocentric combination as well as other frames of reference combinations. This thesis differs from that of Yang and Olson in that we are interested in environments that represent realistic locations. In the spatial collaboration tasks, egocentric navigation is constrained to a floor plane, as opposed to movement in three dimensions throughout the space, and the task involves object manipulation rather than data exploration.

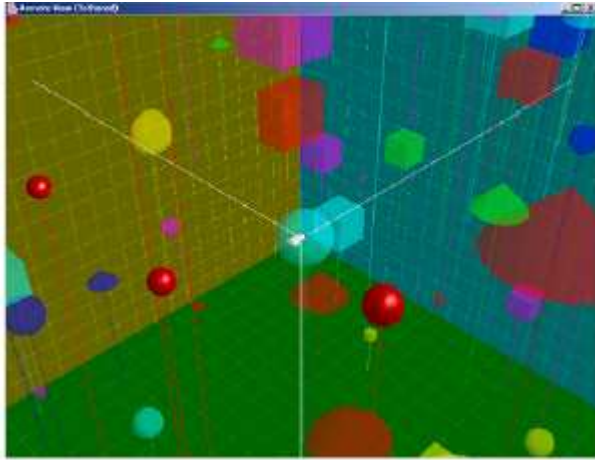


Figure 2.11 Collaborative navigation in a CVE

An exocentric user navigates an egocentric partner (white object at center) through a scientific data set to investigate the effects of the frame of reference combination.

Lastly, there has been some prior work with understanding and supporting social awareness within CVEs. The University of Nottingham and King's College London has reported on issues with collaborators understanding one another's locations and viewports into the environment. Their interest was in object-focused interactions in desktop virtual environments. Using the MASSIVE-2 CVE, they observed that two users with egocentric viewpoints experience awareness problems when trying to arrange furniture in a single room (Hindmarsh, Fraser, Heath, Benford, & Greenhalgh, 2000). The participants could not see their partner's avatar and the objects of discussion with the same viewpoint. They also had difficulty seeing what their partner was doing and understanding what their partner could see. To address these issues, peripheral lenses and explicit representations of actions and viewing fields were implemented in a second version (Fraser, Benford, Hindmarsh, & Heath, 1999). This provided greater awareness to the participants, but also introduced new problems because the design used distorted views.

Dyck and Gutwin have also recognized the need to support awareness in three-dimensional environments (Dyck & Gutwin, 2001; Dyck & Gutwin, 2002). They propose a set of techniques to provide awareness information about each collaborators location. The techniques address four aspects of location awareness including which direction a collaborator is facing, what objects are in a collaborators viewport, where each collaborator is positioned with the space, and how far away a collaborator is. The collection of techniques provided improved awareness in comparison to simply using avatars, but often added clutter to the environment in their use of colorful visual cues. Similar to the work of Yang and Olson, the techniques are demonstrated within a three-dimensional dataset that does not represent a physical place.

The work reported in this thesis continues to investigate this issue of awareness in CVEs. It observes collaborators interact with one another while performing object manipulation tasks. It also explore new approaches to supporting social awareness by modifying radar views, tailoring

some of Dyck and Gutwin's techniques for realistic places, and observing awareness effects with frame of reference combinations.

2.4 Summary

In summary, this literature review illustrates that spatial collaboration is multifaceted with issues investigated by geographic information systems, groupware systems, augmented reality, and collaborative virtual environments. This demonstrates the range of possible spatial collaboration activities and provides a context for the limited focus of this work. The review also highlights prior work that relates the specific problem statement of this research. Prior geography work encourages the use of shared, interactive representations and prompts an investigation of interaction techniques. The GIS solutions presented focused on representing the space, but discuss few features for collaborative interactions. Also, both previous groupware work and collaborative virtual environment efforts emphasize the issue of social awareness. These projects discuss the importance of knowing where other collaborators are working. They offer some initial approaches to providing this awareness information, but further work is required to ensure the techniques are both useful and usable. Lastly, the collaborative virtual environment research suggests that there is advantages to having collaborators use different frames of reference. This claim is not validated with a realistic spatial collaboration task and desktop technology, though. In the examples, it applies to more immersive virtual environment technology and scientific data sets.

Chapter 3

3 Preliminary Work

We pursued a number of projects prior to the later planned studies of this work. These projects relate to the topic of spatial collaboration and provide an understanding of its challenges. They investigate ways to enhance distributed, spatial collaboration activities and motivate the later phases. One project reviews a design space for spatial representations, another applies a map interface to a collaborative environment, and the last offers an initial prototype to support distributed, spatial discussions.

3.1 Classification of Interactive, Map Software

One such project was a survey of the state-of-the-art of interactive map software (Schafer, 2001). Map software takes advantage of people's familiarity with maps and offers potential solutions for spatial collaboration. To gain a better understanding of the differences among map software applications, a classification was developed. This classification enables a comparison of map applications, revealing differences between software that seems very similar, and highlighting similarities with software that appears very different. The classification focuses on the interactions supported by the map software and the features available to a user.

Forty different map applications were surveyed during the creation of the classification of interactivity (see Appendix A). Web-based map programs, map software developed for various research projects, and off-the-shelf software products were examined. This sample provided a good cross-section of interactive map software, as the user tasks range from sharing travel information, to viewing hydrographic data, and creating a military battle plan. For each application, a description of the possible user interactions was recorded, which was then generalized into a classification of interactivity.

The classification consists of three dimensions: navigation techniques, support for collaboration, and data sources available (Figure 3.1). Navigation techniques refer to the way a user can change the viewpoint of the map, such as through zooming and panning. Support for collaboration refers to an application's ability to allow multiple people to communicate ideas. For example, the map software may allow users to chat in real-time, or write comments on the map for later viewers. Data source availability refers to the options available for displaying different data on the map. It investigates whether or not the user's actions can have any effect on the data displayed and if so, how much control is given to the user.

In terms of navigation techniques, map software can simply provide multiple map images, or it can use some form of panning and zooming. With multiple images, hotspots, such as Web-based imagemap configurations, link the multiple maps. Using panning and zooming, navigation occurs through either discrete interactions or continuous interactions. With discrete interactions, the map display refreshes after the interaction is completed. For example, a user could click on the map and drag it to pan across the map area, but the newly exposed area is only rendered after the mouse is released. With continuous navigation, on the other hand, the display provides continuous feedback to the user. The classification distinguishes between panning and zooming

support, considering four possibilities for each: no support, discrete interactions, continuous interactions, or both discrete and continuous interactions.

Support for collaboration is decomposed into asynchronous and synchronous support, depending on whether collaborators can exchange ideas by using the software at different times, and by using it at the same time (Ellis, Gibbs, & Rein, 1991). Within each type of collaboration, there are four levels of interactivity. At the most basic level, there is no support for collaboration, meaning that there are no features that assist multiple users working together. The other levels incorporate the different means for supporting collaboration: collaboration using an external resource, such as chat tools or email, collaboration via the map interface, such that multiple users can share ideas directly with the interface, and collaboration using both an external resource and the map interface.

The third dimension of the classification, data source availability, consists of four possibilities. The first level is the most basic and does not give the user any data options. The map uses only one data source and always displays that same data source. A map image on a Web page is a good example of this data source level. At the next level, a map application uses multiple data sources, but can only display one source at a time. For example, a display with a New York state map could allow users to see different county statistics such as population and density of people, but it only displays one statistic at a time. The third level is termed non-interactive data sources. In this case, the user does not have any data options, but the software incorporates different data sets and has an underlying algorithm for their use. This feature is commonly activated when a user zooms in on a specific area of the map. More data and details are often displayed on the map when the magnification increases. At the last level, the user has total control over the data sources used as he/she chooses which data sets to display. This is commonly referred to as layers, and is widely used in Geographic Information Systems (GIS).

Given this classification, it is interesting to consider the advantages and disadvantages of the different techniques. For instance, discrete navigation interactions can be frustrating to a user, but they also enable precision. Discrete panning allows the viewpoint to be changed with distance measurements, and discrete zooming allows the viewpoint to be scaled by specific values, such as 2x. When considering continuous interactions, their use of free-form movement allows a user to obtain any number of viewpoints, as opposed to the fixed number available with discrete interactions. On the other hand, continuous interactions may be difficult to control. If the area of the map is large and the scale range is extreme, interacting with the map could lead to undesirable viewpoint adjustments.

Upon analyzing the different types of collaboration support, we discovered that a map application that only supports asynchronous or synchronous collaboration can be restrictive. With just asynchronous collaboration, the map software does not allow for chance encounters. On the other hand, with simply synchronous collaboration, more than one user has to be using the application at the exact same time. It is possible to explore the map individually and then use another form of communication, but it can be difficult to specify a spatial concept without using the map application that generated the idea. Similarly, when collaboration is only possible through external resources, this can limit the communication. When activities are not map intensive, such as discussing general trends in the display, the external resources may be

sufficient. At other times, a user may want to illustrate his/her idea by physically pointing out an area on the map, or annotating on the map for others to see. Collaboration using just an external resource does not allow these interactions to occur.

The forty map applications in the original survey were dispersed across the classification. Many of the map applications surveyed did not support collaboration, however. This was due to the fact that many applications were located on the Web, which typically requires additional software in order to support collaboration. It is interesting to explore the classification for novel combinations that include this collaboration support (Card, Mackinlay, & Robertson, 1990). One of the novel combinations we identified supports asynchronous and synchronous collaboration using an external resource and a layers data source. This is unique because most map applications that feature layers offer no support for collaboration, and those that do often focus on asynchronous collaboration using the map interface. This particular combination allows multiple users to share ideas and explore different combinations of data sources. It enables users to work independently with the map and then contribute their ideas either in real time or over time.

By modifying this last combination, other unique combinations exist. Adding support for synchronous and/or asynchronous collaboration using the map interface allows users not only to discuss the different views of the map, but also to collaborate on the map display itself. For example, the software might allow users to share annotations on the map, or enable users in different locations to view the same layers, but navigate the map independently. Adding these extra components opens up a range of possible tasks.

This map survey project helped formulate the research questions of this work. It revealed common navigation interaction techniques, as well as less frequently implemented collaboration features. Exploring the classification of interactivity also encouraged an analysis of the different interactions and the unique ways they could be combined. This prompted further questioning about their advantages and disadvantages in collaboration activities. As a result, some of the techniques observed in this study are compared and evaluated in the other phases of this work.

<p>Navigation Table Techniques</p>	<p>Multiple Maps OR Zooming Support</p> <table border="1" data-bbox="810 394 1449 719"> <tr> <td>Discrete and continuous interactions</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>Continuous interactions</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>Discrete interactions</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>No Support</td> <td></td> <td></td> <td></td> <td></td> </tr> </table> <p>No support Discrete Interactions Continuous Interactions Discrete and continuous interactions</p> <p style="text-align: center;">Panning Support</p>	Discrete and continuous interactions					Continuous interactions					Discrete interactions					No Support				
Discrete and continuous interactions																					
Continuous interactions																					
Discrete interactions																					
No Support																					
<p>Support for Collaboration</p>	<p>Asynchronous Support</p> <table border="1" data-bbox="676 909 1401 1323"> <tr> <td>External resource and map interface support</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>Map interface support</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>External resource support</td> <td></td> <td></td> <td></td> <td></td> </tr> <tr> <td>No Support</td> <td></td> <td></td> <td></td> <td></td> </tr> </table> <p>No support External resource support Map interface support External resource and map interface support</p> <p style="text-align: center;">Synchronous Support</p>	External resource and map interface support					Map interface support					External resource support					No Support				
External resource and map interface support																					
Map interface support																					
External resource support																					
No Support																					
<p>Data Source Availability</p>	<table border="1" data-bbox="568 1547 1267 1585"> <tr> <td>One data source</td> <td>Multiple data sources viewed separately</td> <td>Non-interactive</td> <td>Layers</td> </tr> </table>	One data source	Multiple data sources viewed separately	Non-interactive	Layers																
One data source	Multiple data sources viewed separately	Non-interactive	Layers																		

Figure 3.1 Interactive, map software classification.

The diagram indicates the three dimensions of the classification and the numerous possibilities for map software.

3.2 Map-Based Navigation in MOOsburg

Another project we conducted explored the use of a map-based navigation tool in a collaborative, community environment (Schafer, Bowman & Carroll, 2002). MOOsburg is a collaborative environment based on the town of Blacksburg, Virginia whose goal is to support community development within the town (Carroll, Rosson, Isenhour, Van Metre, Schafer, & Ganoë, 2001). Ideally, users will build distinct locations within the environment, similar to rooms, and establish a community network through synchronous and asynchronous communication. The project reported here developed a navigation tool for MOOsburg that incorporates a map interface.

In traditional MUDs and MOOs, users interact with one another and with the objects in the environment through simple textual commands. The interface is a scrolling sequence of commands and descriptions that pertain to the activities occurring in a particular room of the environment. Movement is accomplished with textual commands such as "go north." Users navigate from room to adjacent room using their knowledge about the layout of the environment. To move large distances across the environment, multiple commands are repeatedly used.

Map-based navigation provides one alternative to these traditional interfaces. It uses a graphical map of the environment to support navigation. Users do not have to develop a mental model of the spatial layout, and the map provides direct access to rooms. Such a graphical overview of the environment can also provide visual cues to guide users in finding interesting activities. It can indicate rooms with many objects, rooms with many users, or rooms that are visited most often, etc. In MOOsburg, map-based navigation also establishes a connection between the real place and the virtual environment. The map provides a familiar view for the people local to the area, possibly encouraging a sense of community.

The usage scenarios for MOOsburg are fairly different from those of traditional MOOs. In MOOsburg, users will want to find information and collaborate with other users. The tool developed supports these tasks. The map interface uses different visualization techniques to provide information about the landmarks. A red spot indicates the user's current location on the map. Spots colored light blue indicate the presence of other users, while dark blue corresponds to places currently without visitors. We also represent the number of users at a location through the size of a spot, where larger circles correspond to more visitors.

The map interface for MOOsburg also needs to address the transformation from cognitive maps of Blacksburg to a software-based, graphical representation. People need to recognize the map as their town and be able to relate to it. Including landmarks on the map interface is one way to make this connection. The design can also be tailored to the Blacksburg community and make use of community knowledge such as prominent places in the town. Thorndyke and Hayes-Roth explain that people understand spaces using different types of knowledge (1982). Landmark knowledge is a familiarity with the visual attributes of an environment. Route knowledge is a familiarity with paths through an environment; and survey knowledge is a familiarity with an overview layout of an environment, similar to a map. The community members of Blacksburg, who have a cognitive map through first-hand active navigation, are likely to have developed landmark knowledge and route knowledge. Similarly, Lynch proposed that people use paths (e.g. walkways and passages), edges (e.g. walls and fences), landmarks (e.g. church spires), nodes (objects with similar characteristics), and districts (distinct sections) to build their

cognitive maps (1960). Placing these objects on the map will also help people to recognize locations and understand the representation.

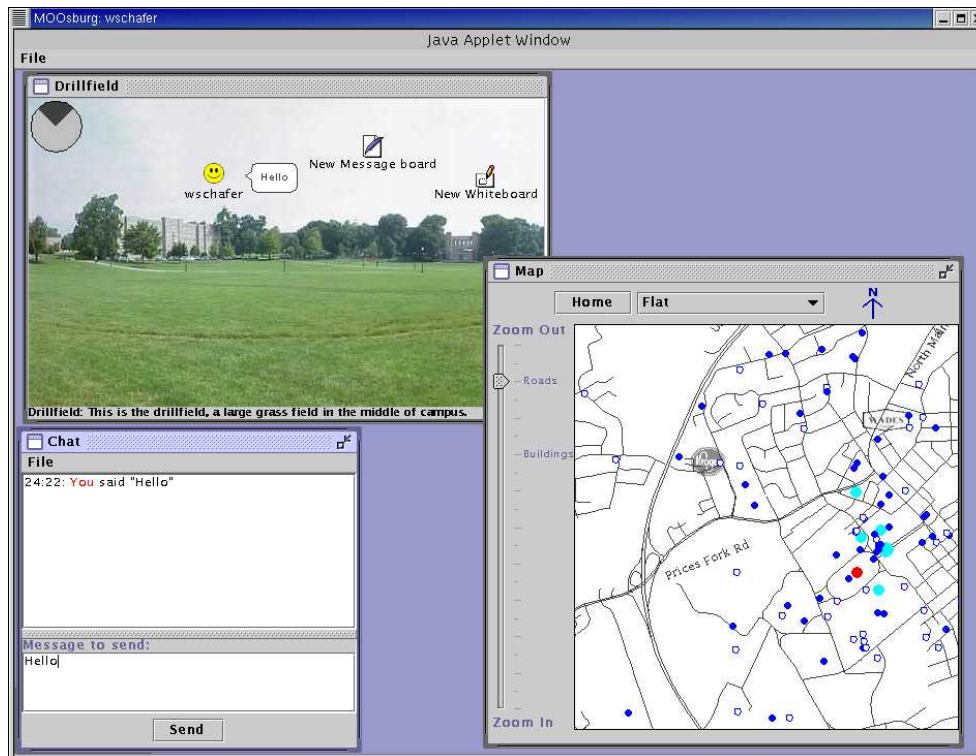


Figure 3.2 MOOsburg

The map-based navigation tool (lower right) allows users to visit different virtual spaces. User 'wschafer' is currently visiting the Drillfield.

The map supports four main types of interaction—applying projections, using a single click, dragging the mouse, and zooming. The data displayed consists of a set of vector coordinates, allowing transformations with different mathematical functions (Arc Tan, Hyperbolic, Parabolic x^2 , and Parabolic $x^{1.5}$). These functions produce different projections that create a fisheye view (Figure 3.4). Fisheye views use focus+context techniques to present details as well as a compressed view of the overall structure in one image (Furnas 1982; Greenberg, Gutwin, & Cockburn, 1998; Lamping, Rao, & Pirolli 1995). When users click on the interface, the map is redrawn with the location of the click placed in the center of the map. Dragging the map produces a panning motion, revealing parts of the map not previously visible. Lastly, users can change the zoom factor of the map continuously through a slider widget. Figure 3.3 and Figure 3.4 illustrate these map interactions.

Using a prototype of the map-based navigation tool, we conducted a formative evaluation with representatives from three different user groups in Blacksburg: middle school students, college students, and seniors (see Appendix B). This allowed us to explore how to tailor the design of the map for the community members, as well as investigate the usability of our interaction techniques. The participants were presented with paper maps and asked to point out locations in order to learn about their perceptions of the town. They also interacted with the map prototype

to test out the clicking, dragging, and zooming interactions. Similarly, participants pointed out locations with paper maps with fisheye projections, and interacted with the fisheye projections in the map prototype.

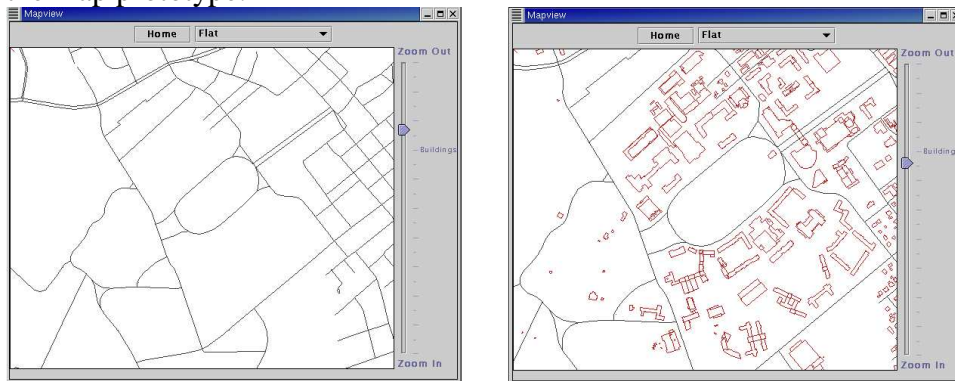


Figure 3.3 Zooming with the map-based navigation tool

With the flat view, the map does not use any projections. Dragging the slider on the right changes the scale and adds additional details.

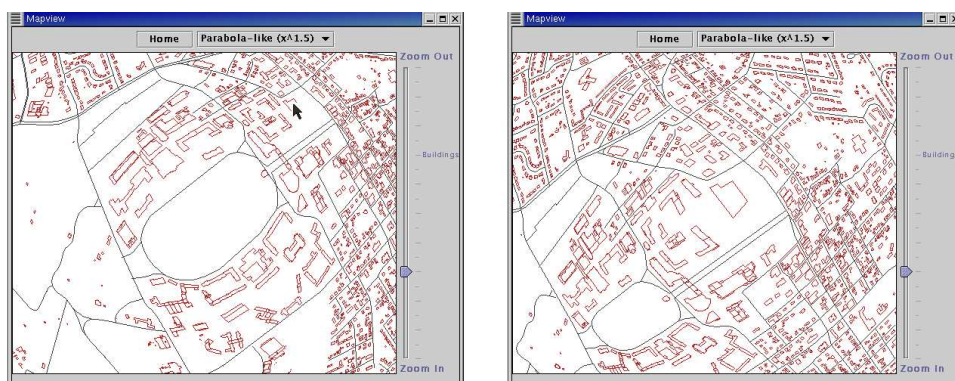


Figure 3.4 Clicking with the initial map-based navigation prototype

Using the Parabolic fisheye projection and clicking on the map, the display animates so that the clicked location is positioned at the center of the map.

This evaluation revealed simple modifications that we made to increase the prototype's usability. These included placing the zoom slider widget to the left of the map so that it would not be confused with a scrollbar, adding road labels, and using levels of detail to support the navigation process of narrowing in on a location. The clicking interaction was also changed from a single left-click to a right-click with a pop-up menu to reduce confusion about the re-centering.

This evaluation also provided a better understanding of how people use interactive map software that represents a familiar place. These results are characterized as a set of eight lessons learned:

1. Individual users have different perceptions of the same place. Most of the eighth graders did not recognize a major intersection in town when we gave them the road

- names. They may have been familiar with this intersection, but known it by some other way, such as the intersection with the video store.
2. Individual users are familiar with different areas of the same place. Some of the users lived in the southeast section of town and were familiar with various places and roads in that area, while other users had little knowledge of that section of town.
 3. Different user groups have knowledge of different landmarks in town. Most of the college students did not recognize the places thought to be major landmarks in town, while none of the eighth graders or the seniors had trouble with these landmarks. The grocery stores were an exception as they were recognized as by all of the participants.
 4. User groups differ in navigation techniques. All of the participants with drivers' licenses referred to road names throughout the tasks, while all of the eighth graders relied on landmark knowledge.
 5. Recognizable landmarks are key for navigation tasks. While the users navigated with the map prototype, they were observed to be looking for landmarks that they recognized in order to orient themselves with the map.
 6. Users are drawn to large buildings. These buildings are prominent on maps, and users will try to use them during navigation.
 7. Users learn about other landmarks while working with the map. Uniquely shaped buildings identified in earlier tasks were referenced by the same users in later tasks.
 8. Users need visual reminders of how to navigate the map. None of the users used all of the interaction techniques to complete the tasks. Many had to relearn how to interact with the map after looking at paper mockups.

This work with this map-based navigation tool has touched on a number of concepts related to spatial collaboration, including awareness, interaction techniques, and spatial knowledge. The tool aims to support social awareness among collaborators interacting from different locations. It uses a spatial approach to help the collaborators discover where other collaborators are working. Knowing where others are working is also an issue when the map interface directly supports collaboration. The map-based navigation tool also examines unique interaction techniques. Many of these techniques are compared and evaluated in the other phases of this work. Lastly, the experiences with everyday people working with an interactive spatial representation of a familiar place provide valuable insight into the field of spatial collaboration. The evaluation observed people searching for recognizable landmarks, focusing on large buildings, and using different approaches to navigate the map.

3.3 Collaborative, Spatial Discussion Prototype

Another project we conducted explored the use of a collaborative, spatial discussion prototype for a local environmental group. The Stroubles Creek Focus Group is interested in the protection and restoration of Stroubles Creek in the town of Blacksburg, Virginia. The organizations involved in this group collaborate through monthly meetings where they discuss the latest issues and efforts surrounding the creek. These meetings are beneficial, but scheduling conflicts occur. Members of the group also talk with one another outside of these meetings through email and phone use. However, this communication typically occurs between two people, and the results are not usually transmitted to the other members. Furthermore, each organization collects different types of data on the creek, but they do not have a process to share this information. The

meetings coordinate activities among the organizations, while the data collection results remain local to the organizations gathering the information.

This focus group's collaboration needs are unique in that they have a common interest in a physical, spatial entity. One way to support their stream-based collaboration is through an interactive map application. A map interface provides a familiar representation of the creek and allows for an exchange of spatial ideas. Using map software to support collaboration appears to be a good idea, but there are many questions to ask about the collaborative software. For starters, what kinds of collaboration tasks can be supported with map software, and should the communication be synchronous or asynchronous? More specifically, which interaction techniques are appropriate for collaborative map tasks, and can a specific design be easy to use for both experts and non-experts of stream ecology?

After meeting with each of the organizations of the focus group, we devised a number of scenarios to highlight the group's interactions with each other and their possible uses of a collaborative map tool. We discussed these scenarios with the group to obtain a more accurate description of their activities. For example, one scenario involved a discussion of the use of a common set of terms. This scenario was based on a previous conversation of the group members and reflected how a similar conversation could occur using a collaborative map tool:

Ray and Lyn both have an interest in Stroubles Creek. Ray has been working with the Water Center to conduct an assessment of the creek and Lyn is the coordinator for Save Our Streams (SOS) efforts in Blacksburg. Both are accustomed to looking at maps of the watershed such as topographical maps of the area, GIS stream layers, or other types of cartographic representations. Yet, Ray and Lyn have different terms for the various sections of the creek. Using a map object in MOOsburg, the two agree on a common set of identifiers and begin to communicate more about the different sections using the map as a reference point.

Following this discussion of scenarios and realistic group activities, we generated a list of ideal requirements for the collaborative map tool. This list acted as a guide in the design process and included statements like:

- The map object needs to have features that are easy to use and can be learned quickly.
- The map needs to provide a way for the group to point out locations, such as the straight pipes and a new creek threat.
- The map must provide a way for the groups to select and discuss areas or sections of the stream, such as the areas assessed by the SOS project and the area referred to as the West Branch.
- The map must provide a way to access historical information relating to the creek.
- The map needs to distinguish between different data collection sites in terms of the type of data collected, the organization responsible, the dates of collection, etc.

This requirements list was lengthy, and a number of paper prototype designs were attempted before narrowing the scope of the project. The decision was to focus on a tool that would enhance the group's distributed discussions. This goal was provide a spatial context for the discussions that typically occurred through email and phone. The tool allowed both

asynchronous and synchronous communication to occur, and served as a central location for the group to discuss stream issues outside of their face-to-face monthly meetings.

The prototype design supports navigating the map interface through panning, zooming, and changing the data layers displayed. It also includes a text-based chat feature and shared drawing tools. The combined use of these features allows the collaborators to carry on a discussion using both text and spatial references. The shared annotation tools allow the collaborators to exchange sketches and scribbles, while the chat feature allows them to contribute textual information and maintain a history of their discussion. This design combines the two features through a technique that enables text messages to be linked with annotations in a collaborative tool (Figure 3.5).

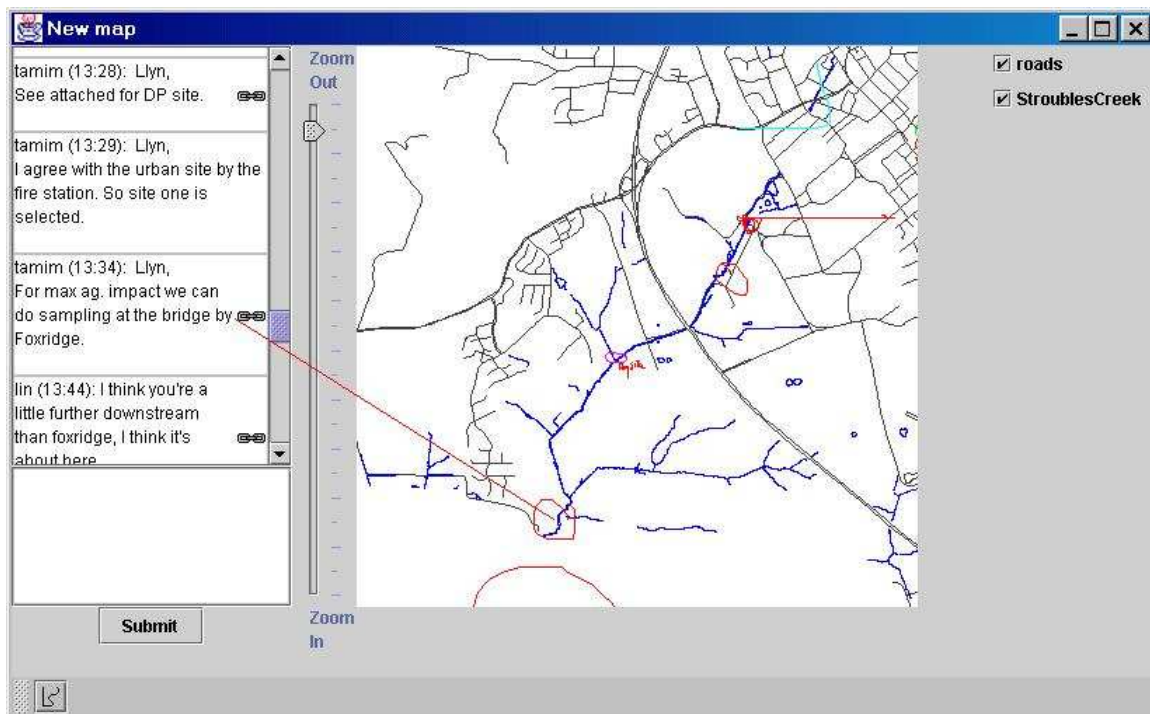


Figure 3.5 Spatial discussion prototype

Collaborators "tamim" and "lin" are carrying on a spatial discussion about locations along Stroubles Creek. The comments with icons have associated drawings. A mouse positioned over an icon draws a line to the corresponding drawings.

This prototype is unique in the range of spatial collaboration activities it can support. Collaborators can contribute synchronously, similar to a chat tool, or they can carry on discussions over time asynchronously. The use of a spatial discussion also encourages a variety of activities. Collaborators might use this tool to discuss different naming conventions for areas of the stream and determine a common set of terms. They also might exchange information about data collection sites or future collection plans. The prototype also lends itself to discussions about new stream observations and possible causes due to upstream activities.

An initial evaluation of the prototype was conducted with two members of the Stroubles Creek Focus Group to learn about the usability and usefulness of the tool (see Appendix C). The two members were asked to work with the prototype to identify four locations along the stream. The locations would be part of a workshop that would provide an overview about the creek's current condition. The group members sat facing one another in such a way that their computer monitors provided a visual barrier. They were asked not to talk, but rather use the tool for all their communication. At the end of the session, each member completed a short list of follow-up questions before a group discussion occurred.

This evaluation revealed some interesting observations. Overall, the tool was useful since the text augmented the drawing, as expected. The text allowed the users to discuss the drawing locations and ask questions about the drawings. In particular, the text-based chat feature was easy to use, while the map created problems. Both participants understood the text tool quickly, and did not resort to talking out loud during the experiment. The map, on the other hand, was at times more of a hindrance than an aid. The participants struggled with where to make their drawings because they lacked an understanding of the map. The map lacked features that the participants recognized. One user even pulled out her familiar topographical map for comparison. She explained that the topographical map provided contours that are helpful in understanding the creek's path, unlike the interactive map. The interactive map also lacked imported landmarks like bridges, government buildings such as the fire station and post office, and the sewer treatment plant. This made it difficult for the collaborators to identify locations and landmarks, thereby denying them a common point of reference.

The communication was also slow paced, as one user would often have to wait as the other typed out his/her response and possibly added a drawing. Time was required for both steps. Neither group member was efficient at typing and each struggled in deciding where to mark the map.

The evaluation also presented a number of additional features that could be added to the tool as well as some usability problems with the current features. The participants wanted a way to delete annotations, as they would often draw something and then change their mind about where it was located, or the markings they used. They also suggested a text tool would allow them to label places on the map as opposed to writing with a freehand drawing tool. At the end of the session, the participants also wanted to save and print their conversation and map so they had record of it. There was also some confusion in knowing when someone was typing and drawing. The typing and drawing appeared only after the statement or drawing was completed. The drawing feature had another issue in that the participants would draw in different areas and not notice their partner's addition.

This project with the spatial discussion prototype and the Stroubles Creek Focus Group motivates additional research on distributed, spatial collaboration. It provides one example of how a shared, interactive representation can support spatial collaboration activities. It also encourages an investigation of other possibilities in terms of the activities supported. Given the initial list of requirements, what other designs could support this group and their efforts together? Can the current prototype be improved so that it is more useful? What techniques would allow the collaborators to be more aware of what each person is working on, and where in the space they are focused?

Chapter 4

4 Two-Dimensional Map Study

In the following sections we describe a study of navigation and radar view technique combinations and their effects on spatial collaboration (see Appendix D for materials). The empirical study uses a common collaboration task completed by knowledgeable, but not necessarily technical users. In the following sections, we present the navigation and radar view awareness techniques implemented, the task and interface used, as well as the experimental design and procedure. The results from the study are followed by a detailed discussion of the findings and implications for supporting distributed, spatial collaboration.

4.1 Study Details

In thinking about spatial collaboration and the tasks involved, it is natural to consider supporting these activities through a collaborative, two-dimensional application, or a collaborative map. Map software, in general, has become more common, especially with increasing use of the Internet. For example, electronic maps are no longer available just to the institutions and people who purchase the software. Anyone with Internet access can visit a website such as Mapquest (www.mapquest.com) or Maps-On-Us (www.mapsonus.com) and interact with a map application to generate individualized maps, determine routes, and find location information.

It also known that collaboration involves both individual and shared efforts (Dourish & Bellotti, 1992; Gutwin & Greenberg, 1998). For instance, a participant's thought process often involves expressing ideas as well as understanding their partners' ideas. One approach for supporting this mixed-focus collaboration is to allow each user to navigate the representation individually and be able to perform actions not visible to the others, while providing awareness information about the group effort. For example, a group of people designing a town park from distributed locations might want each person to have an individualized map of the park where everyone sees the same set of participant-created annotations. Our representation uses this approach of allowing users to explore independently, which is termed *relaxed what you see is what I see* (Stefik, Bobrow, Foster, Lanning, & Tatar, 1987).

4.1.1 Variables of Interest

Our investigation focuses on supporting navigation and providing awareness information between the distributed participants. These variables can have an effect on spatial collaboration both independently and collectively. This study examines two extremes in navigation techniques and two alternatives to radar view techniques.

4.1.1.1 Navigation Techniques

In considering how the users will view the representation, the simplest approach is to represent the space of interest with a single image that fits on a typical computer screen. This type of display might work well for some tasks, but spatial collaboration involves a large, detailed spatial entity and many tasks will require these details. Navigation techniques provide a way for representations to display varying levels of scale. They present a "window" into the entire representation and implement an interface for controlling the user's view (Tolsby, 1993). This notion of a "window", also referred to as a "viewing window", should not be confused with a

window interface widget. It is a user's view of the representation at any one time or one's viewport into the spatial representation.

Navigation techniques can vary in the number of viewing windows available to the user and, similarly, the amount of control the user has in switching between these windows. Given a single representation, one navigation technique might offer only a few windows into the space. This limits the way a user can navigate. He is restricted to the windows provided by the technique and he has limited control over what can be displayed. Alternatively, a navigation technique could offer numerous windows into the representation. This provides the user with many navigation options and gives him greater flexibility in adjusting the display. A technique that offers these additional windows may enable a user to manipulate their view to better solve a spatial problem. On the other hand, in thinking about distributed collaboration, having more viewing windows creates collaboration issues. With more windows there are more ways for the multiple, individual displays to differ. Consider the case of just two users. There are numerous possibilities for how the displays can differ. One window might be in a completely different area than the other, or it might be somewhat close, or it might be slightly overlapping in one corner, or it might be mostly overlapping along one side, or it might be completely enclosed due to different scales, etc. This range of differences can create problems for the collaborators. It is not obvious what spatial features are visible on a collaborator's screen or how one user's display compares to the others'. This leads to time and effort by the group to coordinate the individual displays and find similar references. Limiting the number of navigation windows available, on the other hand, reduces this coordination. There are fewer ways the displays can differ and the number of navigation steps required to align the windows is reduced. Mathematically, there is also a greater chance that two windows will be the same, eliminating the coordinated window problem altogether.

Our experiment compares two styles of navigation. Each offers a different number of window options. The discrete navigation condition provides only a few windows into the space, while continuous navigation is much less constrained, offering many windows. In order to examine the effect of window differences, the discrete condition also uses non-overlapping windows. Because the windows do not overlap at a particular scale factor, the differences between individual displays are restricted even further. Two windows at the same scale will either be identical or have nothing in common. Placing these restrictions on discrete navigation may not correspond to the most useful technique for solving a spatial problem, but it allows us to explore whether users have greater coordination problems due to differing displays. Continuous navigation does not implement this constraint.

This comparison between discrete and continuous navigation is similar to previous research on text presentation. With a scrolling presentation, the text changes line-by-line on the screen, while a paging presentation changes the text all at once, similar to turning the page in a book. Previous studies have found that paging enables users to build a better mental representation of the text (Piolat, Roussey, & Thunin, 1997; Schwarz, Beldie, Pastoor, 1983). If this finding also applies to spatial navigation, this would imply that a non-overlapping discrete condition would support a more accurate understanding of the space. This is not obvious, however, because text tasks are inherently different than map-based tasks. Our study explores this idea with discrete and continuous navigation techniques.

4.1.1.2 Navigation Implementations

Both the continuous and discrete navigation are implemented using zooming and panning interactions. These interactions are typical of map software and they allow the viewing window to change in scale and shift horizontally and vertically across the representation.

In discrete navigation, the map representation is divided into ninety-one non-overlapping windows. There are three discrete zoom levels with different windows available at each level. At the most zoomed-out level, the entire map is displayed. At the middle zoom level, there are nine possible windows corresponding to a division of the representation into a 3x3 grid. At the most zoomed-in view, each square of the middle zoom level grid is further divided into nine more windows. Figure 4.1 below depicts these ninety-one windows into the representation:

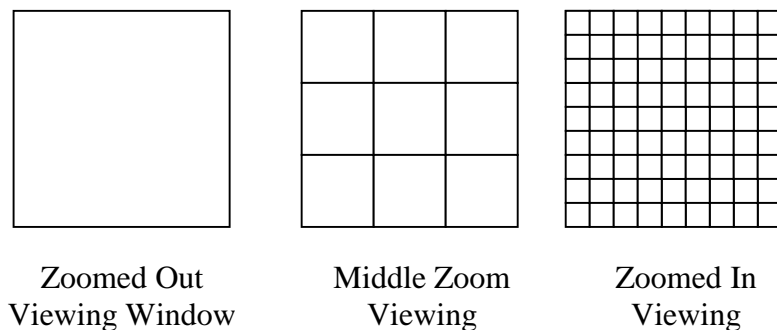


Figure 4.1 Discrete navigation technique

This technique uses three discrete zoom levels with non-overlapping windows into the representation, leading to ninety-one possible viewports.

This navigation technique provides a simple paging approach to viewing a representation. There are relatively few windows into a potentially very detailed and complex space and users are limited in their navigation to jumps between viewpoints. Each location in the representation is only viewable in three viewing windows, corresponding to the three zoom levels.

Because there are a limited number of windows and very few of them overlap, the differences between any group of displays is minimized. For example, if there are two participants and both are using the same zoom level, they either have the exact same window into the space or a completely distinct window. If the participants are using different zoom levels, there are only three possibilities. One person could see the entire space while his partner is zoomed in on a location; one person could use the middle zoom level and see the area his partner is zoomed in on; or the two viewing windows could have nothing in common.

These three options may make it easier for a pair of users to explore individually and then align their views for discussions. Each person knows that there are only a limited number of windows into the display that his partner could be looking at, and the differences between their views can easily be determined. Also, if their task pertains to a particular area in the space, there are probably only a few windows that can provide a view of that location.

In contrast, the continuous condition allows a user to navigate many more windows with much greater control over his display. The representation is divided into numerous overlapping windows. They overlap along the horizontal axis, vertical axis, and in between zoom levels creating the illusion that any view of the space can be obtained. Figure 4.2 portrays just a few of the possibilities:




Figure 4.2 Continuous navigation technique

This technique provides many windows into the representation. Users control the scale and position of their window through panning and zooming interactions.

In this condition, the number of windows into the space may seem infinite, but they are actually limited by the user interface. For example, scrollable interface widgets are ideal for adjusting the display, but they have definite end points and their movements are inherently discrete. Thus, they cannot provide an absolute continuous navigation. Similarly, the representation must adhere to the pixels available and the precision they provide. So using any style of interaction, the space can only be enlarged or shrunk to a certain degree of precision. Our approach is to provide a feeling of seemingly continuous control through the use of scrollable widgets with extreme end points and many discrete levels. This navigation technique provides each user with flexibility in what he can view. The display can easily be zoomed in to see small details as well as enlarged and panned to see the surrounding area.

4.1.1.3 Radar View Techniques

Radar views are a popular technique for providing awareness information in shared representations. They display a miniature of the entire workspace and indicate each collaborator's current viewport. This allows each user to visually understand his view with respect to the space and with respect to the other users' views. Yet, this simple technique could be less useful for more complex spatial representations. In spatial collaboration a representation is often navigated using panning and zooming techniques as well as levels of detail. The difference in these interaction techniques can present usability problems for radar views.

When users change the scale of their view through zooming, the size of their viewport representation in the radar view grows and shrinks. Large viewports are not a problem, but small viewports can be potentially difficult to use. A viewport representation may be too small to visually comprehend. Also, it may be difficult to determine if multiple, small viewports overlap.

Levels of detail in a spatial representation can also cause usability problems. Radar views often provide an abstraction of the shared workspace rather than display all of the details. This approach requires users to discuss their individual views and not just examine the radar view. Yet, in a spatial representation that incorporates zooming, the details in an individual's view

often vary with the zoom level. For example, a representation of a town might be zoomed so that building outlines are added to the display. This level of detail technique is useful for a single user, but it can lead to problems in a collaborative system. Two users could be viewing the same area with different representations. When the awareness information in a radar view only offers an abstraction of these views, the users have to resolve the differences through potentially lengthy discussions.

We designed a fisheye radar view approach to address these potential problems. This technique maintains the basic design, adding a fisheye projection feature that supports the zooming interactions and levels of detail used in complex spatial representations. Fisheye projections are techniques that offer details and an overview within the same representation (Furnas, 1982; Lamping, Rao, & Pirolli, 1995; Sarkar & Brown, 1992). Mathematical functions magnify certain areas of the representation while keeping the scale of the rest of the representation somewhat constant (Leung & Apperley, 1994).




Figure 4.3 Fisheye projections

Applying fisheye techniques magnifies and enlarges areas in a spatial representation.

The use of fisheye views to support collaboration has previously explored by Greenberg, Gutwin, & Cockburn, but their approach is different from our idea (1998). In an initial investigation, they found that applying fisheye projections within a shared space could potentially help users maintain awareness information. Their results suggested that fisheye techniques could be used as an alternative to radar views.

In our design, fisheye techniques are used within a radar view. Fisheye projections are applied to each collaborator's viewing window so their viewports are enlarged in the radar view. This magnification has two consequences. First, it increases the size of small viewport representations. Viewport representations that were too small to visually comprehend with traditional radar views are now magnified in the enhanced version. This allows overlapping views to easily be discerned and it provides a more usable miniature.

Enlarging the viewport representations also allows more detail to be displayed in the radar view. Radar views often only provide abstractions of the shared workspaces because too much detail creates a cluttered miniature. Yet, enlarging areas of the radar view allow more detail to be displayed. One approach is to use the extra space in each user's viewport representation to render more of the individuals' views, including the levels of detail that may not be displayed on the

others' screens. This will provide the users with greater awareness about each other's displays, while maintaining the information about their viewport positions. Thus, this technique reduces the potential for collaboration problems.

4.1.1.4 Radar View Implementations

The figures illustrate the traditional radar view (Figure 4.4) and the fisheye radar view (Figure 4.5) compared in this study. The user is working with a representation of the town of Blacksburg, Virginia and his viewport is outlined with a rectangle in the radar view. In the enhanced version, the viewport representation is magnified, creating a bubble effect.

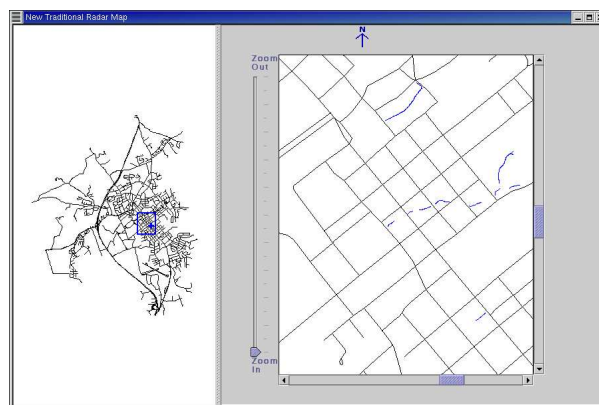


Figure 4.4 Traditional radar view

The user has zoomed into detailed area on the right. A traditional radar view is displayed on the left. A crosshair indicates the user's cursor location, although it is difficult to see.

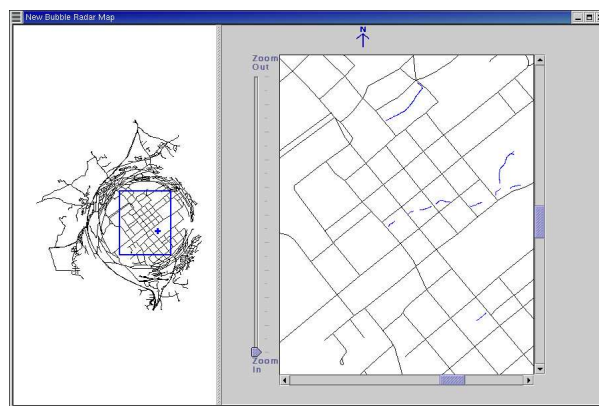


Figure 4.5 Fisheye radar view

Navigating to the same location; the fisheye radar view more clearly indicates the viewport details and current cursor location.

The mathematical function used to create this effect applies a linear transformation to each coordinate in the map. Points located within the circle of radius R that circumscribes the user's viewport are magnified, points in a larger circle of radius E surrounding the viewport are demagnified, and the remaining points are not altered such that:

$$\begin{aligned}
F(a) &= a / \sqrt{D} && (r < R) \\
&= \frac{a(R)}{r\sqrt{D}} + \frac{0.5 a(r-R)}{r} && (R < r < E) \\
&= a && (r > E)
\end{aligned}$$

This transformation is applied to the x and y coordinates separately as each are placed in the function for **a** where:

- **D** is the ratio of the user's viewport size to the size of the entire map ($D < 1$)
- **r** is the distance between the point to be transformed and the center of the user's viewport
- **E** is the radius of the larger circle, which corresponds to the distance required in order to demagnify by a constant scale of 0.5
- **R** is the distance between the center and one corner of the user's viewport

More specifically, the larger radius **E** is determined by the user's viewport size indicated by **D** and **R** such that $E = R(2 / \sqrt{D} - 1)$.

A weighted averaging technique handles multiple users to produce a multiple bubble effect. Using this approach, each point is determined by applying a weight to the results of each individual transformation. Our implementation follows Keahey's algorithm (Keahey, 1997; Keahey & Roberson, 1996).

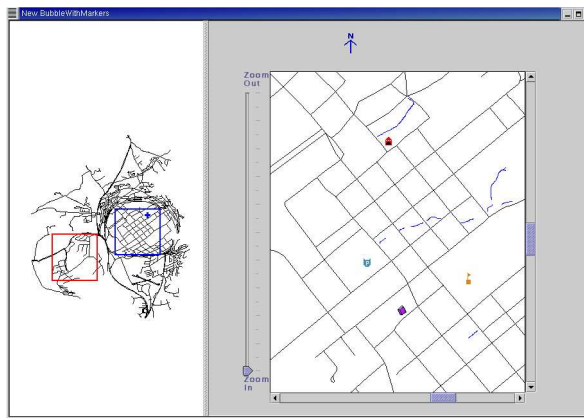


Figure 4.6 Multiple user fisheye radar view

A weighted averaging algorithm supports two users collaborating with the fisheye radar view, creating a multiple bubble effect on the left.

4.1.2 Task

In collaboration, people frequently “get together” to share ideas, have discussions, and make decisions. Similarly, in this experiment two users have separate ideas, which are provided by the researcher, and their task is to make a decision. The users are asked to work together to integrate these spatial ideas using a collaborative map tool.

The specifics involve positioning road signs and traffic lights on a town map. Each user is given a criteria statement about the location for a light/sign such that when read together a general solution is understood (Billinghurst, Kato, Kiyokawa, Belcher, & Poupyrev, 2002). For

example, one statement may say that people turn before they get to the hospital, while the other statement indicates that people headed south miss the turn. This reveals that a road sign should be placed along the road to the hospital north of the landmark.

The road sign tasks were intentionally open-ended to foster a more realistic spatial collaboration activity. Traffic light criteria, on the other end, led users to specific intersections on the map and often involved less discussion. This distinction is important as it investigates two common types of tasks in spatial collaboration: discussion and precision tasks. In collaboration activities, often the goal is to share and discuss ideas, not necessarily to make decisions. This corresponds to the road sign task where the users are given a partial set of spatial criteria. This situation causes them to discuss and share ideas in order to make a high-quality decision. On the other hand, spatial collaboration tasks can also be very precision-oriented such as pinpointing a set of data sites. This need for precision is reflected in the traffic light task. Users are asked to not only determine the intersection from the criteria, but to mark it with a circular ring icon so that the intersection is centered within the ring.

In collaborative work, people bring different ideas to a problem and during collaboration these ideas are often combined, expanded, and evaluated. To simulate this our task gives each participant a different set of criteria. The pair discusses these ideas and eventually reaches a solution that they feel satisfies the criterion. This solution is then specified as each user marks his/her individual map with the agreed upon location. This is desirable because it encourages both users to participate. Each user has both a portion of the problem definition and they have to individually indicate the group decision. Thus, the problem engages both participants at the beginning and end of the task.

Our experiment also explores the common understanding that develops among the participants while completing the task. In collaboration, often a group will believe they have reached a consensus when in actuality the participants' understandings differ. Having each participant mark the solution enables us to investigate this occurrence in distributed, spatial collaboration. The position differences provide one indication of the group's shared understanding and their level of agreement for a particular spatial problem. Typically, a usability-engineered groupware system would have incorporated more visual cues about where each partner was placing their marker or used a single, shared set of markers for this task. This enables the group to negotiate the precise positioning and establish a common solution. Our focus, on the hand, was to not provide explicit awareness information for the markers in an effort to investigate the larger issue of group consensus. Solutions to spatial collaboration problems do not always correspond to marking one location. Our task was simplistic in this respect, but provided an easy measurement for investigating which conditions lead to greater agreement, an indication of the shared understanding of the solution.

4.1.3 Interface

During the experiment, users worked with similar maps displaying roads and landmark icons corresponding to familiar places such as hospitals, schools, airports, etc. This data was acquired from the Internet and it provided the users with reference points. All of the data is displayed at each zoom level in order to eliminate any confusion or biases due to adding and removing levels of detail.

Maine towns were also used in order to eliminate any effects of prior spatial knowledge. Many of the users had never been to Maine and the few that had been a visitor in the towns did not recognize the town maps. In this way, the users had equal knowledge about the places they were navigating and discussing during the experiment.

Figure 4.7 shows a screenshot of the experiment in progress. Each of the landmark icons is indicative of the place it represents and they are generally familiar symbols (e.g. a capital 'H' set against a blue background corresponds to a hospital, an open book represents a library, etc.). Road signs were marked on the map using a star shape icon and traffic lights used a circular ring. These icons stood out on the road maps and were easily distinguished from the landmarks. The open area of the circular ring allowed the users to see the intersection they were marking and to be precise in the icon positioning. The star icon for the sign marker, on the hand, made the groups responsible for the level of precision, as they had to determine their own strategies to ensure they were marking the same location. To position a marker, users simply dragged the icons from the right-hand side onto the map area. Each icon was associated with a set of map coordinates once it was positioned on the map, so that the panning and zooming the map also caused the to icon move. The markers were not fixed to the map though and users could reposition them in completing the tasks.

The map itself supported zooming and panning interactions through familiar scrollbar and slider widgets. A scrollbar below the map allowed users to shift the map right and left, while another to the right moved the map up and down. The zoom slider was placed to the left of the map so as not to be confused with a panning interaction, which typically occurs to the right of a workspace (Schafer, Bowman & Carroll, 2002). In "continuous" navigation, the slider and scrollbars were implemented using 100 discrete steps, which gave the appearance of continuous panning and zooming. On the other hand, in discrete navigation, the scrollbars and slider snapped to the available positions. This occurred when the scrollbar thumb was released after being dragged, when the arrows to either side of the scrollbar were pressed, and when the open space surrounding the thumb was clicked.

Lastly, the radar view indicated each user's mouse cursor location using a telepointer. Gesturing is an important aspect of same place spatial collaboration and our implementation supports these physical actions in a distributed setting (Tang, 1991). Using a crosshair to represent each participant's cursor whenever the mouse was positioned within the map area allowed the users to point out spatial features and gesture ideas (Hayne, Pendergast, & Greenberg, 1994). The color of the telepointer corresponded to the viewport indication so that each user in the pair was represented by a specific color. The crosshair was not displayed when a user had a selected a marker and was dragging it across the map. This prevented the users from simply align their telepointers in positioning a marker. Instead, the groups were encouraged to discuss and agree on a location prior to physically positioning the marker. Once the marker was dropped, the crosshair reappeared in the radar view.

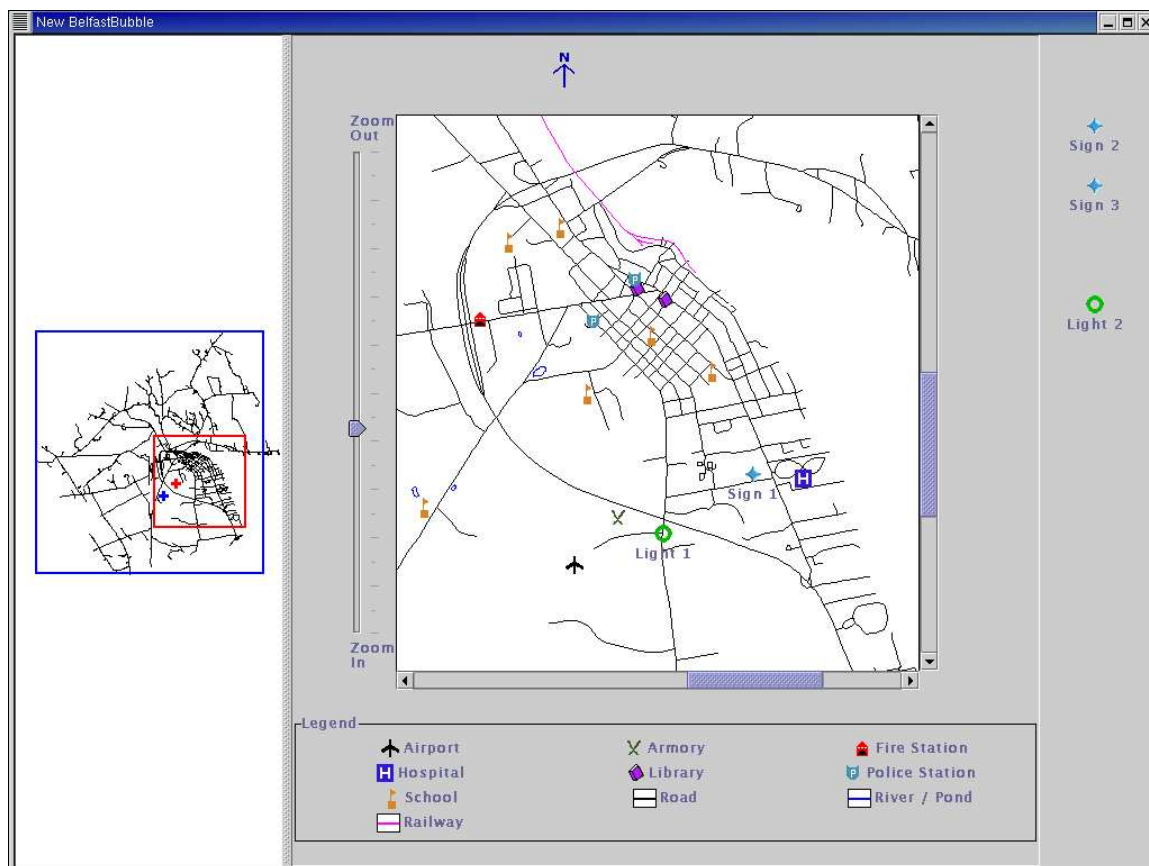


Figure 4.7 Two-dimensional map study prototype

One participant is viewing the entire town map while this user has zoomed in on a downtown area. A sign marker has been positioned near the hospital and a light marker has been placed around an intersection near the airport.

4.1.4 Experimental Design

The experiment used a 2x2 within-subjects design to evaluate four user interfaces. These interfaces corresponded to the four possible combinations of the two navigation conditions and the two radar view techniques. Each participant worked with the same partner to complete the four tasks and the ordering of the combinations was counterbalanced using a Balanced Latin Square. This allowed us to measure the groups at different stages in their familiarity with one another and in their familiarity with the combining criteria task.

The groups used the same set of Maine towns in completing the experiment. Each user interface condition was paired with a different town map. The maps were very similar in size and complexity in order to make the tasks as similar as possible. The specific marker trials of the towns were also designed to be comparable and pilot studies were used to ensure the criteria were clear and logical. Each group positioned three lights and two sign markers with each town map, generating repeated metrics.

Sixteen random pairs of people completed the experiment. Most were Blacksburg, Virginia residents, and a number of interests were represented including Biologists, English majors, Computer Scientists, Human Factors majors, and Ecologists. Twenty males and twelve females completed the task. The average age of the participants was 30, with the youngest person being 21 and the oldest 53 years old. Participants were randomly paired, resulting in six male-male pairs, two female-female pairs, and eight male-female pairs. Most participants did not know their partner prior to the experiment and others were simply acquaintances.

Each pair of users followed a similar procedure. After being introduced and reading through a set of instructions, the group completed the four trials. Each user sat in a separate room with an audio channel connecting the users and the experiment administrator. Users worked with identical PCs and could freely talk and hear responses without worrying about headphones or microphones. Before collaborating with each Maine town, the pair explored the upcoming interface with a familiar map of Blacksburg, Virginia. Then, after completing the Maine task, the users individually rated the treatment condition. At the end of all trials, a final questionnaire asked each user to rank the four conditions with respect to collaboration, navigation and preference. Informal debriefing discussions were also conducted before the group left.

4.2 Experimental Results

A number of different metrics were used to compare the different navigation and radar view combinations. A survey was administered at the end of the procedure as well as after each interface trial. The final questionnaire included questions that compared the four conditions, while the task rating sheets measured the participants' immediate feelings about each condition with respect to perceived effort, navigation support, and collaboration support.

Software logs maintained where each user positioned their markers. This allowed a calculation of the difference between each user's marker position for the same light or sign. The time required to position each marker as well as the total task completion time were also recorded.

Lastly, the sessions were videotaped and the discussions were partially analyzed. Each act of positioning a marker was categorized as using one strategy or another. The frequency counts of these strategies of use were analyzed as well as the strategy chosen for a particularly complex task. Frequency counts of pointing were also observed.

4.2.1 Final Questionnaire Results

In the final questionnaire, users ranked the four treatment conditions based on how easy/hard it was to collaborate with a partner, preference, and how easy/hard it was to navigate the map. The sum of ranks for each condition is displayed in the table below, where lower scores indicate higher placement in the ranking and, in turn, greater ease of collaboration, preference, and ease of navigation:

Table 4.1 Navigation and radar view ranking results

	Continuous		Discrete	
	Traditional	Fisheye	Traditional	Fisheye
Ease of Collaboration	58	66	98	100
Preference	57	51	103	109
Ease of Navigation	48	53	101	118

Users ranked the four conditions where a rank of 1 corresponded to the superior condition. The continuous conditions were ranked significantly lower. Within the same navigation condition, the sum of the rankings for the two radar views was statistically equal for ease of collaboration and preference.

Using the Chi-Squared Goodness of Fit Test revealed a significant difference in the four rankings with respect to ease of navigation (chi-squared = 25.05, df = 3, $p < 0.001$), preference (chi-squared = 51.38, df = 3, $p < 0.001$), and ease of collaboration (chi-squared = 68.21, df = 3, $p < 0.001$). Comparing the differences of the means, both continuous conditions were associated with significantly lower rankings in all three questions. The only significant difference between the fisheye and radar view conditions was within the discrete, ease of navigation rankings. The discrete, traditional condition, with a score of 101, was found to be significantly less than the 118 score for the discrete, fisheye condition.

4.2.1.1 Straightforward Questions

The final questionnaire also asked two sets of straightforward questions focused on the navigation and radar view conditions. Both sets of questions had users indicate which design was better for collaboration, which was easier to use, and which the user preferred. Three choices were provided for each set of questions, including the two treatment conditions and a neither (about the same) option. The results from the twenty-eight subjects who answered these questions are summarized in the tables below:

Table 4.2 Subjective navigation results

	Discrete Navigation	Continuous Navigation	About the Same
Which was better for collaboration?	8	16	4
Which did you prefer?	0	27	1
Which was easier to use?	0	28	0

Final questionnaire results when users were asked to choose between the two navigation conditions. Users were divided on which was better for collaboration, but clearly favored the continuous navigation in terms of ease of use and preference.

Table 4.3 Subjective radar view results

	Traditional Radar View	Fisheye Radar View	About the Same
Which was better for collaboration?	13	9	6
Which did you prefer?	14	8	6
Which was easier to use?	18	5	5

Final questionnaire results when users were asked to choose between the two radar view conditions. Statistically, users did not favor either condition for collaboration or in terms of preference. They found the traditional radar view easier to use.

Using the results where one technique was chosen over the other, there was not a statistical difference in the navigation condition in terms of collaboration (chi-squared = 2.67, df = 1, $p < 0.10$). Yet, users identified the continuous technique as both easier to use (chi-squared = 28.00, df = 1, $p < 0.001$) and preferable (chi-squared = 24.14, df = 1, $p < 0.001$). For the radar view section, there was not a statistical difference for the collaboration question (chi-squared = 0.73, df = 1, $p < 0.50$) or the preference question (chi-squared = 1.64, df = 1, $p < 0.30$). However, users considered the traditional radar view easier to use (chi-squared = 7.35, df = 1, $p < 0.01$).

4.2.1.2 Equal Contribution

The last part of the final questionnaire investigated whether each user contributed equally to the task as users rated their partner's effort in comparison to their own using a 7-point scale. The participants indicated an equal level of effort more often than they indicated any other effort level (chi-squared = 11.57, df = 1, $p < 0.001$). The responses from the twenty-eight participants who answered this question are summarized in the table below:

Table 4.4 Participant effort differences

	More effort 1	2	3	4	5	6	Less effort 7
Partner's effort level in comparison to own efforts	0	2	2	23	1	0	0

Final questionnaire results when users were asked to rate their partner's effort in comparison to their own efforts. Users indicated an equal level of effort within the groups.

4.2.2 Strategies of Use

Each group's session was captured using video recorders during the study and reviewed afterwards. One observation of these recordings relates to the two strategies groups would use in deciding the final positioning of marker. Frequently, a pair would have an intricate discussion about the shape of the roads with respect to where they were going to place the marker. This was the discussion strategy. Other times they would use the radar map to point out various spots when deciding on a location. For example, one user would indicate the spot and the other would confirm the location using his cursor or by mentioning a few nearby landmarks. This was the pointing strategy.

Each sign and light marker task was coded with one of these strategies by the experimenter as they reviewed the videos. When the pair used the telepointers during their decision-making this was classified as a point strategy. No use of the telepointers and a complete reliance on speech counted as a discussion strategy. This analysis did not require multiple coders because the two categories were straightforward.

Analyzing the frequency counts of all the sign and light tasks, users favored a discussion strategy overall and the number of discussions was greater with the traditional approach ($F = 14.24$, $df = 1$, $p = 0.004$). Using a pointing strategy occurred more often with the fisheye design. The average counts and standard deviations for each condition are displayed in the table below:

Table 4.5 Radar view and marker strategy results

	Discussion strategy	Pointing strategy
Traditional	2.36 (s.d. = 0.78)	0.64 (s.d. = 0.78)
Fisheye	1.93 (s.d. = 1.22)	1.07 (s.d. = 1.22)

Average frequency counts and standard deviations for strategies used in positioning the three sign markers. Discussions occurred more often with the traditional radar view and the pointing strategy was greater with the fisheye design.

Examining a similar set of results for the navigation conditions reveals no differences in the strategies. The average counts and standard deviations for each condition were:

Table 4.6 Navigation and marker strategy results

	Discussion strategy	Pointing strategy
Discrete	2.15 (s.d. = 1.15)	0.85 (s.d. = 1.15)
Continuous	2.15 (s.d. = 0.85)	0.85 (s.d. = 0.85)

Average frequency counts and standard deviations for strategies used in positioning the three sign markers. Users did not favor a discussion or a pointing strategy within the navigation conditions.

4.2.2.1 Complex Marker Task

The criteria for each town involved positioning at least one sign in a congested area of the map, or an area with many roads very close to another such as a downtown area. Using a Conchran Q Test to determine any differences in the strategy used with respect to the four conditions, users favored a pointing strategy more often with the fisheye, discrete navigation combination ($Q = 13.8$, $df = 3$, $p < 0.005$). The total frequency counts for the pointing strategy are displayed in the table below:

Table 4.7 Complex task and pointing strategy differences

	Discrete	Continuous
Traditional	1	0
Fisheye	7	3

Frequency counts for using a pointing strategy when positioning a sign marker in a congested area of the map. Pointing was more prominent with discrete navigation and a fisheye radar view map.

4.2.2.2 Confirmation Actions

Reviewing the video recordings, participants were also observed with respect to confirmation statements. During the experiment one user would frequently point out a location to their partner by placing their cursor on the point of interest and making a verbal comment. When the partner responded to this statement and confirmed the location by placing their cursor on the same location, this was noted as a confirmation action. The experimenter tallied counts of confirmation actions for each town task. A confirmation action was only counted when the partners matched their telepointers in the radar view. Again, this analysis did not require multiple coders because an explicit action was required.

Analyzing the frequency counts, there was a trend for more confirmations in the discrete navigation than the continuous tasks ($F = 3.84$, $df = 1$, $p = 0.072$). The table below summarizes the use of confirmation statements, displaying the average counts and standard deviations for each condition:

Table 4.8 Confirmation action results

	Discrete	Continuous
Traditional	5.71 (s.d. = 1.73)	4.50 (s.d. = 2.90)
Fisheye	6.57 (s.d. = 3.20)	4.86 (s.d. = 3.16)

Average frequency counts and standard deviations for confirmation actions in positioning both light and sign markers. Users confirmed a location by matching their cursor with their partners' significantly more often with the discrete navigation.

4.2.3 Task Rating Sheets

At the end of each treatment, users rated the condition or interface they just used. This rating sheet contained three sections: Perceived Effort, Navigation Support, and Collaboration Support. Answers in each section were based on 7-point, Likert-like scale and averaged for an overall perceived effort, navigation, and collaboration rating. The results from each section follow.

4.2.3.1 Perceived Effort Ratings

The perceived effort questions asked how hard the task was to complete, how much effort the task required, and how hard they had to concentrate. Conducting an Analysis of Variance, users overwhelmingly indicated that continuous navigation corresponded to less effort ($F = 32.67$, $df = 1$, $p < 0.001$). There was also a significant interaction between the radar view and navigation conditions ($F = 4.46$, $df = 1$, $p = 0.044$). In the discrete case users rated the fisheye tasks with greater effort ratings. Yet, in the continuous case the fisheye was associated with lower ratings.

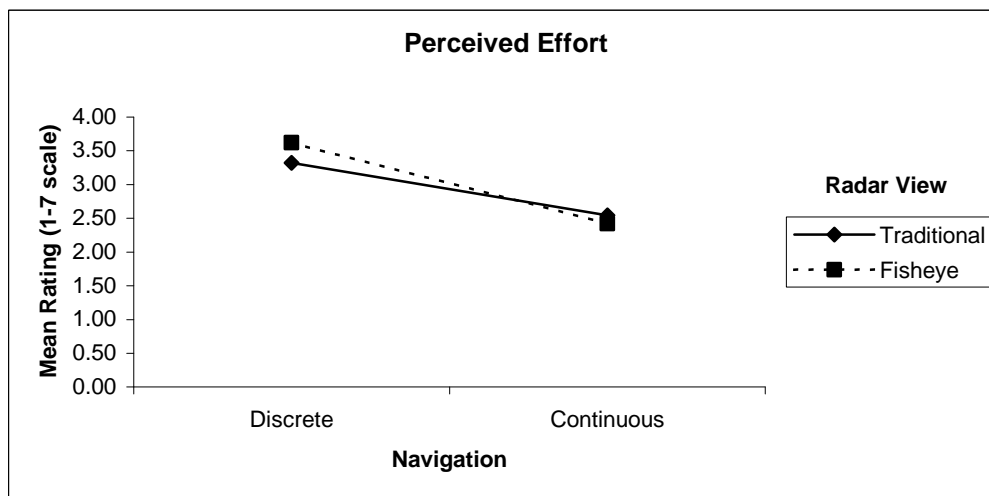


Figure 4.8 Perceived effort interaction

There was an interaction between the navigation and radar view conditions with respect to perceived effort ratings. Users felt that the fisheye view required less effort with the continuous case.

4.2.3.2 Navigation Support Ratings

The navigation questions focused on the user's experience in adjusting the map, finding locations, and align one's view with his partner's. Continuous navigation was associated with higher ratings, implying that the continuous interface was easier to navigate ($F = 108.07$, $df = 1$, $p < 0.001$). Results for the traditional radar view also involved higher ratings than the fisheye radar design ($F = 22.62$, $df = 1$, $p < 0.001$).

There was also a significant interaction between the navigation and radar view conditions ($F = 13.32$, $df = 1$, $p = 0.001$). The mean for the fisheye view with discrete navigation was much less than its traditional radar view counterpart, but in the continuous case the fisheye mean was closer to the level of the traditional radar mean.

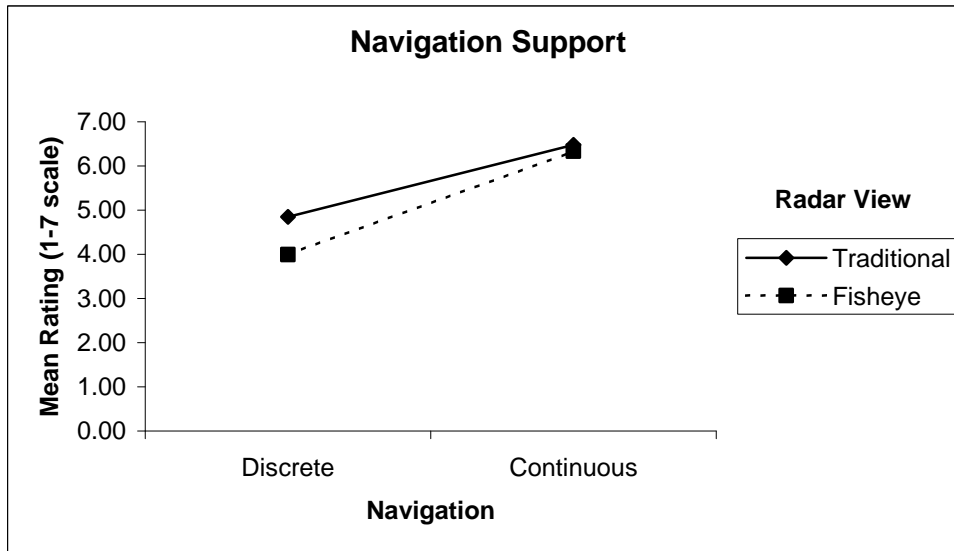


Figure 4.9 Navigation support interaction

There was an interaction between the conditions for the navigation support questions. Continuous navigation received higher navigation support ratings and the fisheye technique was associated with higher ratings when combined with continuous navigation.

4.2.3.3 Collaboration Support Ratings

The collaboration questions explored the communication between the two users, whether they could easily understand one another, how easily they could communicate their ideas, and how helpful they found the radar view. Analyzing the average collaboration rating, there was not a significant difference in the results, although there were differences in some of the individual questions. Users responded in favor of continuous navigation when asked about how easily they could work with the other person ($F = 5.42$, $df = 1$, $p = 0.027$), how easily they could understand their partner ($F = 7.24$, $df = 1$, $p = 0.011$), and how easily they could communicate their ideas ($F = 10.84$, $df = 1$, $p = 0.003$).

4.2.4 Task Completion Times

At the conclusion of the study, the video recordings were also used to determine the amount of time each group required to complete the tasks. An analysis was conducted on the amount of time spent on an entire town or treatment condition, as well as the average times used to position the light and sign markers within a town. Users took significantly more time on the towns with the discrete navigation technique ($F=9.29$, $df = 1$, $p = 0.012$). The average task completion times and standard deviations for a town are below:

Table 4.9 Task completion time results

	Discrete (min:sec)	Continuous (min:sec)
Traditional	15:14 (s.d. = 5:41)	14:39 (s.d. = 4:37)
Fisheye	16:51 (s.d. = 4:14)	14:11 (s.d. = 3:55)

Users collaboratively positioned three sign and two light markers with each condition, spending a significantly greater amount of time with the discrete navigation conditions.

The time involved with positioning a light marker was also significantly greater with the discrete navigation technique ($F = 9.28$, $df = 1$, $p = 0.011$). Another interesting finding was a non-significant interaction between the conditions for the time used in positioning a sign ($F = 3.60$, $df = 1$, $p = 0.079$). The fisheye condition required much more time than the traditional radar view for discrete navigation and less time than the traditional in the continuous case.

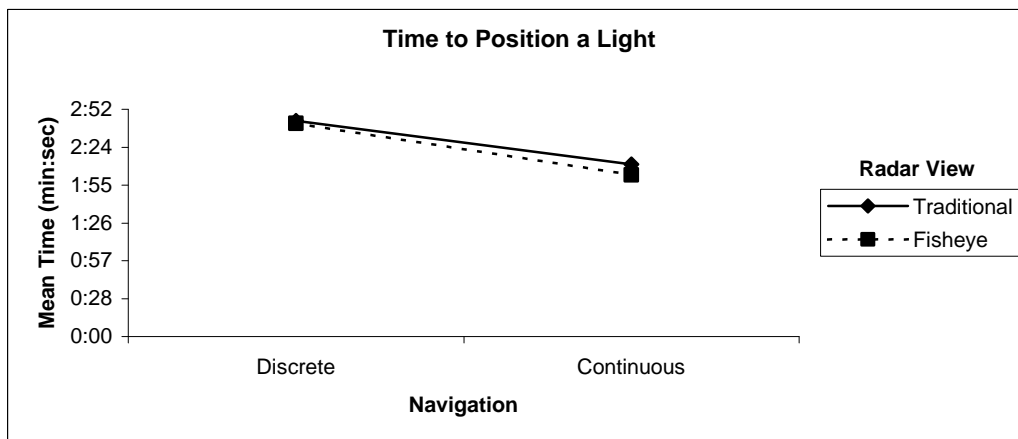


Figure 4.10 Time to position a light marker

The time required to position a light was significantly greater with discrete navigation.

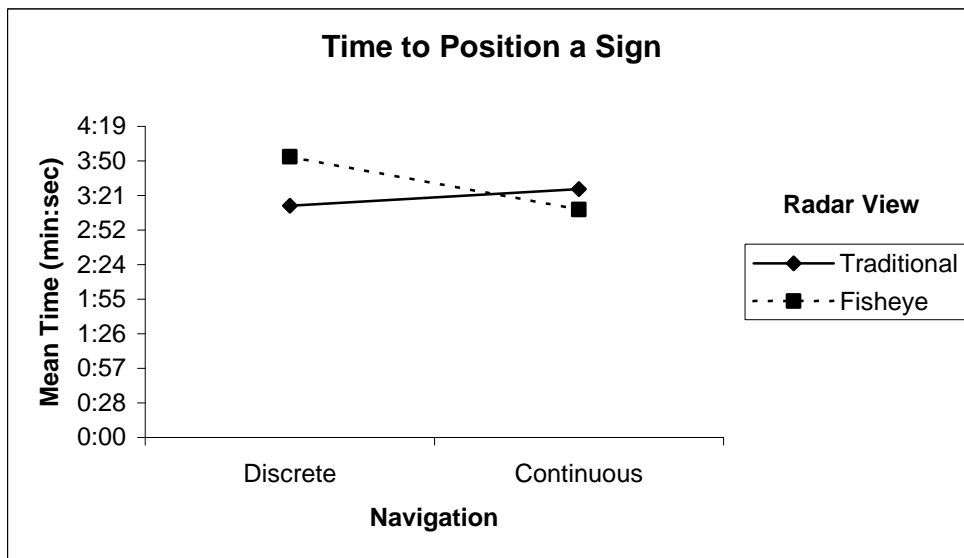


Figure 4.11 Time to position a sign marker

Non-significant interaction between the conditions for the time required positioning a sign. Users took less time to mark a sign with the fisheye radar view in the continuous case.

4.2.5 Agreement Metric

Using the log files generated during the experiment, an analysis was conducted to see if the conditions had an impact on how close the users placed their markers to one another. A statistical analysis of the distances between all light and sign marker pairs did not reveal any differences. Yet, subjects placed the traffic light markers closer together in the discrete condition ($F = 5.19$, $df = 1$, $p = 0.039$).

4.3 Discussion of Results

These results provide insight into the advantages and disadvantages of the two navigation and radar view techniques. We observed usability problems with the techniques as well as positive and negative effects from the different combinations. A discussion of these results follows.

4.3.1 Usability issues with discrete navigation

One of the main findings from this evaluation is that users overwhelmingly prefer continuous navigation and they also complete tasks faster with this technique in comparison to the discrete approach. In the results, the participants completed the tasks significantly faster with the continuous navigation and positioned the light markers quicker (Section 4.2.4). They consistently rated the discrete navigation as more difficult to navigate and requiring more effort in the final questionnaire and task rating surveys. And, they consistently reported continuous, not discrete, navigation as the preferred approach. During the evaluation, we also observed users struggling particularly with the discrete navigation conditions.

One explanation of this finding is that the discrete technique requires users to perform different spatial transformations than those required by the continuous technique. In our implementation, there was a minimal overlap between the viewpoints. Viewpoints available within the same

zoom level were non-overlapping and only changes to the zoom level displayed any overlaps. This requires a single user to spatially comprehend each discrete movement by maintaining a mental representation of the initial viewpoint in order to understand the new viewpoint. For example, in panning a user had to make a mental image of the map features that extended off the edge in the current viewpoint. Then, after switching the viewpoint, they could match up their understanding of the old viewpoint with the features along the similar edge in the new viewpoint. Similarly, in zooming, users needed to remember what they were looking at beforehand when trying to understand the new viewpoint. It was also important to spatially comprehend that zooming in adjusts the map to the center of the previous viewpoint, while zooming out adjusts the viewpoint to a preset position in a higher zoom level.

The effect of these mental transformations were apparent during the study as users repeatedly zoomed out or switched back and forth between two adjacent viewports in order to comprehend the space. The tasks used in the experiment were designed to be somewhat challenging and so many of the references provided could not be viewed within the same viewport. This caused users to spend more time navigating with the discrete condition because they had to spatially comprehend the references as well as the overall space in order to solve the problem.

Similarly, the discrete condition also forces the users to make use of less than ideal views. Discrete navigation involves a set of predefined viewpoints and often these viewpoints cause interesting map features to be physically located near the corners and the edges of viewpoints. Such positioning limits the view of the area surrounding the feature. Our implementation was particularly frustrating to users because each point on the map was only viewable from one zoomed in viewpoint. Often, the users desired a more complete view of the feature that was only available through navigating to the adjacent viewpoints and by building a mental model of the surrounding area.

4.3.2 Approaches for collaborating with discrete navigation

These usability issues with discrete navigation were often observed with just a single user. Yet, the discrete conditions also had an effect on the pairs of users. This navigation condition caused the groups to take a different approach to completing the tasks, an approach that required more time and effort.

In order to deal with the more complex discrete conditions, the users would often engage in collaborative navigation to complete the tasks. For example, one user would frequently mark the agreed upon location at a zoomed out level by keeping their cursor over the spot and having the radar view share this location with his/her partner. The other user would use this telepointer as a guide to navigate to the corresponding zoomed in view. The first user could then navigate their map to match his/her partner's viewpoint by observing the window differences in the radar view or by aligning his cursor with his partner's cursor location in the radar view. In this way, neither user would get lost and they both would end up with the designated location on their displays. The discrete case provided them an easy way to match up their viewpoints with its limited view options. Yet, this strategy required a good deal of effort on behalf of the group.

Another approach that was less successful involved zooming in on a particular location or landmark. At the most zoomed out view, the users would agree on the location and each would

independently try to navigate to the zoomed in view. This was often problematic because one user would inevitably find the desired view faster, leaving his partner to try and locate it on his own or rely on the faster user's description of the view. Often we observed that the lost user did not see the identifying landmarks the partner mentioned. For example, in one group we watched the person trying to match his partner's view navigate to the correct viewpoint multiple times and yet not recognize that it was the same view as their partner. This approach was often very time consuming as it took time for the pair to reach the same viewpoint.

The groups also struggled with discrete navigation in making sure all of the criteria of the tasks were met. In some instances the groups were successful, as each user examined a different viewpoint and the radar map was used to point out and discuss the references. In other cases this effort was less mutually collaborative. Very often we observed that one partner took a leadership role. This person typically had the better discrete navigation skills and could create a mental map of the space with greater ease. Yet, this leader was faced with the difficult task of describing their spatial thought processes to his partner so he/she could see the same idea. Often times the partner listening was not interested or did not take the time to understand the intricacies of the explanation. He would simply mimic the leader's actions and match cursor locations using the radar map. This resulted in a poor understanding of where and why one user wanted to place a marker. It was also undesirable because the task was to collaborate and yet the interface likely caused some solutions to be the product of one user.

Many of the results from this study agree with these observations. The groups required more time in completing the discrete tasks and the users rated discrete navigation poorly with respect to ease of navigation (Section 4.2.1), perceived effort (Section 4.2.3.1), and preference (Section 4.2.1 and 4.2.1.1). However, there were also some positive effects of discrete navigation. The rating sheet and final questionnaire results provide one indication. In the collaboration section of the Task Rating sheet, users did not consistently rate the discrete tasks as providing less collaboration support. A few questions had significant results in favor of continuous navigation, but overall the two navigation conditions were rated statistically equal. Similarly, an analysis of the straightforward question on the final questions asking which navigation approach was better for collaboration did not reveal a difference between the two conditions.

Analyzing some of the other results, discrete tasks involved more confirmation actions (Section 4.2.2.2) and, when working within a congested area on a discrete map, users favored a pointing strategy (Section 4.2.2.1). This indicates that the groups used more pointing approaches to collaborate under the discrete conditions. One explanation for this relates to ability to have the same viewpoint. Discrete navigation enables users to easily align their viewpoints where there are a limited number of possibilities. With two viewpoints aligned, both users know they are looking at the exact same view so they freely point out features to one another. It is impossible for their collaboration problems to be a result of differing views as every feature one user sees or points out is definitely displayed on his partner's screen. In this way, increased pointing with the discrete navigation is a positive effect as users were taking advantage of the easily aligned viewpoints.

Our implementation also led to greater agreement among the participants where the paired light icons were positioned closer in the discrete case. One explanation is that the users zoomed in

completely with the discrete conditions but not with the continuous cases, leading to less precise marking with continuous navigation. This result could also be due to inherent precision opportunities in discrete case. When each person's view is using the same underlying transformation function to convert a pixel to a point in the map coordinates, it is possible to mark the exact same location and achieve a zero difference. Two different transformation functions, on the other hand, can result in very small differences when the users follow extremely precise procedures. This is the case with the continuous navigation as the views most likely are not exactly aligned and are not using the same transformation.

4.3.3 Advantages of continuous navigation for collaboration

The continuous technique better supports this particular collaboration task. The continuous approach has less usability problems, and our results clearly indicate that users find it easier to navigate (Section 4.2.1 and 4.2.3.2), preferable (Section 4.2.1 and 4.2.1.1), and requiring less effort (Section 4.2.3.1) in comparison the discrete approach. Using this technique, users can adjust the display with both large and fine grain movements enabling them to maintain a context of the space without having to perform mental transformations. Similarly, users can easily examine more of the surrounding area where a range of adjustments allows the user to keep an interesting map location within his view. Other map tools also commonly use continuous navigation or discrete techniques that provide many overlapping windows, so users may be more comfortable using this technique.

Many of the results from this experiment also agree that continuous navigation is preferable for collaboration. In the final questionnaire, users ranked the continuous conditions as providing greater ease of collaboration and navigation. Similarly, users answered many of the collaboration questions on the Task Rating sheet in favor of continuous navigation, even though each of these questions specifically asked about how easily one could work with their partner. Users also performed the town collaboration tasks faster using continuous navigation. To summarize, all of the results collected were either in favor of continuous navigation or the two conditions were statistically equivalent.

One explanation for these results is that the greater ease of use enables an independent exploration of the space. When users can navigate the map on their own, the collaborative effort is more salient as both people are contributing and the task itself is perceived as easier. Also, as the users understand their own view, they are better prepared to understand their partner. They do not struggle to comprehend their individual map during the task, but rather can focus on what their partner is saying and collaborate efficiently. The views produced with discrete navigation, on the other hand, are at times very confusing to the users and this distraction most likely limits the ease with which they can work with their partner.

It also could be easier to understand a partner in continuous navigation, because users can point out locations without having to align their individual views. Granted, the results indicate that groups used more pointing strategies with the discrete condition. But in discrete navigation there are less ways the views can overlap. The viewing windows must be identical or one view has to be zoomed out and include the other window. In this condition, very often the users had to explicitly align their viewpoints before they could share ideas. The overlapping feature of

continuous navigation, on the other hand, enables multiple users to be able to share map locations even if their viewpoints are different.

4.3.4 Other Navigation Possibilities

Considering the positive and negatives of both the discrete and continuous navigation techniques we evaluated, it is interesting to reflect on other navigation possibilities. Limiting the discrete navigation to a few viewpoints was beneficial to spatial collaboration because the participants could easily achieve the same view and point out features, but navigating with no overlaps had a negative effect on their spatial understanding. Repeating the study using overlapping windows might reverse this second finding, but this is not obvious. Keeping the number of viewpoints the same and simply using an overlapping design as in Figure 4.12 would provide some spatial references as the users navigated to adjacent viewpoints. Yet, this approach still would require spatial, mental transformations and taking note of the features along one side because the differences between viewpoints remain great.




Figure 4.12 Discrete navigation modification

A modified version of the discrete navigation technique with overlapping windows still requires users to perform spatial, mental transformations with only a few viewpoints.

Our continuous navigation implementation, on the other hand, does not suffer from this problem. It allows the users to navigate with both minimal and large viewpoint shifts. However, it suffers from difficulties in aligning multiple viewing windows and ensuring the collaborators know they are referencing the same features. This is what the discrete navigation technique supported with its limited viewpoints. Thus, this encourages an integration of the techniques to support both flexible navigation and aligned viewpoints.

One design possibility is to modify the current continuous navigation to enable synchronized viewports. This would allow the collaborators to easily align their individual views and be confident in the similarities. Yet, such a technique also presents new design issues. How long does the synchronization last? What happens when the users navigate? How is the synchronization decision negotiated? And, most importantly for spatial collaboration, will the synchronization disrupt an individual user's spatial navigation thoughts, causing him to become disoriented or forget where he was previously working and/or what he was doing?

Another approach is to change or improve the awareness techniques in order to better complement continuous navigation. Our investigation strictly focused on using a radar view to provide awareness information and examined two approaches to this technique. Yet, the radar views did not solve all of the awareness requirements of the task. The collaborators continued to have issues with what their partner could see and what they were directly looking at. One possible improvement is to provide additional awareness information on the map interface itself.

Subtle indication of telepointers and viewport bounds in the primary interface would make the awareness information more salient. Again, this solution presents new design issues, as a spatial representation can already be complicated with many details and it is important that the awareness cues are not confused with the map data.

4.3.5 Negative effects of the fisheye, discrete combination

In analyzing at the interactions that occurred between the navigation and radar view conditions, it is interesting to consider the discrete navigation, fisheye radar view combination. This pairing produced many poor results. In the final questionnaire, users ranked the discrete, fisheye tasks as the hardest to navigate. Similarly, in the task ratings, the discrete, fisheye tasks required the most effort and provided the least navigation support. Lastly, there was a trend in the task completion times for positioning a sign marker where the discrete, fisheye condition took the longest amount of time.

There are a couple of explanations for these results. One way this combination creates problems relates to the fisheye's distortion. The fisheye technique affects the size and shape of the viewpoint, not providing an accurate view of the users' zoom levels or the closeness of their viewing positions. During the experiment, users in the discrete condition would often direct each other on how to navigate their viewpoint. This is not as straightforward with the fisheye's distortion. It is not as obvious how two viewpoints differ and nor the navigational steps to take to align viewpoints. Users would often rely on trial and error to deal with this situation.

The discrete-fisheye combination is also problematic because the radar map is less useful in understanding the space. In general, discrete navigation does not provide a context for the space with its non-overlapping viewpoints. Because unlike continuous navigation, users have to build a mental map of the layout as they navigate. Combining this navigation technique with a traditional radar view provides users with a visualization of the entire space, reducing the need for a mental map. Using a distorted radar view with discrete navigation, on the other hand, limits the usefulness of the visualization for understanding the layout. Users must rely on their mental image of the space and determine the layout by observing the viewpoints displayed through navigation.

4.3.6 Performance and usability of the fisheye design

The results of this study do not clearly distinguish between the fisheye and traditional radar view design, yet they indicate the potential for a fisheye design and encourage further design and evaluation. In the interaction analyses, the fisheye design was associated with less effort (Section 4.2.3) and faster completion times in positioning a sign (Section 4.2.4). It also corresponded to greater use of a pointing strategy (Section 4.2.2) implying that the awareness tool is usable, efficient, and helpful for spatial collaboration.

It is particularly interesting to consider that the use of the discussion strategy was significantly greater with the traditional radar view, while the pointing strategy was used more frequently with the fisheye radar view (Section 4.2.2). This indicates that the users relied on spoken discussions when the radar view did not provide them the awareness information that they required, particularly in deciding where to position a marker. This was directly observed in the experiment, where one participant would try to point out a location and the other would be

unsure of the location, asking for a more detailed description of the spot or suggesting a different spot that perhaps both participants could easily recognize. The fisheye approach, on the other hand, enables one user to point at a location and the other user to understand where his partner is pointing with minimal verbal description. With its enlarged and pronounced viewport representations, the fisheye design affords pointing interactions. Each user knows that their partners have a magnified version of their view so they often take advantage of this feature and express ideas through pointing.

It is also interesting to consider the interaction effects observed with respect to perceived effort ratings and task completion time. These results indicate that the utility of the fisheye approach differs with varying navigation conditions. Continuous navigation was statistically either the superior technique or equal to discrete in all of the analyses. Considering the fisheye's performance with the continuous navigation implies that it better supports collaboration, although the differences were not statistically significant.

Regardless of the differences in strategy and the interaction effects, many of the users in this study were unsure about the usefulness of the fisheye design. This was apparent in the final survey as users were divided on which condition was better for collaboration and which they preferred (Section 4.2.1 and 4.2.1.1). During the trials, some users noticed the difference the bubble made while working together to decide where to place a marker. This group realized that the bubble was magnifying their viewport representations and they often expressed positive comments. Other users saw the traditional radar map as more intuitive and simply more pleasant to look at. To these users, the fisheye technique added another element of complexity for the collaboration activity.

The few negative results in the experiment for the fisheye approach were that the users found the traditional radar view easier to use (Section 4.2.1.1) and that the traditional condition provided greater navigation support (Section 4.2.3.2). This again may be because the traditional radar map is more intuitive to novice users than the fisheye view and its distortion effect. Possibly with further use, this and other user ratings would change. However, the fisheye condition can also create problems with respect to how users point out features to one another. The transformation of the fisheye view may be confusing in that it is not clear whether the telepointer's position has also been distorted. Also, when one user points out a location to the other, it may be more difficult to determine where they are pointing if the telepointer falls within the de-magnified part of the distortion directly adjacent to the magnified area. This can happen when one user is zoomed in and the other is not.

To improve the fisheye radar view design, we are exploring ways to modify the transformation function and add visual cues. One improvement would be to determine a more appropriate technique for combining multiple fisheye magnifications. The current approach of using a weighted average of the individual fisheye transformations is problematic in that the presentation does not always correspond to expectations, where 3D distortion techniques offer one alternative (Carpendale, Tigges, Cowperthwaite, Fracchia, 1998). Also, adding visual cues to the fisheye display could better inform the users about zoom level differences and the transformation's effects. One approach to reduce the confusion is to associate the space with a checkered grid or simply a line grid (Carpendale, Cowperthwaite, Fracchia, 1997). Drawing such a grid so that it

does not overwhelm the map content and it is visible behind the magnified portions of the radar view could better indicate the magnification and compression effects.

4.4 Implications

In conducting this experiment, a number of lessons were learned about spatial collaboration, in general, as well as distributed, spatial collaboration situations. We enumerate some issues to consider in designing spatial collaboration applications and review some implications for future collaborative map software programs.

4.4.1 Issues in designing for spatial collaboration

In the experiment, the participants' differences in spatial abilities, thought processes, and problem solving approaches were clearly apparent. We anticipated some of these differences and planned to explore their influence on the collaboration task, but they were highly variable even in our controlled study. For instance, some users understood how the discrete navigation worked and offered suggestions for improvement, while others clearly did not understand this technique as they completed the tasks using trial and error.

Many pairs of users also approached the tasks differently than we expected. We provided the participants with a paired listing of criteria for the sign tasks and the light tasks. Yet, some groups choose to complete the light and sign trials in different orderings and a few groups had considerably different interpretations of the criteria. One group even located all the references in the criteria on the map before stepping through the individual tasks. We had not anticipated these approaches and often had to take them into consideration when compiling the data. Another example of spatial differences was with the notion of precision. One group had an extreme definition of precision and were observed counting pixels. Other groups would simply make sure their crosshairs were aligned in the radar view, while others positioned the markers so the tips of the icons marked the same locations. Again, this range in variability was not expected, but it reveals how people tackle spatial problems differently.

During the experiment, we also observed many groups defining their own terms and procedures as they progressed through the trials. For example, one group repeatedly referred to placing the tip of the marker in the "crotch of the Y" formed by two roads. Both participants knew what this meant and they could easily align their marker positions using this funny phrase. This strategy was less successful with another pair as they became confused when there was more than one "alien firehouse". The group used this phrase to refer to a firehouse that looked like it had antennae formed by nearby roads. The problem was that another firehouse had a similar look and the group did not realize that at times they were talking about two different locations. This reveals the need for collaborators to have a shared understanding of the space. These groups established an understanding through fun terms that referenced specific locations on the map. Similarly, spatial collaboration solutions need to consider how this shared understanding of a space will develop in an activity. Designs should support a common understanding and reduce confusion, such as the problem with two "alien firehouses".

We also observed different group dynamics in our pairs of participants. Some groups had a natural leader who took charge and helped the other person through the tasks, while other groups

took a more collective approach. In one instance, a very capable user let his partner lead the way because the partner was not listening to his ideas. It was easiest for this user to let his partner lead the effort and offer his confirmation. Some of the group dynamics were attributed to differences in spatial abilities and other times the group had a natural leader or a set of team players. In any case, styles of collaboration can differ between groups and this is important to consider in designing future spatial collaboration solutions. Our distributed design lacked features that allowed the users to help one another when someone became lost. Yet, it was supportive of different styles of independent and collective work as the software did not limit the users to specific roles and it enabled both users to explore the representation.

Lastly, this study revealed that collaborators might become very focused on a spatial visualization while working together. For instance, during the experiment we observed the participants looking very intensely at the map while completing the tasks. The trials would often start with each user reading their criteria out loud and occasionally the pair would consult the criteria to confirm specifics, but most of the users' time was spent navigating, discussing the space, and agreeing on the marker location. Groups would intensely work with the collaborative map to solve the problem we had given them. Other tasks may not result in a similar behavior, but it is important to recognize that some spatial collaboration tasks can be very representation intensive.

4.4.2 Issues in designing distributed, spatial collaboration solutions

In conducting this experiment, we also learned about issues specific to a distributed setting of spatial collaboration. For instance, our task turned out to be very audio-intensive. Even with visual cues provided in the radar map, one user had to listen closely as the other described his idea or read his clue. We also observed non-native English speakers struggling to communicate with their partners. Users would often express their ideas verbally and language issues often caused pairs with a non-native English speaker to require more time in conveying their ideas. Lastly, our participants would often have to stop everything they were doing and listen to their partner. They found that they could not comprehend an idea and explore the map simultaneously. This heavy reliance on auditory information may be specific to our task, but it makes a good point about designing for distributed, spatial collaboration. Namely, it is important to consider how the collaborators will communicate spatial ideas. Designs should support the expression of ideas that involve distance, relative direction, and precise positioning. Our design was limited in the ways people could share ideas. The participants could either express themselves verbally or point and have their partner follow their telepointer on the radar map. These restrictions were partly due to the metrics we wanted to collect, but future designs could also implement telepointers within the primary map, indication or sharing of markers, and simple annotation tools.

Through this experiment, we also realized the complexity of developing one's own ideas with an individualized interface and also staying aware of what your partner(s) are thinking and investigating. Some of our groups faced collaboration issues when the participants explored the space in parallel. One user would often figure out part of the task and report back causing the other user to lose their train of thought. Supporting this parallel style of work seems more complicated in spatial collaboration as users' ideas are more spatial in nature and often require visualization. Distributed, spatial collaboration solutions should keep this complexity in mind

and provide ways for users to coordinate parallel activities. One idea is to allow users to create bookmarks for their viewpoints so they can easily return to a previous view. Another feature that allows one user to quickly navigate to a partner's view may also facilitate this coordination.

4.3 Lessons learned about using representations for distributed, spatial collaboration

In addition to learning more about the issues with distributed, spatial collaboration, this experiment revealed some key insights about using a map representation in a distributed setting. Earlier we discussed how important it is for the collaborators to establish a shared understanding. Similarly, users need to have a common understanding of a representation. During the experiment, we observed one user pointing out a location and having his partner match up his cursor to be sure they were talking about the same location. Other times, users would align their viewpoints to establish a shared understanding, even if they were working with continuous navigation. This demonstrates how essential a common understanding is for collaboration with a representation.

This experiment also revealed how providing multiple representations is not necessarily a good solution. For instance, using the radar map and the individual viewpoint was not as straightforward as we had hoped. We observed users forgetting about the awareness information available in the radar map and relying on a discussion with their partner to determine information. Users also commented that they did not notice or use the view to the left. Switching between the two maps was problematic for our users and further work should investigate whether this approach is useful in other contexts.

4.5 Conclusions

This research explored the concept of distributed, spatial collaboration and investigated different techniques for supporting this style of collaboration. In particular, our experiment compared two ways to navigate a representation (continuous and discrete) and two ways to use radar views to provide awareness information (traditional and fisheye views). The results from this empirical study reveal that discrete navigation is difficult to use when the windows into the representation do not overlap. This lack of overlap causes people to spatially comprehend the navigation movements in comparison to continuous navigation, which is more intuitive. As a result, the discrete navigation tasks took longer to complete and the groups would work together in order to navigate the space. Continuous navigation, on the other hand, supported the spatial collaboration activity, as users did not struggle with the navigation allowing each participant to explore the space individually. There were particular problems with the discrete-fisheye combinations in the study. This condition further limited the users' understanding of the space, as they could not rely on the radar view for a visualization of the entire space. Also, the advantages of the fisheye technique are not yet clear. Some results indicated its usefulness, while others suggested a redesign.

The contributions of this work include empirical evidence of the advantages and disadvantages of different navigation techniques and different radar view techniques. Using both qualitative and quantitative measures, we present the effects of each technique examined and initiate an exploration of alternatives. This work also offers a novel awareness technique. Radar views are enhanced with a fisheye approach to reduce issues with changes in scale and level of detail. This

design addresses important aspects of navigation with a complex spatial representation and it encourages the development of other design ideas. Lastly, this research presents a unique collaboration task and variety of metrics to evaluate spatial collaboration. The task of combining criteria encouraged both users to participate and enabled the pair to solve both open-ended and more precise problems. The strategies of use observed, the analysis of confirmation actions, and the questionnaires provided valuable information about the collaboration process and the users' satisfaction. Future spatial collaboration studies can reuse this task and these metrics to investigate other design solutions.

This work looks at how people collaborate on spatial problems in a distributed setting. It offers some lessons learned about this collaboration context and encourages further exploration. We will continue to explore novel techniques that support distributed, spatial collaboration as well as investigate the issues it presents. For example, we wish to look at other designs for discrete navigation as well as the fisheye radar view. We are also interested in exploring the ways people spatially collaborate. This study focused on very specific task of combining criteria collaboratively and marking map locations. Future work will explore the variety of tasks people engage in and possibly offer a classification of these. Such an understanding will help guide the design of spatial collaboration solutions.

Chapter 5

5 Three-Dimensional Virtual Environment Study

The objective of this phase is to investigate how frames of reference affect spatial collaboration. Frames of reference have been explored in previous aviation research. Investigating different frames of reference with respect to flight navigation, the displays differed in their viewpoints on the space (Wickens & Preveet, 1995; Wickens, 1999). Similarly, when interacting with a virtual environment, the user's view of the environment depends upon his/her frame of reference. One might be situated within the spatial representation so that the user has a first-person view of the environment. This is an egocentric frame of reference. Alternatively, one could look at the space from an external, or exocentric, frame of reference. The user might be located above looking down on the space or below and looking up at the space or have some other angled, external view of the space.

Many virtual environments permit the user to navigate or “fly” into and out of the space, thereby blurring the distinction between strict exocentric and egocentric frames of reference. This navigation provides multiple perspectives on the space and creates a flexible environment for single user applications. Our research, on the other hand, constrains the user's navigation to either an exocentric or egocentric frame of reference. This allows us to examine the direct effects of frame of reference.

In the following sections, we describe a study of similar and different frames of reference and their effects on a task involving spatial collaboration (see Appendix E for materials). We review the related research and present the details of experiment conducted, including the CVE implemented and the specific collaboration task. Participants work in pairs to collaborate on an object manipulation task. We observe their collaboration activities and the participant's awareness of each other's work effort. The results are followed by a discussion of the findings and the broader implications for supporting distributed, spatial collaboration.

5.1 Spatial Collaboration Experiment

In collaboration, people are typically engaged in both individual and shared efforts (Gutwin & Greenberg, 1998; Dourish & Bellotti, 1992). A participant's thought process often involves expressing ideas, as well as understanding the ideas of the other collaborators. One approach to support this mixed-focus collaboration is to allow each user to navigate the same three-dimensional representation individually, while providing awareness information about each collaborator's location and actions. This independent navigation is termed relaxed what you see is what I see (Stefik, Bobrow, Foster, Lanning & Tatar, 1987). Our CVE uses this technique. Avatars provide a visual indication for each user's location and orientation. Objects within the representation are shared so that when an object is moved, all participants witness this manipulation from their own display.

5.1.1 Variables of Interest

Our investigation focuses on frame of reference combinations and their effects on both the overall collaboration session and the participant's awareness of each other's location and activities. The study we conducted examined two frames of reference, and had the participants play two distinct roles. Both of these roles can have an effect on any given spatial collaboration task.

5.1.1.1 Exocentric-Egocentric Combinations

In examining frames of reference, our experiment investigated one technique that provided an exocentric perspective, and another technique that used an egocentric perspective. These two frames of reference enable an exploration of the extreme effects that frames of reference combinations can have on spatial collaboration. The two perspectives are also interesting with respect to the possible combinations. Given the presence of two collaborators, participants could both have an egocentric frame of reference, an exocentric frame of reference, or a combination of an egocentric and exocentric view.

The effects of different egocentric and exocentric combinations on a spatial collaboration task can vary greatly. Using the same frame of reference could be advantageous because both collaborators will have a similar view of the representation, thus establishing a common reference system for the task. The collaborators might approach the task together and take turns interacting with the environment. However, if the collaborators use different frames of reference, this could provide more insight into the task. Each collaborator would have a unique view, possibly making him/her a purveyor of information, or a specialist for certain interactions within the environment. For instance, the participant with the egocentric viewpoint might contribute precise object positioning, knowledge of small details, or a first-person perspective of the space. The user with the exocentric perspective might add large-scale object movements or knowledge of the spatial layout. Differing frames of reference might also enable a group to solve complex spatial problems more efficiently because the unique viewpoints provide multiple views of the same problem.

Frame of reference combinations can also have an effect on the collaborator's awareness of one other's efforts. Awareness has many dimensions, which include knowing where one's partners are located, what they can see, what they are currently focused on, and what they are doing (Gutwin & Greenberg, 2002). When two collaborators use the same frame of reference technique, their similar views might provide for greater awareness because one collaborator can easily align his view with what his partner sees. A disadvantage of this scenario occurs when a user cannot see the other's avatar, making it difficult to discuss the environment together, or to coordinate their individual views of the space. A group of users may have problems knowing where everyone is located and what actions each user is performing.

Given different frames of reference, more situations may exist where the collaborators can see their partners' location and orientation. For example, a user with the exocentric view easily will be able to see the avatar of person with an egocentric view. At the same time, a user with the egocentric view can always know his exocentric partner has an external view of the

representation. However, the frames of reference might be such that it is difficult to discern what features of the space each partner can see, or what actions each is performing.

Our study was designed to explore the differences in exocentric-egocentric combinations. We investigated the effects of each combination on a spatial collaboration task, and, in particular, examined the collaborator's awareness of his/her partner's location and activities within the desktop CVE.

5.1.1.2 Frame of Reference Techniques

The egocentric and exocentric techniques constrain the user's navigation to within the corresponding frame of reference. This ensures that the frame of reference assignments are constant throughout the experiment. The egocentric frame of reference restricted the user to a first-person view of the space. One could navigate forward and backward, and in essence take side steps to the left and right. These movements occur within the same plane, so that the user always remains at the same height. There is no collision detection so that users can walk through the objects that are at the same height or taller. One can also look around by changing his/her orientation. The user can look left, right, up, and down. The interactions support a 360-degree view of space as one can look left/right and turn all the way around. Looking up and down, a user can also change their viewpoint a complete 360 degrees, but this is less useful as the environment appears upside down if he/she rotates beyond looking directly up or directly down.

This egocentric navigation technique provides a fairly simple and usable egocentric frame of reference. A user with this perspective can easily see the objects in the environment and their positions, including height, rotation, and distance from other objects. This frame of reference also enables a user to see an object or room from different viewpoints. For instance, one can repeatedly navigate to view the same room from each of its entrances. Similarly, with multiple egocentric users, each participant can contribute to the collaboration from a different viewpoint. One person might be to the left of the object of interest, while another person looks on from behind. Located in the same area, the group members have a common set of objects to reference as they work together. On the other hand, the participants can also navigate so that they are positioned in different areas of the environment without any objects in common. For example, the collaborators could be located on opposite sides of the space. This can cause problems in the collaboration as the collaborators may struggle in understanding one another. They will most likely be less aware of one another's locations and actions because of the differing viewpoints. The group may even encounter issues as they try and find one another. In not knowing where a partner is located, one collaborator could navigate in the wrong direction, placing him/her further away from the partner.

An egocentric user can also look up, above the environment, to see an exocentric user's avatar. Such an avatar typically is typically not seen from the egocentric frame of reference, but with orientating upward it is viewable. This allows an egocentric user to be aware of an exocentric user, but it does not provide a constant indication of the exocentric partner's presence, or detailed information about the collaborator's viewpoint.

The exocentric frame of reference, on the other hand, is restricted to an external view of the space. In our implementation, the user navigates along the surface of an imaginary half sphere

that surrounds the top half of the space. The interactions support both circling around the environment to view the space from different sides, and changing the height of the viewpoint angle by moving up and down along the half sphere. At the highest position, one sees a top-down view of the environment, and circling around simply rotates the environment. Using this technique, the user's viewpoint is always oriented toward the bottom, center of the representation and the viewpoints are arranged so that the user can see the entire space (Figure 5.1). Therefore, an exocentric user has a more global perspective of the space.




Figure 5.1 Exocentric viewpoints

The exocentric frame of reference enables a user to navigate a half sphere surrounding the representation, where one is always oriented toward the bottom, and center of the representation.

There is no support for zooming to limit a user from seeing a more egocentric perspective. Occlusion also prevents a user from seeing every object in the environment from every viewpoint, as objects in the environment block one's view of the other features in the environment. Coupling this interaction technique with a complex task ensures that the exocentric users cannot simply use a top-down view. For instance, using a task that involves the height of objects, they must navigate the various viewpoints to see the differences.

This exocentric frame of reference technique provides a user with view of the entire representation. One cannot necessarily see the details of the environment, but rather can see the spatial layout, the distances and directions between areas, and the ideal paths through the space. As in the egocentric technique, a user can navigate through a range of viewpoints and view the environment from different perspectives. Looking from the top down provides a view of all the different areas in the environment. Looking from an angle can be beneficial in comprehending the horizontal aspects of a representation, such as elevation changes, or the details on a building wall. As in the example with multiple egocentric viewpoints, multiple exocentric users can also collaborate from different viewpoints and take advantage of the similar objects in their views. However, the participants can also navigate so that their viewpoints have very few objects in common. For instance, two users could navigate down different sides of the half sphere so that occlusion restricts what each can see. As with the egocentric users, this can cause problems in the collaboration with respect to awareness and how well they can understand and communicate with one another.

With an external viewpoint, an exocentric user can easily look inside the representation and see an egocentric user's avatar. Such an avatar cannot be seen from every viewpoint due to occlusion, but it is always visible from the top-down view. This allows an exocentric user to stay aware of an egocentric user. The avatar clearly indicates the egocentric user's location and, if it has facial features, it designates the user's orientation. The face on an avatar indicates which direction the user looking. This information can be used in understanding a collaborator's view

of the environment, including what he/she can see and what one is looking at directly. It also can provide awareness of what actions the collaborator is performing, such as navigating towards an object, moving an object to his/her left.

5.1.1.3 User Roles to Encourage Awareness

Awareness is a critical issue for collaboration, and an important aspect to explore in a distributed, spatial collaboration setting. In order to create an evaluation task that required the participants to be aware of one another within the CVE we used roles. Each participant was assigned a distinct role for the task. The two roles employed were a director and an actor, similar to the work conducted by Gutwin and Greenberg (1999). The director knew the specifics of the task and his/her responsibility was to see that the task was completed correctly. The actor, on the other hand, was charged with performing the actions required.

In this way, the directors needed to be aware of where their partners were located and what they could see in order to provide the instructions. The director also needed to know what the actor was doing to ensure proper task completion. Likewise, the actor often needed to be aware of his/her partner because the director frequently provided directions relative to his location and orientation. For example, a helpful director might provide instructional phrases such as “in front of me”, “follow me”, or “it needs to go to my left”. All of these statements demand an understanding of the director’s perspective. A director could also use statements such as “beside you”, or “near the fireplace”, but these required less awareness of the director from the actor.

Given these roles, our CVE implementation enabled both participants to navigate the spatial representation, but only the actor could perform the actions. This encouraged collaboration and participation as the task responsibilities were divided between both users. Furthermore, the tasks were designed so that one user could not complete the task alone.

5.1.2 Task

In collaboration, one participant often has information that he/she wants to express to the other participants. As the group listens to this idea, members will often confirm their understanding by rephrasing the idea, or performing some relevant action. This experiment recreates this situation for a spatial collaboration activity. One user has a spatial idea about an object’s location, which is provided by the researcher, and the objective is to convey this thought to another user. The person with the spatial idea is the director, and the person listening and confirming the idea is the actor.

Our task has the users collaborate within a three-dimensional space and discuss a new location for an object. Three-dimensional environments are ideal for investigating spatial layouts. They enable the participants to move around and view spatial designs from many perspectives (Pekkola, 2002). In our task, the setting is an interior design of a house, which contains familiar elements such as doorways, windows, furniture, and appliances. The director is provided with a new position for a specific furniture object, and he/she must convey this idea to the actor. The actor is required to reposition the object to demonstrate comprehension. Working together, they discover the environment and the specifics of the task (Figure 5.2 and Figure 5.3).

The director-actor pairs completed four variations of this activity in four similar, but unique virtual environments. The specifics of each task were fairly complex. These included hanging a picture frame on a wall, placing a basketball in between two shelves, moving a chair behind a bookshelf, and positioning a table next to another to form one long table. Each of these required the actor to move the object in the x, y, and z dimensions, as well as rotate the object around the y-axis. This level of difficulty allows a more accurate exploration of distributed, spatial collaboration.

The tasks and corresponding environments were designed so that they were similar in complexity. The structures of the houses were created using mathematical reflections, yet the rooms were different in each environment. Different furniture objects were also used to create unique living rooms, bedrooms, offices, dining rooms, and bathrooms. Users were also required to navigate and adjust their orientation. A user with an egocentric frame of reference started the task in a room different from the one with the furniture object that needed to be repositioned. Also, the task involved moving the object to an opposite corner of the house. The distance between an egocentric user's starting position and the object's initial position, and the distance between the starting and ending position of the object were similar in each environment. The task was also complicated in that the environment contained two similar furniture objects in the initial room of the object to manipulate. This required the director to be specific in giving instructions as both objects were located near each other, and both could be selected and moved by the actor.




Figure 5.2 Egocentric, actor condition

An egocentric actor has selected the orange basketball, which can be moved with buttons in the pop-up window. The floating head avatar represents the egocentric director's position and orientation.




Figure 5.3 Exocentric, actor condition

An exocentric actor and an egocentric director work together to position a picture frame centered above the bed and flush against the wall.

5.1.3 User Interface

The users interacted with the environment using a desktop computer. This technology is prominent in homes, offices, and some public places. It is also currently used for everyday collaboration through the use of email, instant messaging, and online forums. Similarly, it is likely to be utilized for many spatial collaboration activities in comparison with more expensive and more immersive displays. For instance, people currently use desktops to email online map links.

All of the interactions were performed using a standard mouse and keyboard. In the egocentric technique, the up/down arrow keys would move the user forward and backward according to his current orientation, while the left/right arrow keys corresponded to the perpendicular side steps. Movements were constrained to an x-z plane that was six feet tall, measured from the house floors. Clicking and dragging on the window enabled the users to change their orientation. Dragging the mouse to the left rotated the viewpoint around the y-axis, causing the user to see more of the environment to the left. Dragging it upward enabled the user to look down, and dragging it diagonally would affect both orientation dimensions.

The interactions for the exocentric technique mostly involved a keyboard. The left/right arrow keys allowed a user to circle around the environment, while the up/down arrows changed the height of the viewpoint angle. Navigating through the top of the sphere, the user's orientation was adjusted to keep an upright viewpoint as opposed to viewing the space upside down.

Selection and manipulation interactions were the same for both the exocentric and egocentric techniques. Right-clicking over an object selected it, and this was indicated by a red-lined bounding box. A popup window was displayed with buttons to move a selected object forward, backward, left, right, up, and down, and rotate the object left and right around the y-axis (Figure 5.2 and Figure 5.3). Both forward/backward and left/right movements were based on the user's viewing angle. Selecting forward/backward buttons moved the object in and out of the screen, and left/right moved the object across the screen. There was no collision detection such that the object could pass through all the other objects in the environment, including the floors and walls. When the object was positioned within a predefined level of precision, it changed to a bright yellow color to indicate task completion.

5.1.4 Software Design

The software used in the experiment was written using the C programming language. The Simple Virtual Environment Toolkit, or SVE, was used to generate the virtual environment programs and interactions (Kessler, Bowman, & Hodges, 2000). This code interfaced with the Content Object Replication Kit, or CORK, which provides data sharing and is written in Java (Isenhour, Rosson, & Carroll, 2001).

5.1.5 Experimental Design

Thirty-two people participated in the experiment, including computer scientists, biologists, human factors experts, and ecologists. Of this group, sixteen males and sixteen females completed the task with five male-male pairs, five female-female pairs, and six male-female pairs to balance the effects due to gender. The average age of the participants was twenty-five, the youngest person was nineteen and the oldest fifty-four years old. Many participants did not know their partner prior to the experiment and others were simply acquaintances. Almost all of the participants had experienced a virtual environment before. Many had viewed a demonstration of an immersive virtual environment with the CAVE, or a Head-Mounted Display, and six of the participants played three-dimensional games on a regular basis.

Each participant followed a similar two-session procedure including: a practice and an evaluation. During the individual practice sessions, participants read the instructions and completed a set of practice trials. This allowed them to learn the interactions required to navigate their viewpoint, select an object, and manipulate the object selected. All participants were exposed to the same three trials, which involved being an exocentric and egocentric actor, as well as an egocentric director. If the participants struggled with one of the tasks, or did not feel comfortable with one of the perspectives, a second practice trial was completed with the particular condition. The same researcher trained each of the participants and contributed to the practice tasks as a partner would in an evaluation session. This practice allowed the users to experience the collaborative virtual environment before completing the timed tasks two to four days later.

In the evaluation session, the two participants were randomly paired. After being introduced and reminded how to interact with the software, the pairs completed four unique tasks together. They were seated back to back, so as not to be able to see one another's face, or each other's screen. They were asked to not turn around, but to talk freely about the task. Before collaborating with each interior house design, the pair was notified of their different roles and

perspectives. Then, after completing the furniture manipulation task, the users individually rated the treatment condition. Informal debriefing discussions were conducted with the pairs at the conclusion of all tasks before the group left.

The experiment used a 2 x 2 factorial design in blocks such that each subject played the actor and the director for two tasks. Each group worked with the four possible assignments of roles and frames of reference (Table 5.1). The ordering of these combinations and the assignment of the users to the roles varied with each pair. The four furniture manipulation tasks were also randomly assigned to each condition to limit any confounding effects from the tasks not being similar.

Table 5.1 Frame of reference, role conditions

		Actor	
		Egocentric	Exocentric
Director	Egocentric		
	Exocentric		

Both frames of reference were assigned to the actor and director roles, which corresponded to the four tasks completed by a pair of participants.

5.2 Experimental Results

Three different metrics were collected in order to compare the different frame of reference combinations. At the end of each task, participants completed a questionnaire that inquired about their collaboration during the task, the effort involved in completing the task, and their awareness of their partners during the task. The survey also had an extra question that specifically asked about the partner’s location while completing the task. This provided another measure of awareness that did not involve subjective option. Lastly, the software automatically collected timing data to reveal how the pairs of participants spent their time. Tables with mean values are provided for each interaction.

5.2.1 Questionnaire Results

The questionnaire completed at the end of each task contained three sections: Collaboration Rating, Perceived Effort, and Awareness Rating. Each section consisted of three or four questions, and answers were based on a 7-point, Likert-like scale. The Analysis of Variance results for each section follow.

5.2.1.1 Collaboration rating

The collaboration rating consisted of three statements shown in Table 5.2. Participants rated their level of agreement with each statement. A low rating corresponds to “strongly agreed” and a high rating disagreed.

Table 5.2 Collaboration rating statements

1. I feel the collaboration session overall went great. We had no problems and did not struggle to complete the task.
2. I could easily understand my partner’s ideas.
3. I could easily communicate my ideas.

Participants agreed with statement 1 more when the actor had an egocentric view ($F = 2.81, p = 0.0973, df = 1$). The mean value for the egocentric frame of reference was 2.25, in comparison with an exocentric view, with a mean of 2.72. Users also agreed with this statement more when the director had an egocentric view ($F = 2.99, p = 0.0872, df = 1$), with the means being 2.25 and 2.73, respectively.

There was a trend for an interaction for statement 2 with the director and actor view assignments ($F = 3.33, p = 0.0713, df = 1$). Table 5.3 below indicates the mean values for the conditions. Comparing these means using Least Significant Differences, the egocentric-egocentric combination is significantly less than the others. In other words, subjects understood one another more clearly when both partners had egocentric views.

Table 5.3 Mean ratings for collaboration statement 2.

		Actor View	
		Egocentric	Exocentric
Director View	Egocentric	1.90 * (s.d. = 0.98)	2.69 (s.d. = 1.62)
	Exocentric	2.46 (s.d. = 1.46)	2.22 (s.d. = 1.07)

In response to the statement “I could easily understand my partner’s ideas”, the egocentric-egocentric combination results are significantly different.

5.2.1.2 Perceived effort rating

The perceived effort section also included three statements where a lower rating corresponded to agreement and higher numbers meant disagreement (Table 5.4).

Table 5.4 Perceived effort statements

- | |
|--|
| <ol style="list-style-type: none"> 1. The task was easy to complete. 2. The task required little effort. 3. I did not have to concentrate very hard to do the task. |
|--|

Averaged together for an overall perceived effort, there was an interaction trend in the participants’ role and the director view assignments ($F = 2.92, p = 0.0913, df = 1$). Table 5.5 shows how the actor’s response with the exocentric director corresponded to significantly greater perceived effort than the other combinations.

Table 5.5 Mean ratings for the perceived effort statements

		User Role	
		Actor	Director
Director View	Egocentric	2.29 (s.d. = 1.11)	2.50 (s.d. = 1.28)
	Exocentric	2.90 * (s.d. = 1.53)	2.43 (s.d. = 1.23)

Considering all three perceived effort statements, actors responded significantly different when collaborating with an exocentric director.

Analyzing the individual perceived effort statements, there was a significant interaction for statement 1 with the frame of reference assignments ($F = 4.08$, $p = 0.0466$, $df = 1$). Subjects felt the task was more difficult to complete when both users had exocentric views. Table 5.6 provides the mean values for each condition.

Table 5.6 Mean ratings for perceived effort statement 1.

		Actor View	
		Egocentric	Exocentric
Director View	Egocentric	2.50 (s.d. = 1.32)	2.12 (s.d. = 1.21)
	Exocentric	2.22 (s.d. = 0.97)	2.82 * (s.d. = 1.77)

Responses to the statement “The task was easy to complete” were significantly higher for the exocentric pair combination.

Participants also responded differently to the second perceived effort statement “The task required little effort”. The participants felt the task required less effort with an egocentric actor ($F = 2.92$, $p = 0.0910$, $df = 1$). The mean for the egocentric view was 2.47, and with an exocentric view it was 2.95.

In the third statement, there was a significant difference between the actor’s response with an exocentric view and the director’s response with an exocentric actor ($F = 3.55$, $p = 0.0628$, $df = 1$). Users felt the task required greater concentration when they played the role of an exocentric actor, and less concentration as the director with an exocentric actor. Table 5.7 provides the mean results for each condition.

Table 5.7 Mean ratings for perceived effort statement 3.

		Actor View	
		Egocentric	Exocentric
User Role	Actor	2.58 (s.d. = 1.26)	3.70 * (s.d. = 1.47)
	Director	2.73 (s.d. = 1.22)	2.05 * (s.d. = 1.55)

Exocentric actor responses to the statement “I did not have to concentrate very hard to do the task” were significantly different from director responses with an exocentric actor.

5.2.1.3 Awareness rating

The awareness section of the questionnaire included the four questions listed in Table 5.8. Participants responded to each question on a 7-point numerical scale, where 1 corresponded to “always” and 7 was labeled as “never”.

Averaging the participant’s responses for the four questions, there was a significant effect of the director’s view assignment ($F = 13.03$, $p = 0.0005$). Participants indicated that they had a greater awareness of their partner’s location, viewpoint, and actions in the egocentric director combinations. The egocentric mean was 3.08 and the exocentric mean was 4.44.

Table 5.8 Awareness questions

1. How often did you know where your partner was located?
2. During the trial, how often did you know what your partner could see?
3. During the trial, how often did you know what your partner was directly looking at?
4. During the trial, how often did you know what your partner was doing? (Standing still, navigating/moving, moving an object, etc.)

Examining the individual awareness questions, there was an interaction trend for the first question (“How often did you know where your partner was located?”). Participants responded different depending on their role and the director’s assignment ($F = 2.81, p = 0.097, df = 1$). Comparing the means listed in Table 5.9 using Least Significant Differences, directors with an exocentric frame of reference indicated having the least awareness, followed by actors who worked with an exocentric director. The greatest amount of awareness corresponded to both egocentric directors and actors who worked with egocentric directors.

Table 5.9 Mean ratings for awareness question 1

		User Role	
		Actor	Director
Director View	Egocentric	2.89 * (s.d. = 1.88)	2.72 * (s.d. = 1.90)
	Exocentric	3.98 (s.d.= 2.55)	5.17 (s.d. = 2.33)

With an egocentric director, both participants indicated an equal level of awareness, which was greater than the other conditions.

Participants also indicated awareness differences for the second question (“How often did you know what your partner could see?”). Examining the effects of the actor’s frame of reference, pairs with an egocentric actor indicated greater awareness with a mean of 3.42 ($F = 2.85, p = 0.095, df = 1$). With an exocentric actor, the average response was 4.02. Similarly, there is significant interaction for the second question as participants responded differently based on their role and the actor’s assignment ($F = 5.25, p = 0.0245, df = 1$). Comparing the mean values listed in Table 5.10, egocentric actors indicated that they frequently were aware of their partners’ viewpoint, while exocentric actors were the least aware. The directors responded similarly with either type of actor and their awareness level was in between that of the egocentric and exocentric actors.

Table 5.10 Mean ratings for awareness question 2

		User Role	
		Actor	Director
Actor View	Egocentric	3.03 (s.d. = 1.45)	3.81 * (s.d. = 1.68)
	Exocentric	4.31 (s.d. = 1.95)	3.74 * (s.d. = 2.11)

Directors indicated an equal level of awareness with exocentric and egocentric actors. This was significantly less than that expressed by an egocentric actor and significantly more than that indicated by an exocentric actor.

Lastly, there is an interaction trend for the third awareness question asking about a partner's specific gaze. Participants again responded differently based on their role and the actor's assignment ($F = 3.46$, $p = 0.066$, $df = 1$). Egocentric actors indicated the greatest level of awareness and their responses were significantly different from those of exocentric actors and directors working with egocentric actors. **Error! Reference source not found.** provides the mean values for each condition.

Table 5.11 Mean ratings for awareness question 3

		User Role	
		Actor	Director
Actor View	Egocentric	3.77 (s.d. = 1.62)	4.57 (s.d. = 1.96)
	Exocentric	4.46 (s.d. = 2.04)	4.06 (s.d. = 2.16)

Egocentric actors were more aware than exocentric actors and directors with egocentric actors.

5.2.2 Awareness Question Results

One additional question on the survey directly measured the participants' level of awareness. The question asked them to recall where their partner was located during the task. To prevent the users from learning this question, four variations were used in the four tasks. The versions differed in the time frame of the question and in the reference points provided. For instance, one question asked about which wall the partner was closest to when the object was selected and another asked about which object the partner was closest to at the end of the task. The question variations were balanced with the treatment conditions so that each question was asked four times for each condition. The ordering of the questions was also counter-balanced across the groups.

Participants did not respond to the question when their partner was exocentric, resulting in a total of sixty-four responses. Twelve of these were correct; two of the participants provided the correct answer twice. Examining the conditions used in providing the correct answers, there were no significant differences between the perspective combinations or in the users' role assignments.

5.2.3 Timing Results

The software automatically logged timing data for each task. The timing started when one user clicked the large "Start" button covering both displays and ended when the correct object was

positioned within the level of precision and unselected by the actor. Participants were asked to perform the tasks as quickly as possible.

5.2.3.1 Time to Select

Examining the time required for the participants to select the correct object, the egocentric pair required significantly more time than the exocentric-egocentric combinations ($F = 4.06$, $p = 0.054$, $df = 1$). The egocentric pairs also took more time than the exocentric pairs, but the difference was not statistically significant. Table 5.12 indicates the mean times for each condition.

Table 5.12 Mean times for selection

		Actor View	
		Egocentric	Exocentric
Director View	Egocentric	99.65 sec (s.d. = 69.67)	39.08 sec (s.d. = 22.23)
	Exocentric	51.26 sec (s.d. = 42.23)	70.01 sec (s.d. = 53.14)

The egocentric pair combination required significantly more time to select the correct object than both the exocentric-egocentric combinations.

5.2.3.2 Selection and Coarse Positioning Time

There were also differences when combining the time required to select the object with the time used in positioning it within a coarse level of precision. Egocentric actors required significantly more time for the selection and manipulation than the exocentric actors ($F = 8.58$, $p = 0.007$, $df = 1$). The mean time for an egocentric actor was 159.12 seconds, and with an exocentric actor it was 105.83 seconds.

In conducting the study, each participant played the role of the director twice. If we assume that there was not a user effect with regards to which participant played this role in the four tasks, there was an additional significant difference in the times listed in Table 5.13. When both participants had an egocentric frame of reference, the pairs required significantly more time to select and coarse position the object.

Table 5.13 Mean times for selection and coarse positioning

		Actor View	
		Egocentric	Exocentric
Director View	Egocentric	188.19 sec * (s.d. = 121.59)	98.25 sec (s.d. = 52.04)
	Exocentric	130.06 sec (s.d. = 68.57)	113.42 sec (s.d. = 78.03)

Egocentric pairs required significantly more time to select the correct object and position it within a coarse grain level of precision.

5.2.3.3 Coarse Positioning Time

Analyzing only the time required to position an object with a coarse level of precision also differed significantly. Again, egocentric actors required much more time than their exocentric

counterparts ($F = 8.07$, $p = 0.008$, $df = 1$). The mean time for an egocentric actor was 129.79 seconds, while for an exocentric actor it was 82.33 seconds.

Assuming that there were no user effects due to the director assignment, the egocentric pair required significantly more time than both the exocentric pair and the exocentric actor-egocentric director conditions. The egocentric pairs also took longer than the egocentric actor-exocentric director for this part of the task, but this difference was not significant. The mean times for each combination are listed in Table 5.14.

Table 5.14 Mean times for coarse positioning.

		Actor View	
		Egocentric	Exocentric
Director View	Egocentric	147.63 sec (s.d. = 76.56)	83.57 sec (s.d. = 47.42)
	Exocentric	111.96 sec (s.d. = 51.23)	81.09 sec (s.d. = 55.86)

The egocentric pair combination required significantly more time than both the combinations with an exocentric actor to position the object within a coarse grain level of precision.

5.2.3.4 Fine Positioning Time

There was a significant difference in the time required to finely position the object after it had been located within a coarse level of precision. Egocentric actors required significantly less time than their exocentric counterparts ($F = 3.17$, $p = 0.086$, $df = 1$). The mean time for an egocentric actor was 64.03 seconds, and that of an exocentric actor it was 93.75 seconds.

Looking at the mean times and assuming no effect of director assignment, the exocentric actor-egocentric director combination required significantly more time than the other conditions. Table 5.15 lists the mean times for each condition.

Table 5.15 Mean times for fine positioning

		Actor View	
		Egocentric	Exocentric
Director View	Egocentric	59.43 sec (s.d. = 45.17)	123.35 * sec (s.d. = 117.86)
	Exocentric	66.63 sec (s.d. = 50.72)	64.15 sec (s.d. = 69.70)

The exocentric actor-egocentric director combination required significantly more time to finely position the object.

5.2.3.5 Positioning Time

Assuming no effect of director assignment and analyzing the time used to finely position the object after selection, the exocentric pair required significantly less than time than both the exocentric actor-egocentric director and egocentric pair conditions. The exocentric pairs also took less time than the egocentric actor-exocentric director, but this difference was not significant.

Table 5.16 lists the mean times for each condition.

Table 5.16 Mean times for positioning

		Actor View	
		Egocentric	Exocentric
Director View	Egocentric	207.06 sec (s.d. = 76.21)	206.93 sec (s.d. = 143.88)
	Exocentric	178.59 sec (s.d. = 83.23)	145.24 sec (s.d. = 103.14)

The exocentric pairs required the least amount of time for positioning the object, including coarse and fine movements. This was significantly different from the egocentric director combinations.

5.2.3.6 Total Time

Overall, the egocentric pair condition required significantly more time than both the exocentric pair and the egocentric actor-exocentric director conditions, assuming that there was no effect due to director assignment. The egocentric pair also took more time than the exocentric actor-egocentric director, but this difference was not significant. Table 5.17 lists the mean times for each condition.

Table 5.17 Mean times for task completion

		Actor View	
		Egocentric	Exocentric
Director View	Egocentric	306.70 sec (s.d. = 116.01)	246.01 sec (s.d. = 152.77)
	Exocentric	229.85 sec (s.d. = 111.48)	215.25 sec (s.d. = 126.05)

The egocentric pairs required the most amount of time to complete the task, significantly more than the pairs with an exocentric director.

5.3 Discussion

These results provide insight into how frames of reference affect spatial collaboration. They reveal the advantages and disadvantages of combining the actor and director roles with different viewpoints. They also explore how the frame of reference combinations affected the collaboration. A discussion of the results follows.

5.3.1 Actor and Director Frames of Reference

When analyzing these results it is interesting to consider how the participants' responses on the questionnaire agree and disagree with the timing metrics. One example where the users expressed problems that corresponded to increased time involved an exocentric actor. Participants felt that the collaboration session proceeded more smoothly and with fewer problems when the actor had an egocentric view. They also responded that the task required more effort and greater concentration with an exocentric actor. This can be attributed to the actor's ability to see more of the details of the environment with an egocentric frame of reference. When located in the environment, the user could easily recognize the correct object to move, the target location, and the small differences between an object's current and target location.

An exocentric actor, on the other hand, struggled to position the object with the fine grain movements required in the final manipulation stages. Such an actor cannot see the details of the environment as clearly, and it is important to position oneself to achieve an optimal view of the object and the target space. The timing results support this finding. The pairs with an exocentric actor required more time to finely position an object once it was within a coarse grain level of precision. This indicates that the exocentric actor condition faced object manipulation issues in the later parts of the task.

However, the user's belief that the collaboration session went more smoothly and was less problematic with an egocentric actor is not apparent in the timing metrics. In fact, the participants took longer to complete two aspects of the task when the actor used an egocentric frame of reference. Both the time to position the object with a coarse grain level of precision, and the time to select and position within a coarse grain level required significantly more time with an egocentric actor. The participants did not reveal issues with the egocentric actor in their responses, but the selection part of the task demands additional time and effort because the actor has to find the object by exploring the space. The director can help or hinder this process depending on his/her own knowledge of the space.

Also, in manipulating the object an egocentric actor must rely on his spatial knowledge of the environment. This knowledge is unlike an exocentric actor's and most users cannot determine the shortest path to the target. Such a path often involves moving the object through walls and other objects in the environment. It also involves positioning oneself so that the object moves along the straightest trajectory toward the target. This is not easy to accomplish because it requires that the user knows precisely which way to face, where he/she should position himself/herself with respect to the object's location, and whether the object should be moved forward in a pulling motion, or moved backward in pushing movement.

There is also a discrepancy in the questionnaire responses regarding the director's perspective. Participants felt that the collaboration session occurred more smoothly and with fewer problems when the director had an egocentric view. The actors also indicated that the tasks required more effort with an exocentric director. This again could be due to the ability to see more of the details with the egocentric frame of reference. Providing these details to the director may enable the user to be more involved with the task, placing less of a burden on the actor to determine the specifics of environment. Another possibility is that users in general are less familiar with an exocentric view, and so they are less helpful in providing important information to their partners. However, the timing results do not provide evidence for these issues. The times associated with selection, coarse positioning, coarse and fine positioning, and total task completion time indicate that the combinations with an exocentric director are similar to and even faster than their egocentric director counterparts. This implies that the participants may have encountered problems with an exocentric director, but it did not affect their overall performance.

Similarly, the questionnaire results favored an egocentric director with respect to greater awareness. Participants expressed knowing more often about their partners' location, what they could see, what they were directly looking at, and what they were doing with an egocentric director. This may be because the director is in a better position to see the details of the environment that are highly important for the task, and, in turn, are frequently discussed by the

actor. These details involve identifying the correct object to select, confirming the correct target location, and ensuring the object is correctly positioned. As a result, the director maintained a greater sense of the actor and his actions, regardless of his partner's frame of reference.

Also, with the director physically located inside the environment, the instructions one provides can be more specific because the director can reference his avatar's specific location. However, the timing results support an alternative explanation, as the pairs required significantly more time with the exocentric actor-egocentric director combination to finely position an object after a coarse positioning. This reflects the exocentric actor's difficulties in finely tuning the position, as well as the pair's combination to collaborate effectively. The egocentric director knows the task specifics and can see the object in motion, but this user is less helpful in conveying the specifics required by the exocentric actor. This is most likely due to the director's lack of awareness and the actor's impatience. For example, the exocentric actor might see the correct object and move it to the general target area before the egocentric director has navigated to the location of the object, or the target location. These results demonstrate that an egocentric director can be both beneficial and problematic with respect to awareness issues for spatial collaboration.

5.3.2 Frame of Reference Combinations

In addition to the results of the actor and director assignments, there are also some interesting findings related to the three frame of reference combinations. When both participants used an egocentric frame of reference, the users felt that they more clearly understood one another. This confirms that similar egocentric frames of reference have an effect on collaboration where it provides collaborators with a familiar first-person view of the environment and enables them to share ideas with greater ease. However, the timing results describe a different conclusion. The two egocentric views corresponded to the greatest amount of time for selection, selection and coarse positioning, coarse positioning, and total time. This demonstrates how the two egocentric frames of reference are at a disadvantage in comparison with the other combinations. One explanation is that neither participant can easily see the object at the start of the task, or has an understanding of the environment's layout. They must explore the environment to gain this spatial knowledge and find the object. Once the object is selected, the pair next has to find the target location and move the object through the environment using their spatial understanding.

Our task can be considered somewhat biased in that the actor, and more importantly, the director was not familiar with the environment. This caused the participants to initially discover the environment together in order to locate the object, the target location, and a path from one location to the other. In a more realistic spatial collaboration activity, the director would be familiar with his spatial idea before expressing it, and the actor would likely know about the space under discussion. Our task was the same for each of the frame of reference combinations, but this unfamiliarity could have caused the timings to be greater with the egocentric pair condition. Neither user could see the overall layout of the space, which could have contributed to a larger "discovery" time in comparison with the other conditions.

In comparison to the egocentric pairs, both egocentric-exocentric combinations corresponded to an easier task completion rating and the least amount of time for the selection. This agrees with the hypothesis that when working together the users were able to see the fine grain details as well

as the overall layout of the environment, enabling them to solve the task with greater ease. The users also realized and commented on the potential value of the different frames of reference, leading to favorable survey responses due to intuition as well as experience. Conversely, the surveys also indicated that the pairs had the most problems understanding one another with different frames of reference. Using different frames of reference, the users had to deal with issues that were different from having the same frame of reference. Their views into the space were unique in that an exocentric view of an object looked very different from an egocentric view of the same object. This confirms that different frames of reference can have a negative effect on collaboration.

Particularly interesting are the timing results from the egocentric actor and exocentric director combination. The time required for selection was similar to that of the exocentric actor, egocentric director condition and both exocentric frames of reference. This implies that the director was very helpful in guiding the actor to the correct object. The coarse positioning time was also not significantly different from these two combinations, extending the benefits of the combination even further. This shows that an egocentric actor, aided by an exocentric director, can perform large-scale object manipulations similar to that of an exocentric actor. A similar comparison also exists with the total task completion time, suggesting that this combination supported both fine and coarse grain movements.

Lastly, the paired exocentric frames of reference presented unique issues to the collaborators. In the survey, the participants felt that the task was most difficult to complete with this combination. The precise movements required by the task and the need to distinguish between objects in the environment lead to problems. We also observed participants discussing the layout with respect to their viewpoint using references such as bottom, left corner and forgetting that their partner's view was different. It was also more difficult to use the avatars for collaboration in an exocentric frame of reference. If you see an avatar in the egocentric view, spatial references such as "behind you" are possible, and are much more prevalent. Seeing an avatar in the exocentric view more often confirmed that the other person was present and that this partner had an exocentric frame of reference as well. However, these disadvantages were not revealed in the timing results. In fact, the exocentric paired condition required significantly less than time than both the exocentric actor, egocentric director and the egocentric paired conditions with respect to the time required to finely position the object once it was selected. This is probably due to the limited time required by two exocentric users to perform large-scale manipulation of an object across the space. Both can see the object in respect to the spatial layout and can ensure the object moves along the straightest path toward the target location.

5.3.3 Awareness Information

The awareness results from this study were less straightforward. We envisioned that the frame of reference combinations with an exocentric and egocentric viewpoint would correspond to greater awareness. The survey results, however, revealed differences in the frame of reference assignments for the actor and director and presented interactions between the role played by each user and either the director or the actor's assignment. For example, in response to the question "How often did you know what your partner could see?", participants replied differently with an exocentric and egocentric actor. There was also a significant interaction between the actor and the director's responses and whether the actor was egocentric or exocentric. These differences

were not expected, but they provide insight into how the awareness information varied with the different situations created by the study.

One explanation for the question that asked about “what your partner could see” relates to the different combinations. Paris with an egocentric actor indicated that they had an increased level of awareness over pairs with an exocentric actor. This can be explained by the fact that egocentric actors would have had either a similar, or more detailed, view of the environment in comparison to their director. Exocentric actors, on the other hand, would have known less about their partner’s view because their directors either had a detailed egocentric view, or the actors knew little about an exocentric director’s location and viewpoint.

A similar relationship also exists in the significant interaction where the pair results are decomposed into the actor and director’s responses. In this interaction, the director’s responses are not significantly different for the two types of actors, while the actor’s responses are significant. This reflects that the directors, in general, did not differ in their knowledge of their partner’s view based on the different frames of reference used by their partner. Thus, the interaction reiterates the differences between an egocentric and exocentric actor.

Looking at the results for the related awareness question “How often did you know what your partner was directly looking at?”, there is a similar interaction trend. Egocentric actors indicated the greatest level of awareness and their responses were significantly different from those of both exocentric actors and directors who were partnered with an egocentric actor. In this case, the directors with egocentric actors knew the least about their partner’s gaze because the egocentric actors are most likely focused on a small part of the environment, as opposed to exocentric actors who can see the entire space. This narrow focus would be difficult to discern as both an egocentric and exocentric director leading to a different response from a director with an egocentric actor.

Lastly, there was an interaction trend as participants responded differently when asked: “How often did you know where your partner was located?” Directors with an exocentric frame of reference indicated having the least awareness, followed by actors who worked with an exocentric director. The greatest amount of awareness corresponded to both egocentric directors and actors who worked with egocentric directors. One explanation for this result is that actors who work with an exocentric director either have an egocentric viewpoint, or a similar exocentric viewpoint. In either case, the director is potentially less involved in the collaboration because of his outside perspective. Thus, making the actor less knowledgeable of his location. On the other hand, with an egocentric director, the partner may be more involved in the collaboration as he is physically located in the environment, enabling the actor to know more about the director’s location. This does not explain why an exocentric director is the least aware of his partner’s location, however. It seems that such a viewpoint would allow a user to visibly see where an egocentric actor was located, but not know an exocentric actor’s position. Possibly, the inherent issues with two exocentric viewpoints dominated this result.

5.4 Implications

In conducting this experiment, a number of lessons were learned about distributed, spatial collaboration within a collaborative virtual environment. One of the most salient observations

relates to the participants' ability to collaboratively position the object. After the object was moved close to the target location, the groups would often work together to determine the fine grain manipulations. Using their unique viewpoints, the actor and the director would contribute simple movement suggestions. For example, an egocentric director would often encourage the exocentric actor to move the object up or down, while the actor took the steps to align the rotation using a top-down view. The director often preceded such an activity with phrases such as "I'll help you position it", or "I'll go and see if it's close". This coordination was also prominent with the two egocentric pairs. With both participants located in the environment, the users often positioned themselves so that they had complementary viewpoints of the object. For instance, in placing the picture frame centered above the bed and flat against the wall, one user would navigate to the foot of the bed and face the wall, while the other was positioned along the side of the wall to gauge the distance between the frame and the wall. This demonstrates that users are able to take advantage of their independent viewpoints in a virtual environment. In other words, the actors would rely on their partners' viewpoint and their comments as they positioned the object together. They did not complete the task on their own and navigate to see the objects position from multiple viewpoints. This is not a surprising finding, but it is important point for designing distributed, spatial collaboration solutions. Independent viewpoints allow users to collaborate from different perspectives and enable them to solve complex spatial tasks.

During the experiment, we also observed that the different viewpoints present challenges for distributed, spatial collaboration. In particular, users frequently gave directions relative to their own viewpoint. With two egocentric users, the participant who understood the spatial layout better would often lead the partner to a particular area. The pair would, in essence, play "follow the leader", as one user followed his/her partner's avatar through the space. If the pair became separated, the leader would give directions based on his path, or retreat to find the lost partner. In using this approach, the leader participant guided relative to his/her own viewpoint.

Similarly, with an exocentric actor and an egocentric director, the director typically indicated that the object needed to be moved to his/her left or right. Problems occurred when a participant needed to provide a reference independent of his/her location. For instance, exocentric directors frequently were observed instructing with the phrase "move the object towards me". This was not helpful because the actors typically did not know where the director was located. Providing such a reference relative to the environment or to a partner's location did not come naturally to the participants. Trial and error was often used as the director tried to indicate to the actor which direction to move an object. For instance, one director started with "You're going to have to move it to your . . .", but could not easily figure out the direction in order to complete the sentence. He/She finished his comment with "yeah, that way", as the actor moved the object in one direction. This indicates that users favor directions relative to their own viewpoint and struggle with the mental transformations inherent in different frames of reference. Spatial collaboration solutions need to recognize these issues. They should provide techniques so that users can give clear directions based on their viewpoint, while reducing the mental transformations required in understanding the instructions.

Another interesting finding from the observations was that the users preferred to see their partner's avatar. The avatar provided crucial information about their partner's ongoing presence and position. Combinations with an exocentric viewpoint struggled with this lack of a useful

avatar. With two exocentric viewpoints, there was often confusion about where a partner was located. At the start of the tasks, we observed the director asking where the actor was located and inquiring whether they were supposed to be able to see each other. The group continued the task assuming that they could not see one another, although certain viewpoints did include the avatars.

Similarly, with an egocentric director and exocentric actor, we observed one director asking his partner “Are you directly behind me?” and the actor responding with “I’m above you.” This did not answer the director’s question about where the actor was located, but the director continued on simply knowing the actor was present. On another occasion, an exocentric director repeatedly kept asking his partner “Do you see me?”. This director kept hoping that the egocentric partner would be able to follow his/her directions similar to when the user had egocentric viewpoint. This strategy was not successful though as the egocentric actor did not understand the exocentric director’s location. The users with egocentric viewpoints could always look up to learn about their partner’s location, but this rarely occurred. Rather the groups continued on assuming an exocentric user’s presence, but not knowing their location. With two egocentric frames of reference, the avatars played a different role and were particularly helpful as the participants completed the tasks. Frequently, one user would navigate to a prominent place to mark a location in the environment. This allowed one user to explore the space and have an additional landmark within the environment. Similarly, users would often navigate around corners to guide their partner on which direction they should navigate toward. This use of avatars provided the participants with awareness information in addition to acting as navigation aids. Similarly, spatial collaboration solutions with different frames of reference need to specifically consider how the participants will stay aware of one another. Designs should support knowing both who is present and where they are located. The details of the information provided do not need to be specific, but they should reduce the problem of not having a location context for an exocentric user.

Lastly, during the study we observed some usability problems with the exocentric frame of reference implementation. Using this viewpoint, participants would occasionally question the objects in the environment. These users had trouble identifying the furniture pieces in the world and sometimes began to move the wrong object in completing the task. They also were less certain about the rooms of the interior design referenced in the task than other groups. This reflects a need to design exocentric viewpoints so that the objects and the environment areas are distinguishable. Similarly, the avatar locations and orientations should be salient as well as the selected, and moving objects should be obvious. Future spatial collaboration designs should keep these guidelines in mind.

5.5 Conclusions and Future Work

This research investigated the effects of similar and different frames of reference in a distributed and collaborative spatial environment. In particular, our experiment compared three frames of reference combinations (exocentric-exocentric, egocentric-egocentric, and exocentric-egocentric) using two roles (actor and director) to investigate awareness issues.

The study contributes to the design of collaborative virtual environments as it presents advantages and disadvantages in assigning the actor and director of a task specific exocentric and

egocentric frames of reference. An egocentric actor allows a user to experience the environment first hand, by exploring the space, seeing the details, and recognizing fine grain tasks. This frame of reference was less efficient in finding objects within the space and moving them large distances across the space. The exocentric actor was much faster in accomplishing these tasks. With a view of the entire space, an exocentric user could quickly identify objects and manipulate them along the straightest path toward the target. The exocentric actor struggled, where the egocentric was successful, in positioning an object with fine manipulations. In assigning the director to the different frames of reference, an egocentric director was beneficial to the collaboration because the user could see and contribute with the details of the object manipulation task. The usefulness of such a director was limited, however, as an exocentric actor continued to have issues with fine grain manipulation.

Examining the frame of reference combinations, the two egocentric users took longer to complete the task, but felt that they could understand each other the best. The exocentric pairs often required the least amount of time, but the participants felt that the task was most difficult to complete. The egocentric actor and exocentric director combination was particularly promising. This pairing seemed to take advantage of the different frames of reference in performing the task efficiently.

This work looks at how people collaborate on spatial problems in a distributed setting using a collaborative virtual environment. It offers some lessons about this collaboration context and encourages further exploration. In the future, we plan to continue to explore different implementations of the exocentric and egocentric frames of reference to determine if these results are repeatable. We also would also like to explore the collaborators' behaviors when they are not restricted to one frame of reference but rather can navigate the space using both first person and more external viewpoints. Given the same roles, it would be interesting to see how the users approach similar tasks, and how their performance compares with the results reported here. Lastly, the findings from this study encourage an exploration of other spatial collaboration tasks. Tasks that require varied object manipulations and other forms of interaction could yield different results.

Chapter 6

6 Real-World Application and Validation

This work extends our previous research on spatial collaboration. It allows us to validate our findings from the previous studies, as well as contribute new insights for distributed, spatial collaboration (see Appendix F for materials).

6.1 Introduction

In previous studies, we have explored supporting synchronous, spatial collaboration through novel, interactive, shared representations. A formal experiment was conducted with collaborative map software that used radar views to indicate each collaborator's viewpoint (Schafer & Bowman, 2003). A second experiment investigated multiple frames of reference with a collaborative virtual environment (Schafer & Bowman, 2004). Both of the studies were limited to pairs of users in specific tasks and did not explore spatial collaboration in a realistic setting. One objective of this work is to investigate a more natural occurrence of distributed, spatial collaboration. Working with a small, active research group, we observe their interactions with complex, open-ended spatial problems. The discussions occur over multiple sessions and each member has a genuine interest in activity. Through observations, we note their communication patterns, their strategies and approaches, and their dynamic use of subgroups.

Secondly, this work demonstrates that spatial collaboration solutions do not have to be limited to one type of representation, such as the collaborative map, or the three-dimensional collaborative virtual environment. Multiple, different representations of the space can enable the users to explore different aspects of the space. The problems addressed through spatial collaboration can often be complex and involve the consideration of multiple issues. Using multiple representations is one way to investigate these different issues. In particular, this work integrates two-dimensional and three-dimensional representations. Reusing the favorable techniques from the previous studies, we present a novel prototype that allows the collaborators to switch between representations. We observe the research group conduct a complex task with this prototype and consider how the different representations support the various activities of the task. It is interesting to examine the situations where collaborators chose the two-dimensional or three-dimensional representation throughout the process. Additionally, we investigate how the similar and different displays enable the collaborators to solve spatial collaboration problems. Observing their subgroup formation, we note the representation chosen by each participant as they work closely together on subtasks of the large goal.

In the following sections, we describe the design of our multiple, representation prototype. We present the observations from a qualitative study with the research group and discuss the implications for the design of spatial collaboration solutions.

6.2 Preliminary Studies

In the two prior studies, we investigated synchronous, distributed spatial collaboration with two-dimensional and three-dimensional spatial representations. Both examined navigation and awareness techniques, but differed in their approach and experimental design. Here, we build on the findings from these experiments.

In the two-dimensional experiment, a collaborative map was used to study navigation and awareness techniques (Schafer & Bowman, 2003). Specifically, it studied discrete and continuous 2D navigation techniques, and compared traditional and fisheye radar views. The major findings from this empirical study reveal that discrete navigation is difficult to use, but affords aligning viewports and pointing out features. Continuous navigation, on the other hand, is preferable for spatial collaboration, as users do not struggle with the navigation and can explore the space individually. In terms of the radar view techniques, some results indicated the usefulness of the fisheye design, while other results suggested a redesign. The traditional radar view is also advantageous due to its intuitive design.

The second experiment studied the effects of different frames of reference within a CVE (Schafer & Bowman, 2004). It investigated an exocentric, or bird's-eye, frame of reference, and an egocentric, or first person frame of reference. It also considered two specific user roles, a director and an actor. The major findings from this study reveal that there are tradeoffs in assigning the collaborator performing object manipulations to an egocentric or exocentric frame of reference. An egocentric actor supports the user in exploring the space firsthand, seeing the details, and recognizing fine grain tasks, yet this view is less efficient at finding objects and moving them large distances. The results for exocentric actor are just the reverse. There are also advantages and disadvantages in assigning the collaborator directing a task to an egocentric frame of reference. An egocentric director contributes with the details of the object manipulation task, however, such a director in combination with an exocentric actor, still leads to issues with fine grain manipulation.

In terms of the frame of reference combinations, two egocentric users took longer to complete the task, but felt that they could understand each other the best. The exocentric pairs, on the other hand, often required the least amount of time, but they felt that the task was most difficult to complete. The combination with an egocentric actor and exocentric director was particularly promising, as this pairing took advantage of their different frames of reference.

6.3 Prototype

A primary difference between the first and second experiments was the dimensionality of the representation. Whereas the first experiment concentrated on a two-dimensional map, the second looked at a three-dimensional virtual environment. Many spatial collaboration activities are complex and could benefit from both of these representations of a physical space. Two-dimensional representations show the structure of the space and the locations of features with respect to another. Three-dimensional representations, on the other hand, can portray the height of features, elevation changes, shadows, etc. Our prototype combines these two representations to investigate their use in a collaborative setting.

In this study, participants can interact with either a two-dimensional map or a three-dimensional virtual environment. The two-dimensional map uses continuous navigation from the previous study. It supports zooming and panning the representation with numerous viewports into the space. The three-dimensional virtual environment, or CVE portion, uses an egocentric frame of reference. Users experience the space as a normal person would. They walk and look around the environment, seeing the space as if they were 6 feet tall. Each collaborator chooses to work with one representation or the other, and he/she can easily switch between the two interfaces (Figure 6.1, Figure 6.2, and Figure 6.3)

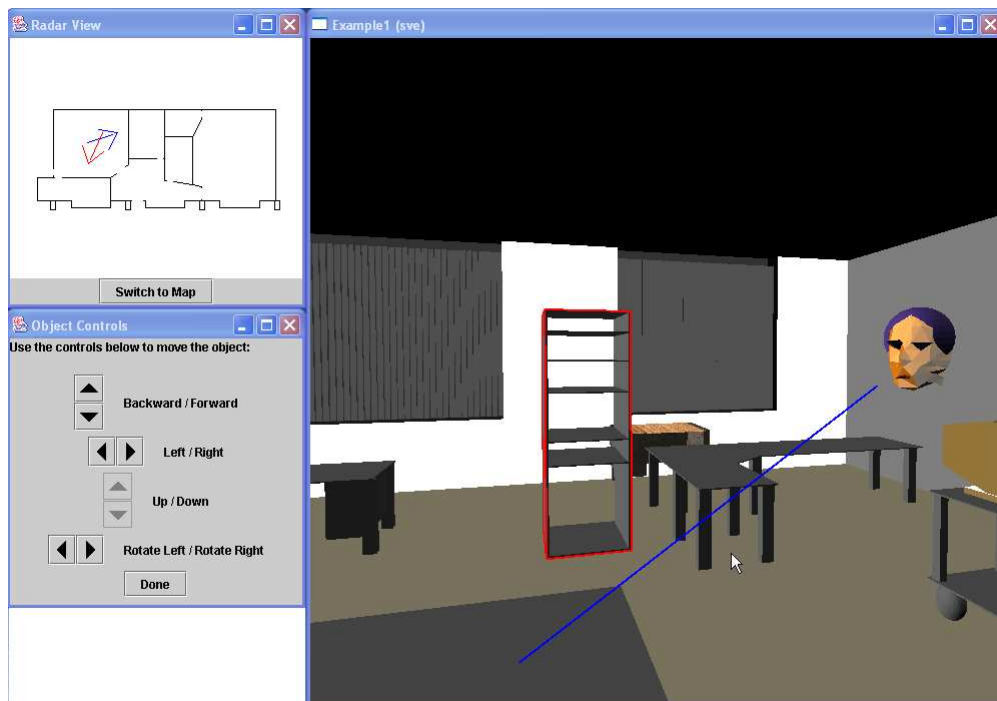


Figure 6.1 Collaboration with the virtual environment representation

A user working in the three-dimensional, collaborative environment has selected the equipment rack outlined in red. This object is movable with the buttons in the window to the left. The floating head avatar represents a collaborator's position and orientation, while the blue line indicates what is being pointing at.

Enabling the collaborators to make this choice in representations allows some participants to view the map, while others look at the CVE. This is an interesting situation in a spatial collaboration activity. The collaborators could use a variety of approaches to coordinate their activities using the unique displays. One possibility is that each collaborator will work with both two-dimensional and three-dimensional representations. Another approach is that some collaborators will primarily interact with the two-dimensional display, while others focus on the problem from a three-dimensional viewpoint. We can also consider how frequently the collaborators switch back and forth in completing a task. Do they switch repeatedly after each interaction, do they remain in the same view throughout, or do they use some combination of the above styles? Our qualitative study investigates this use of the representations.

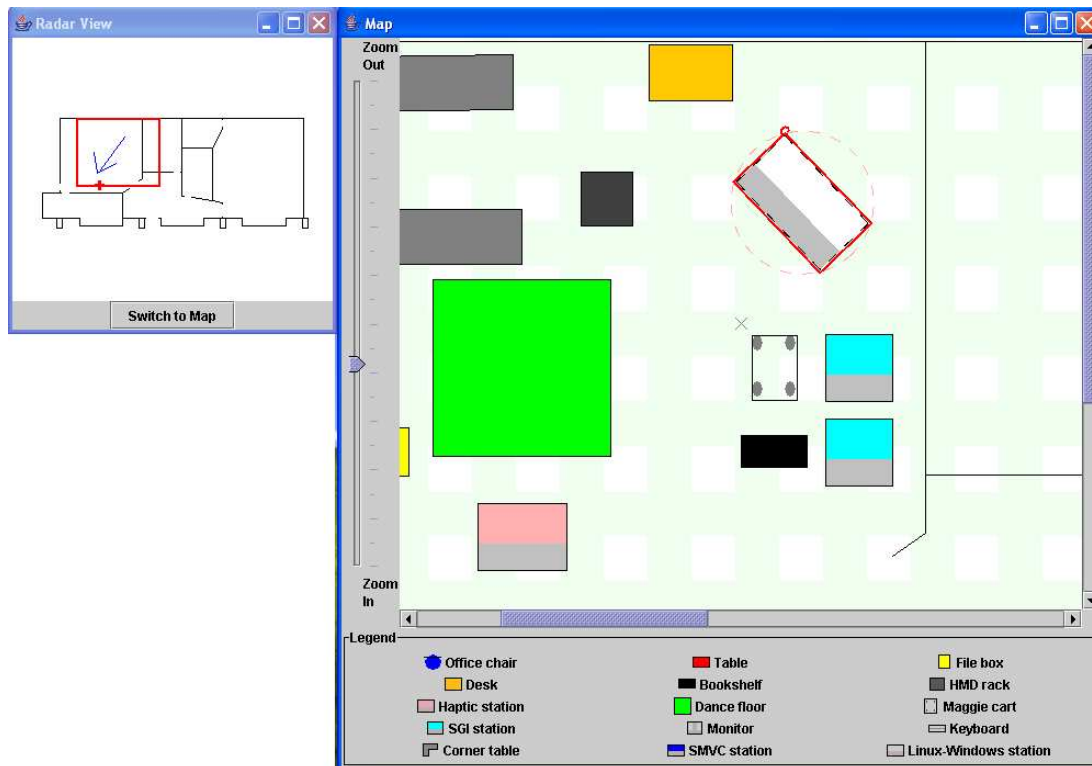


Figure 6.2 Manipulation with the two-dimensional map representation

A collaborator viewing the two-dimensional map is rotating the workstation outlined in red. The radar view indicates a collaborator in the virtual environment.

The prototype also includes awareness features that allow the collaborators to understand each other's work efforts. In particular, a radar view enables each user to visually understand his/her view with respect to the space and with respect to the other users' views (e.g. Begole, 1998; Gutwin & Greenberg, 1998; Smith, O'Shea, O'Malley, Scanlon, & Taylor, 1989). It displays a miniature two-dimensional representation of the entire environment and indicates each collaborator's current viewport of the map. The radar view used in the prototype is unique in that it portrays information for collaborators using both types of representations. It uses a colored rectangle and a crosshair, or telepointer, to indicate one user's viewport into the two-dimensional map and the position of his/her mouse cursor. This telepointer supports gesturing actions typical of same place spatial collaboration (Tang, 1991). For the three-dimensional representation, a v-shaped figure represents one's location and orientation in the CVE, where the angle of the v matches the user's field of view. A line is also drawn to signify the direction in which the user is pointing in the virtual environment. The position of the mouse cursor on the screen is transformed so that a ray is cast into the three-dimensional environment, terminating at the closest object to the user. This ray is then portrayed in the radar view to indicate where the user is pointing.

Within each representation, there are also additional awareness cues about collaborators with a similar representation type. In the two-dimensional map, the colored rectangles and crosshairs of other users viewing the map are shown. These markings are visible when there is an overlap between viewports. If two collaborators are looking at different areas of the map such that their

displays have nothing in common, no markings will appear. Similarly, in the CVE, avatar objects provide a visual representation of each of the other collaborators looking at the three-dimensional view. The avatar represents one's location and orientation within the environment. Colored lines are drawn to indicate the pointing ray corresponding to a mouse cursor position for each user. These markings are visible when the other features in the environment, such as walls and doors, do not block one's view.

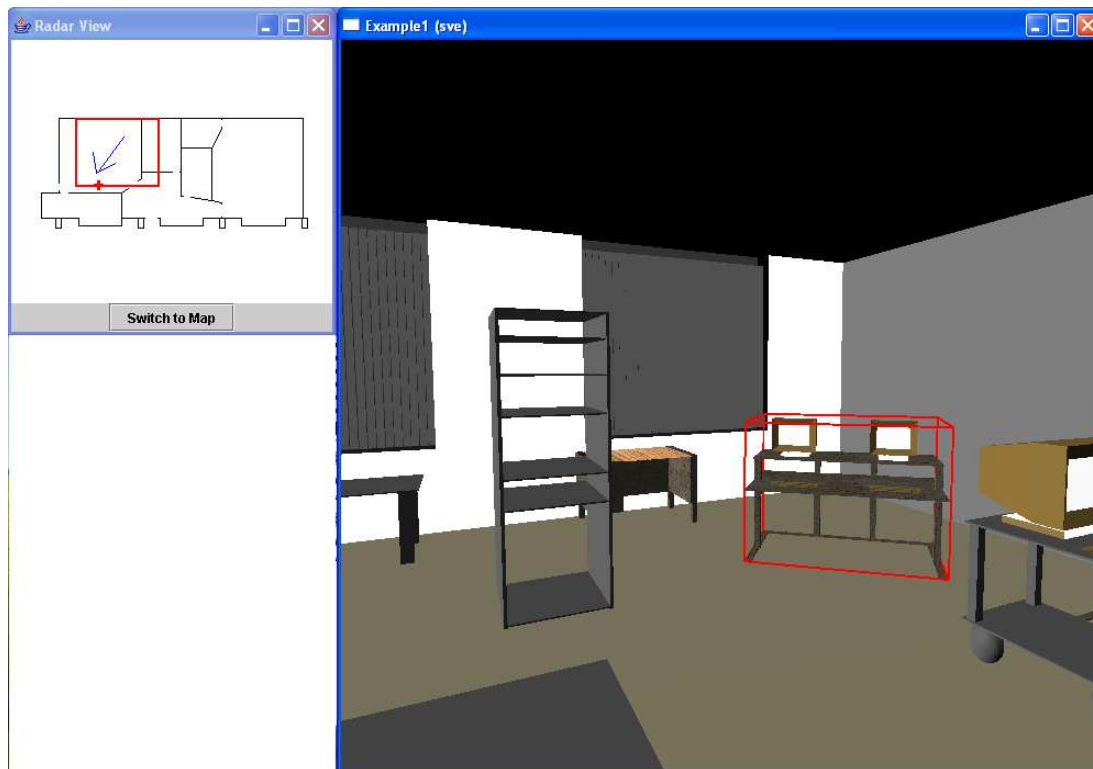


Figure 6.3 Observing manipulations with the virtual environment representation

The collaborator in the virtual environment is observing the rotation manipulation performed by the user with the two-dimensional map.

Each collaborator is assigned a color that is used in the awareness features. This color is used to draw one's markers in the radar map and the two-dimensional map. It is also used to represent the user in the CVE. The avatar and pointer lines both use this color in their markings.

Lastly, the prototype includes a number of moveable objects that correspond to real-world objects. These objects can be manipulated with both types of representations and their movements are shared. When an object's position changes all participants witness this manipulation from their own display, and the changes occur in both the map and the CVE. Interaction is restricted so that only one user can manipulate an object at a time and participants manipulate objects individually, as opposed to as a set, or group.

6.3.1 Interaction Techniques

The prototype incorporates interaction techniques that allow users to navigate both representations, switch between representations, and select and manipulate objects in both interfaces.

6.3.1.1 Navigation

All of the interactions were performed using a standard mouse and keyboard. The map supported zooming and panning interactions through familiar scrollbar and slider widgets. A scrollbar below the map allowed users to shift the map right and left, while another to the right moved the map up and down. The zoom slider was placed to the left of the map so that the users would not to confuse it with a panning interaction, which typically occurs to the right of a workspace (Schafer, Bowman, & Carroll, 2002). The slider and scrollbars were implemented using numerous discrete steps, which gave the appearance of continuous panning and zooming.

In the CVE, the up/down arrow keys would move the user forward and backward according to his/her current orientation, while the left/right arrow keys corresponded to the perpendicular side steps. Holding down the middle mouse button would also navigate the user forward. Movements were constrained to an x-z plane that was six feet tall, but there was no collision detection so that one can walk through walls and doors. Clicking and dragging on the window enabled the users to change their orientation. Dragging the mouse to the left rotated the viewpoint around the y-axis, causing the user to see more of the environment to the left. Dragging it upward enabled the user to look down, and dragging it diagonally would affect both aspects of the orientation.

6.3.1.2 Switching Between Representations

A simple button allows a user to switch between the CVE and the two-dimensional map. When the user presses this button, an animation transitions him/her between the interfaces. The three-dimensional view first orients itself so that the user is looking straight ahead as opposed to up, or down. Then it turns the user around so that he/she is facing the direction that corresponds to the top of the map. Lastly, it smoothly animates from a forward-facing viewpoint to a top-down view of the three-dimensional space, and then switches to the map display. The center of the map parallels the user's prior location in the CVE and the scale is a mid-zoom level.

In order to switch between the map and the CVE, a similar animation occurs in reverse. Using a pop-up menu on the map, the user specifies the desire to switch, where the mouse location corresponds to the new location in the CVE. The map switches to a top-down view of the three-dimensional space and the display animates to a forward-facing viewpoint. Afterwards, the user rotates his/her viewpoint to change orientation.

6.3.1.3 Object Selection and Manipulation

Using the map, objects are selected with a left mouse click and indicated by a colored bounding box. Clicking on a location without an object icon releases the current selection. To position an object, users simply drag the object icons within the map area. Rotation is available through a small, circular handle. Upon clicking this handle, a dashed circle indicates the interaction and the object turns as the user drags the handle.

In the CVE, objects are selected with a right mouse click, and also are indicated by a colored bounding box. Clicking on a location without a selectable object releases the current selection. After selecting an object, a popup window is displayed containing buttons to move a selected object forward, backward, to the left, right, up, and down, and rotate the object left and right around the y-axis. Both forward/backward and left/right movements are based on the user's viewing angle. Selecting forward/backward buttons moves the object in and out of the screen, and left/right moves the object across the screen. There is no collision detection such that the object can pass through all the other objects in the environment, including the floors and walls.

6.3.2 Software Design

The software used in the prototype was written using both Java and the C programming languages. The Simple Virtual Environment Toolkit, or SVE, was used to generate the virtual environment program and interactions (Kessler, Bowman, & Hodges, 2000). The two-dimensional map was written Java and expanded upon the software created for previous map prototypes. Both representations interfaced with the Content Object Replication Kit, or CORK, which provides data sharing and is written in Java (Isenhour, Rosson, & Carroll, 2001).

6.4 Qualitative Study

Both of the previous studies examined spatial collaboration through formal experiments. They used repeated measures, recruited participants with a variety of backgrounds, and included specific tasks with clearly defined stopping points. In contrast, this work explores a more realistic spatial collaboration setting. It examines a small group activity that occurs over multiple sessions. The participants are members of an active research group and the task is both complex and open-ended.

To be specific, the study used four prominent members of the 3D Interaction research group at Virginia Tech. This group meets weekly to discuss concepts and issues related to three-dimensional user interfaces, and the four members we used, three male and one female, are some of the leaders of this larger group. Their primary objective for the study was to explore lab furniture layout designs using the prototype. This is a realistic task for the participants as they were discussing the arrangement of their lab area at the time of the study. The participants were also unique in that they each use the lab differently as members of the larger group. Two of the students also had desk space located in the lab space at the time, giving them a different perspective on the situation, while another student was currently the system administrator for the lab.

Using members of a virtual environments research group could be considered as introducing a bias in the study. These participants are very familiar with computers and the latest technology, including virtual reality. Such characteristics are not typical of most users engaging in spatial collaboration. However, the use of this group allows us to examine an extreme situation of distributed, spatial collaboration: one of expert users who demand a highly useful and usable solution. In this sense, the study offers a rigorous test of the prototype. It also contributes a view of a skilled set of collaborators working on a spatial problem.

The qualitative study was carried out in three phases. The first phase examined the usability of the features in the prototype. With each of the group members as interface designers, their

expertise was helpful in evaluating the usability. The group was given a simple task of moving all the furniture so that it was positioned along the outside walls of the environment. This gave them a chance to become familiar with the interface and provide constructive criticism at the beginning of the study. This phase lasted forty minutes, with twenty minutes spent using the prototype from four different locations and twenty minutes spent discussing its features face-to-face. Next, the group worked on the problem of a new layout design in two different sessions. They were asked not to talk with one another between the sessions about the project, but rather pick up the discussion at the next meeting. The first session lasted twenty minutes and the second session forty minutes. This was the primary phase of the study. At the end of the second session, a focus group discussion lasting fifteen minutes allowed the group to reflect on the problem solving process and their solution in a face-to-face meeting.

During both the usability evaluation and the second phase, the participants were seated at four different computers, which were not located in the lab under discussion. The distributed setting was simulated in a single room by placing dividers between each computer so that the participants could not see one another or each other's screens. The dividers limited their eye contact and use of gesturing, but allowed them to converse verbally without additional technology. This experimental setup was unique in that one researcher could easily observe the interactions of all four participants. Positioning the computers so that they all faced a central location, the researcher would walk around the outside and view the monitors for each participant. Throughout the study, we also recorded the audio conversations and the video of the four displays. Each of the monitors was captured using a scan converter and the four video streams were combined into one video recording to eliminate problems synchronizing the four videos during later analysis.

6.5 Observations

A number of observations were made while the participants worked with the prototype to solve the more complex layout problem. Some of these observations relate to the group's general collaboration approaches, while others relate to their use of the different types of representations. In completing the study, we also noted the participants' use of the tools from the preliminary studies. Specific findings from these previous studies reoccur in this case of distributed, spatial collaboration.

6.5.1 Validating the Previous Studies

The findings from the two prior studies listed both positive and negative effects of features used in the prototype. These features include: the use of continuous navigation in a collaborative map, the use of an egocentric frame of reference in a CVE, and the use of an exocentric-egocentric frame of reference combination. During the qualitative study, we noted the reoccurrence of these findings as a validation procedure.

6.5.1.1 Continuous Navigation Observations

Previously, the continuous navigation technique allowed each collaborator to navigate and interact with the map independently. We observed similar behaviors in this study. There were no instances of one person guiding the other on how to adjust the map, and the participants were able to look at separate areas on their own. During the sessions, the collaborators would pursue different goals and then meet up to discuss the results. In one case, a collaborator suggested two

different functional areas for the space and this prompted two collaborators to move furniture to create these areas. After working independently for a couple of minutes, the collaborators then paused to reflect on their current design.

The observations of this study were somewhat different than the prior one with respect to pointing interactions. Previously, collaborators often struggled to establish a common frame of reference and pointing at a location was accompanied with verbal descriptions to convey the precise location. In contrast, there were limited difficulties in pointing out features when both participants were looking at a map with the prototype. Users knew what their colleagues were referring to and much of the communication occurred through pointing gestures and deictic references such as “here” and “there.” For instance, in discussing a walkway between two locations in the space, one user used his telepointer to gesture as he said: “Walk in through here.” Most of the verbal discussions referenced less spatial concepts such as considering which way a person faces while working and which areas of the room would lead to social interactions.

These comparisons indicate the usefulness of continuous navigation. As users continued to navigate without problems, this confirms a previous positive effect. Our observations also argue against a prior negative effect with respect to pointing. This encourages future collaboration solutions to use continuous navigation.

6.5.1.2 *Egocentric Frame of Reference Observations*

With respect to the egocentric frame of reference technique reused from the prior CVE experiment, we observed many different approaches in this qualitative study. There were some commonalities with the previous study, such as participants getting lost in the space and performing object movements that were time consuming. For instance, one user switched from the map to the CVE while working on the complex task and commented, “I am . . . at a wall . . . whoa, where am I?”. This response was not common in switching between the interfaces, but occasionally users would become somewhat disoriented in the virtual environment. The differences in approaches related more to the use of avatars. Many of the issues and approaches toward using avatars were not repeated. Some of these are positive results. For instance, users did not struggle to provide useful directions to one another when they were each located with the CVE. As well, they did not have issues with spatial transformations, such as confusing left/right due to different viewpoints on the space.

Interestingly, the collaborators in the qualitative study also did not look at each other’s avatars regularly. In one particular case, two users were located within the CVE and both were moving objects with the same goal. The avatars were not visible on either of their displays as they worked, but one user confirmed the other’s presence by saying, “Is that you, Participant X?”, when he noticed someone had selected an object he was interested in. We also did not observe previous strategies, such as following one another through the space, or using the avatar heads as landmarks.

These differences are curious in that they may have been specific to the previous task and the use of participants unfamiliar to the space. Possibly the prior, coerced collaboration task led the users to interact with each other more. In the earlier task, users were required to play specific roles as opposed to the unrestricted structure in this study. Being very familiar with the space

also may have limited these actions as well. Given their knowledge of the space, the users could easily refer to the objects and be confident that their colleagues understood their references.

6.5.1.3 Exocentric-Egocentric Combinations

The use of exocentric and egocentric frame of reference combination was repeated in the qualitative study through the two interface options. The two-dimensional map provided the exocentric perspective, while the CVE offered an egocentric frame of reference. The use and benefits of this unique combination were observed during the study, but the activities were carried out differently than those in the prior work. The exocentric, or map, perspective was used to perform large-scale movements and the egocentric perspective was used to finalize the objects' positions. This agrees with previous findings. The fine-grain movements did not immediately follow the large-scale movements in the qualitative study, however. Very often the large movements occurred, and later in the session someone, potentially a different user, finalized the position. This is because the users worked more independently in the qualitative study and the task was not to position one object, but many.

In the instances where the participants were collaborating with the same object in mind, the participants did not struggle with spatial transformations as in the findings of the previous study. This is similar to the lack of problems with the egocentric frame of reference and can be attributed to the same explanation that users were more familiar with the environment.

6.5.2 General Collaboration Approaches

In general, the collaborators took an ad hoc approach toward the furniture layout problem. Initially, there were comments such as, "What are we going to do?" and suggestions such as, "Should we all go down there and see what's going on?" in reference to the virtual environment. These statements were directed to the entire group, but either no one responded, or one other collaborator agreed to switch to the CVE. Instead, each participant explored the space individually and thought about new design ideas. Eventually, one user made a location suggestion and another participant responded with his/her comments, but the others still continued with an individual approach.

This style of behavior was observed throughout the sessions. Often, it was clear that two people were closely working together. Engaged in an ongoing conversation, they did not reference each other by name, but rather made statements such as, "How about this?" and the other would quickly respond. While their conversation progressed, the other participants remained silent. Sometimes this lack of talking was attributed to being engrossed in a different, individual task, and other times the participants were simply observing and not sharing their own ideas. For instance, one participant said very little throughout the sessions. He would perform adjustments in the CVE, such as rearranging the chairs so that they were paired with monitors and faced the desks. At one point he also commented, "I'm hiding from everyone." This user worked by himself for most of the study. In another example, one user was working with an object in the map and another was located in the CVE. The user with the map commented how the "equipment rack" he believed he was manipulating should be both positioned near the "dance floor" object and rotated so that it faced away from the wall. The user in the CVE heard this and began to move the actual "equipment rack" object to the specified position. When the first user

switched to the CVE and saw his mistake, it was clear that the second was listening to his ideas, as he said: “Is that you, Participant X?”

At a few points in the sessions, one collaborator would verbally attempt to create a second subgroup that was separate from the two participants already working together. Each of these attempts failed. In one instance, one participant asked another if he was moving a chair and asked what his plans were, but the colleague did not respond and so the inquiring participant listened in on the other group’s activities. In another case, a participant in the CVE noticed that she had the matching object held by another participant in the CVE. They began to work on a plan for a good location, but their conversation was interrupted by the other subgroup’s communication and the coordination was abandoned. At other times, a participant would ask the group about an object he/she had selected, but when there was no response, he/she would move it to a location he/she thought appropriate. Again, this shows the independent efforts of the participants.

The subgroups typically changed during error situations. The interruption was one example of such a situation. The subgroups also disbanded when a request to move an object was unresolved, the software was slow, or the participants did not agree with the design that was emerging. Sometimes the participants would offer their ideas or try to make a change, but they frequently chose to work on their own when things did not go as expected.

6.5.3 Strategies Employed

In working on the layout problem, the group adopted two main strategies that they used throughout the task. One of these was to conceive of the space in terms of different areas of functionality. Frequently, in working with the map, the participants would talk about different areas of their design. For example, at one point a user commented: “the students could hang out here and the rest of equipment could be over here.” Similarly, the large “dance floor” object and the necessary equipment became another area of the design. Talking with the group in the final discussion, they also explained their ideas through the different functional areas.

The other strategy employed was to relate the emerging design to their observations and knowledge of the use of the current space. Looking at their design, they would comment on where people would walk, paying particular attention to the space in between furniture and the major walkways through the layout. For instance, one participant mentioned that in the current lab he has about a foot between his desk and the wall so that he can get around the desk. Another participant commented on a row of computers in the current lab layout that are used less frequently because they are on a major pathway. The group also considered working in the spaces they designed. They put thought into which way one would face sitting at a desk, whether there would be a glare from the windows, and whether areas would be used for social gatherings. They also considered the size and location of the windows in the space. Their design did not block the windows with bookshelves or large objects, but rather used this space for short filing cabinets with imaginary plants.

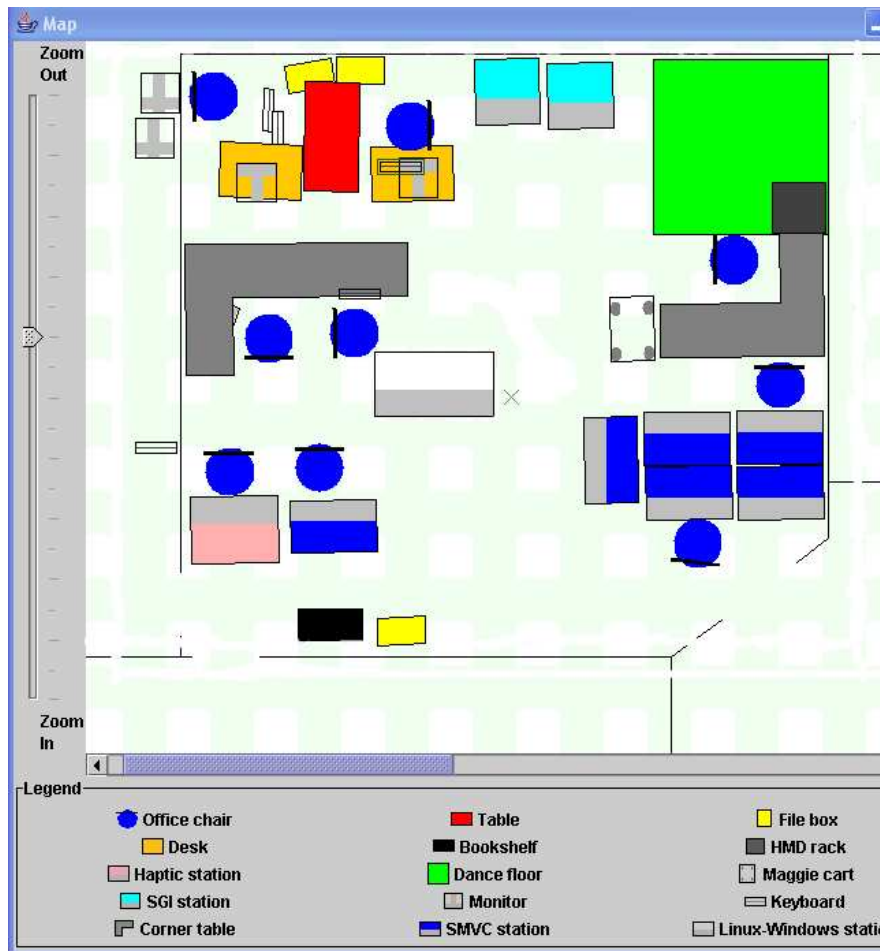


Figure 6.4 Lab design product

The final layout design included different areas for the dance floor, individual desks, and similar workstations. The walkways in between were considered along with the windows of the space.

6.5.4 Uses of the Two-Dimensional Map

In particular, the subgroups would use the map to try out different ideas. After moving a couple of objects, they would stop and reflect on the layout. They would discuss the implications such as which direction someone would face while working, how accessible the space is, how the layout affects the rest of the space, whether the equipment needs were met, etc., and then try another design idea in response. The conversation would include comments such as, “Let’s try this,” “Why don’t you show us what you mean?” and “Yeah, that’s an excellent idea.” The group would frequently work with the large objects and the map using this approach. In both sessions, they spent much of the time repositioning two cumbersome L-shaped desks with the map interface. Once these desks were positioned, they rarely were moved in the CVE.

The map also corresponded to comments about space efficiency. For example, a pair of participants in the CVE noticed there were many workstations of the same type scattered throughout the lab. Switching to the map, they were able to group them so that they took up much less space. Another participant noticed their design and marveled at how it efficient it was.

Lastly, we observed users switching to the map interface in order to comprehend major changes in the layout. For instance, one user working with the map began to reposition many of the objects in a very different design. Another user, who was observing within the CVE, immediately switched to the map when many objects began to move quickly across his display. He knew someone was making a major adjustment and it was much easier to follow this new idea by looking at the same map interface.

6.5.5 Uses of the Three-Dimensional Virtual Environment

Unlike with the map interface, the users did not work closely together when multiple participants were located within the CVE. This interface afforded finalizing the positions of objects in order to experience how the room would feel in actuality. The manipulations performed with the environment were limited to primarily rotations and changes in the objects' height. Participants would rotate a bookshelf so that it was backed against a wall, or rotate a chair so that it sat facing a desk, or position a monitor up on the desk space. After moving objects around with the map, users would frequently switch the CVE to see how it looked in a more realistic setting. This interface allowed them to evaluate their design once they had a layout created with the map that everyone agreed with. For instance, during the session one user announced, "I switched back to 3D and no one can use the computers where they are!" This user noticed that two rows of monitors were lined up so that they faced one another instead of facing outward (see Figure 6.6). Usually, the participants navigated the CVE to look at the spatial relationships between objects, but incorrect rotations, such as the computers, were also quickly spotted.

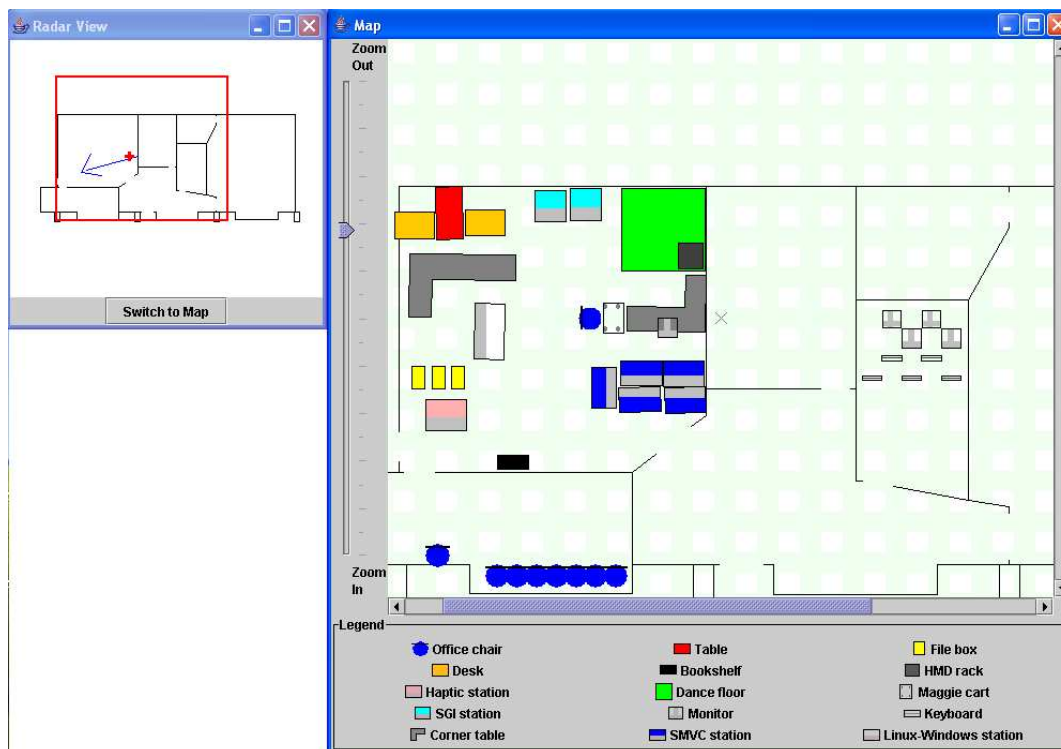


Figure 6.5 Organizing objects with the two-dimensional map

The user looking at the map interface has recently arranged the cluster of blue workstations in the center of the display.

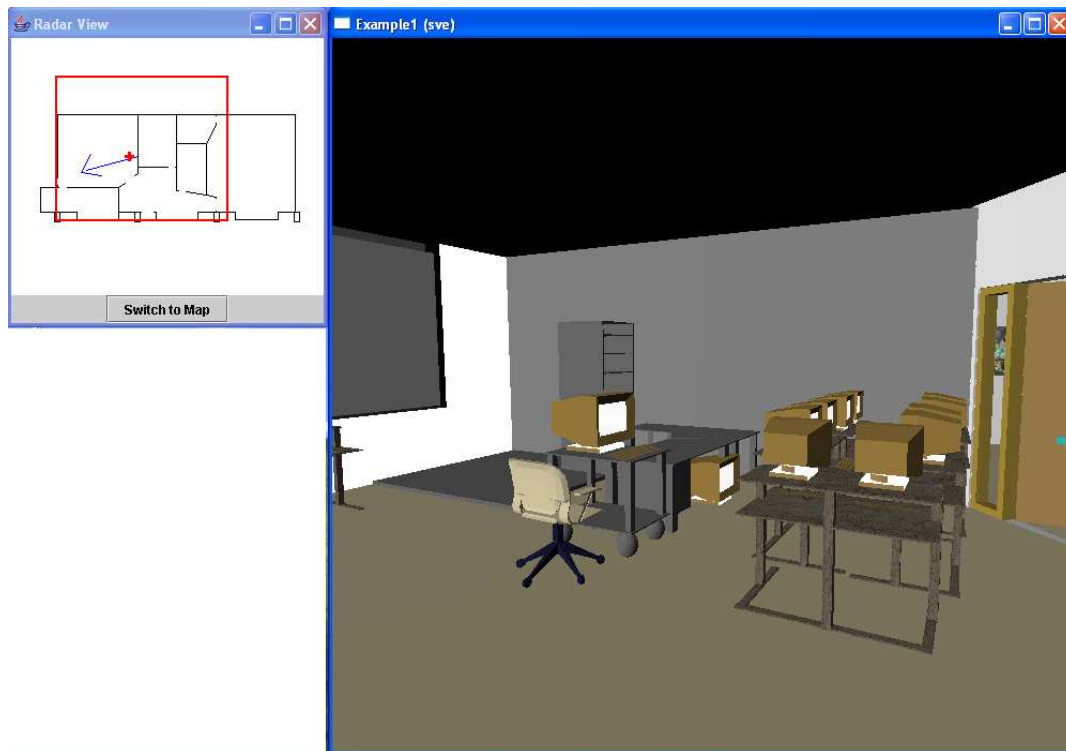


Figure 6.6 Exploring object positions with the virtual environment

The user within the virtual environment realizes the workstations need to be rotated, the chair needs to be aligned, and the monitor under the desk needs to be raised up.

We also observed the participants realizing the magnitude of the problem in the CVE. Users with this interface noticed the objects in disarray more often and would frequently ask “Where should we put Object X?” in an effort to help. As mentioned earlier, we also observed participants in the CVE remark on the number of similar workstations that still had to be incorporated in the design.

6.5.6 Coordinated Activities Using Both Interfaces

On a few occasions during the study the participants would use both the map and the CVE within a subgroup. One participant would work with the two-dimensional map and another would use the three-dimensional environment at the same time. The two representations enabled the collaborators to have two unique perspectives, such that each user would contribute different information. The participants realized this potential early in the sessions. When beginning the task one user asked, “Does someone want to stay in the big room and we’ll send objects to you?” He was implying that one person could work with the objects in the CVE, while the others moved the objects into the room that needed to be designed. One user verbally accepted this position and switched to the CVE, while two others moved objects with the map interface. This coordinated approach was successful until the user within the virtual environment became overwhelmed with the amount of furniture appearing. The collaborators had not considered how quickly one could move an object with the map.

In another case, one user was exploring the space in the CVE and another was simultaneously looking at the map as the map user began to explain why he liked the current layout design. As

he talked through the positive aspects, he began to rotate an object trying to decide which way it should face. The user in the CVE listened in on his description and easily confirmed his last rotation movement with “Yeah, that’s the right direction” (see Figure 6.2 and Figure 6.3). It is unclear if the user working with the map knew of his colleague’s status during the explanation, but he/she may have realized that he/she was in the CVE with this comment. In any case, the map user was able to benefit from someone located in the CVE.

Interestingly, this pair continued to work together with their respective interfaces as they next discussed moving one of the large L-shaped desks across the space. The user with the map quickly relocated the object as the user in the CVE observed and then questioned the rotation of the object. The user looking at the map turned the object in response, and the CVE user confirmed the desired position. The CVE user moved a chair to the desk to complete the area. Again, this demonstrates how the two users were able to work together with two different representations.

Activities taking place between the two representations were not always successful, however. One example during the session occurred when there was a discussion about the amount of space in between two pieces of furniture. The participant looking at the map claimed that there was "a good foot" between the objects and that he could fit in between them, but the user within the CVE was unsure of this assessment. This participant tried to express his/her concern over whether there was really enough space, but the user viewing the map disregarded the comment and did not switch the CVE to see the situation for himself/herself. In this case, there was a breakdown in the coordination that led to a less-than-satisfactory design. In another similar instance, one participant became confused when his/her colleague in the CVE would repeatedly reposition each of the objects he had just moved. These extra movements confused the user looking at the map. The CVE user was moving each monitor and keyboard up from ground level, placing these on top of other objects leading to this confusion. The user with the map could not see these interactions and did not understand them until he/she switched to the CVE interface.

6.5.7 Gesturing Issues

When two participants were using the map, there were few gesturing issues. Occasionally, one user would check to make sure that the collaborators could see his telepointer, or the participant observing the telepointer would ask the other to repeat the gesture. These problems were solved fairly easily.

When the participants used the two different interfaces, gesturing was more problematic. The radar map provided awareness information about all of the collaborator’s work efforts, including their cursor’s position. Understanding the cursor positions was not straightforward, however. On multiple occasions, users viewing the map were unsuccessful in pointing out a location in the virtual environment. For example, one map user stated, “You with the desk, could you move it over here.” The user in the CVE with the desk selected responded with a comment about not seeing the user, and even after the radar view was referenced, the desk was never correctly relocated.

In a similar instance, a user looking at the map pointed out where he thought a computer should be placed. Not understanding this location, the map user encouraged his colleague to de-select the object in the CVE so that he could demonstrate his idea. On another occasion, a user within the CVE had trouble pointing out an object to the other. The CVE user, who was curious about and did not recognize an object, asked if someone else knew what he/she was pointing at. A user looking at the map interface responded, but as he/she could not discern the pointer, he/she switched to the virtual environment to help with the identification.

6.6 Implications

These observations document a complex and realistic example of distributed, spatial collaboration. They also reflect the use of multiple, spatial representations in solving a spatial problem. Both of these have implications for the design of spatial collaboration solutions.

6.6.1 Realistic Setting

Each of the members of the group was engaged in the activity and they contributed through a variety of means. Some participants worked on their own to finalize object positions for the design layout under development, others worked directly with one another to determine the important issues and discuss the tradeoffs in their design. Subgroups were also ad hoc as they formed quickly and dissolved easily when the participants became interested in other aspects of the problem. This description confirms the importance of supporting both independent and collective work in distributed, spatial collaboration solutions. It also strongly encourages the use of awareness techniques that allow the collaborators to maintain an understanding of what each person is doing and how each is contributing to the problem. As each person works independently, information about each person's efforts can ensure that there is coordination between the separate activities. The radar view took a step toward conveying this awareness information. It clearly presented which interface each collaborator was looking at, his/her viewpoint on the representation, as well as the pointing interactions of each collaborator. It did not convey cases where the participants were listening and following along without conversing. It also failed to indicate which aspects of the problem the collaborators were working on, such as the independent participant in the CVE performing the detailed manipulations. This left the collaborators unsure of what their partners were doing and uncoordinated in their group effort.

Also, with so many independent efforts and only one subgroup communicating at a time, the group was unfocused in completing the task. It seemed like the group was not making progress toward a solution, at times, because large objects moved at the beginning would be relocated again toward the middle and end of the sessions. Movements to finalize object positions early in the session also seemed inefficient. As mentioned earlier, one user would frequently stay in the CVE and perform the detailed manipulations. He would often push chairs up to desks, for instance, even though the desks would likely be rearranged later in the session. The group was successful in creating a design layout at the end of the session, but they could have benefited from more process support, or features to help them with the open-ended layout task. Such support could have the effect of making the group feel more coherent and more efficient.

Providing this process support is directly related to the use of awareness techniques. If the users have an understanding of their collaborators' efforts through awareness information, they can also know the group's progress. They will know what areas of the space have been worked on

before, what each collaborator is currently working on, and which subtasks need to be worked on further. This will allow them to be more organized and efficient in completing a task. Again, spatial collaboration solutions need to use more complex awareness techniques than radar views to support this understanding.

Another way to support the users in solving their problem is to add features that enhance the common strategies employed. In the design layout task, the use of different functional areas reflects a need to be able to designate areas of a space and also group objects within an area, so that they can be moved together. During the study, the participants would create these areas and refer to them as they continued, but occasionally an area would be disassembled as a new area was designed. It was unclear if the participants simply forgot about the previous area's functionality, or whether they meant to break it apart. In any case, a novel feature could be added to enhance this strategy. It could allow the users to easily mark and unmark areas, allowing everyone to see the designated area.

Such an area-marking feature would also provide process and awareness support to the collaboration. Through the creation and destruction of areas, the collaborators would have a sense of where their partners are working and what they are doing. For instance, if one participant were to create an area and label it, others would know that he has considered the objects in that area and he is working to associate the objects with a specific purpose. This level of awareness could be beneficial to the others as they work independently and as they choose what action to perform next. Using specific area markings also facilitates a common language between the collaborators and encourages them to refer to their actions with respect to these areas. For instance, in this study, one participant might have offered that he/she was working within a particular area instead of simply trying ideas and commenting about the implications.

Lastly, marking areas could have encouraged more explicit subgoals for the task. As everyone would see the marked areas, this could enable a design discussion about the areas themselves. Participants might consider the functionality associated with the areas and think about alternative organizations. They also could discuss the more abstract area arrangement and the implications of its layout. Thinking about the problem with respect to areas could in turn lead to discussions about the specific objects and their placement within the areas. If such an effect were to occur, the area markings would provide additional process support, facilitating a more defined approach to the complex task.

6.6.2 Multiple Representations

The availability of two representations was both useful and helpful to the collaborators in the study. In particular, it enabled the users to investigate different aspects of the space. Collaborators would use the map to examine the overall layout, while the virtual environment offered a first-person perspective. Other complex, spatial collaboration activities can also benefit from the use of multiple representations. The different dimensional representations, in particular, enable collaborators to view high-level details and three-dimensional spatial relationships.

The prototype also supported a range of activities, as the two representations were used for different tasks. The map afforded large-scale movements. Collaborators would use this

interface to easily explore arrangement ideas and make major adjustments to the design. The virtual environment, on the other hand, allowed the collaborators to experience their design. The participants would switch to this interface to “walk through” the layout they created with the map. They would also observe minor changes such as objects that needed to be rotated or changed in height. There was very little overlap between these tasks. Occasionally the users would rotate an object using the map interface to explore a layout idea, but the final rotations were primarily performed using the CVE.

This demonstrates that each of the representations affords different tasks. Again, this can be helpful for other spatial problems. Many spatial collaboration activities are complex involving numerous steps and numerous types of tasks. Multiple representations that allow collaborators to complete these unique tasks will better address the issues of difficult, spatial problems

6.6.2.1 *Using Different Representations*

In addition to supporting different activities, the combination of the two-dimensional and three-dimensional representations enabled the collaborators to be efficient in manipulating objects. With one user looking at the map and another looking at the virtual environment, they often worked together to position the objects. Each user would contribute different types of information, both of which were important to the task. As each interface afforded different tasks, the users also adopted a natural division of roles. The map user would provide large movements and an overall layout perspective, corresponding to the tasks supported by the map interface. Similarly, the virtual environment user offered the first-person perspective through the CVE interface, contributing object rotation and height details.

The communication between the users with different perspectives was not flawless, however. Often, it was not clear whether the participants knew about each other’s efforts. In a few of instances, the user looking at the map switched to the virtual environment and quickly saw how a colleague had been contributing to his/her task. Sometimes the user noticed a collaborator working in the same area, but was unsure of their specific actions, and other times the collaborator went unnoticed.

There were also issues in expressing ideas when the collaborations used the different representations. In particular, participants struggled to gesture their ideas to one another. In our prototype, the radar map provided the only awareness information between the representations. Users working with the same interface saw viewport rectangles, avatars and telepointers for their collaborators, but similar visual cues were not provided between interfaces. Hence, pointing out a location with the map led to confusion in the CVE as the user needed to understand their viewpoint with respect to the radar map. Likewise, in pointing at an object in the CVE, a map user would have to align the collaborator’s telepointer in the radar view with their two-dimensional view of the space.

Both the communication problems were due to a lack of awareness, limiting the usefulness of the different representations. Previously we discussed using more complex awareness techniques that allow the collaborators to know about each other’s work efforts. Similarly, techniques that convey user activities with different representations will reduce the problems of not knowing about collaborators working in the same area. Also, with respect to gesturing, this study reveals

the importance of conveying pointer information between representations, despite their differences. If similar spatial features can appear in both representations, collaborators need a way to refer to them across representations. This extends Tang's emphasis on pointing and gesturing abilities within a shared workspace (1991).

One way to improve our current prototype in light of these problems is to incorporate additional awareness features. Novel techniques that represent collaborator interactions with the CVE in the map interface, and similarly, map interactions within the CVE would provide more awareness information. Users would not have to rely on the secondary, radar view display, but could see their collaborators' viewport indication and telepointer within their primary interface display. For instance, the map could be augmented with the same collaborator markings provided in the radar view. A v-shaped figure and corresponding pointer line would allow users of the map to know of their partners' location and orientation in the virtual environment, as well as the direction of their pointing.

Adding the map users' information to the CVE is not as straight-forward, however. Map viewports are rectangular areas, which are difficult to represent in a virtual environment. One possibility is to apply lighting effects to indicate these areas, where different color shadings of light could represent different map users. This may have usability problems as map viewports overlap. Another approach is to simply portray the map user's cursor within the CVE. This facilitates pointing and leaves any viewport differences to be verbally negotiated. The virtual environment user's view will most likely be limited by occluding objects and he/she will only see a part of the map user's viewport, anyway. This makes the bounds of the map user's viewport less significant.

Modifying the current prototype with these awareness techniques technique would have definitely lessened the collaborators' problems with gesturing between the two representations. It also could have helped the participants know about their colleague's work efforts within the same area. In looking at the map, CVE users manipulating objects within the same area would be visible. Their position within the virtual environment would be close to the objects being manipulated and this location would be indicated on the map. Likewise, a collaborator in the virtual environment might notice the map-based pointer of a colleague interacting with objects in the same area.

6.6.2.2 *Using Similar Representations*

During the study, the participants would frequently use the two-dimensional map together. They used this interface to discuss the overall furniture layout. They would comment on the primary walkways through their design, the different functional areas of the space, and how different collections of objects divided up the space. They also collaboratively explored different object positions with this interface. One collaborator would move a couple of objects to arrange out a small area, and the others would reflect on the design. The map representation enabled them to have a common frame of reference for these activities. With similar views of the space, they easily understood one another and did not struggle with verbal references to objects or pointing interactions to reference locations.

While different representations allow collaborators to explore different aspects of the space, having the same representation is also advantageous in spatial collaboration solutions. This similarity reduces the problems of miscommunication due to very different viewpoints. This study observed collaborators switching representations to create this commonality. Similarly, other spatial collaboration solutions should enable collaborators to view the same representation. Similar representations also should not be discounted in favor of the multiple perspectives afforded by different representations. Having a similar perspective on the space allows collaborators to simply observe as others express their ideas.

6.7 Conclusions

This chapter explores the concept of distributed, spatial collaboration in a realistic setting and investigates multiple representation techniques for supporting this style of collaboration. In particular, it introduces a novel prototype that reuses favorable techniques from previous studies. Observations of this prototype with an active research group and a realistic task validate the findings from previous studies and offer new insights for spatial collaboration.

The validation results reveal that continuous navigation continues to be useful and usable in distributed settings. Egocentric navigation within the CVE, on the other hand, corresponded to different behaviors. Users with the same egocentric perspective rarely interacted with one another, unlike the previous findings. Also, combinations of the exocentric and egocentric perspectives again led to efficient object manipulations but were often accomplished through a different approach. The fine-positioning manipulations in such combinations did not immediately follow large-scale movements for each object, due to the number of objects in the task.

This study also emphasizes the need for awareness techniques that allow the collaborators to maintain an understanding of what each person is doing, and how each is contributing to the problem. It discusses how supporting the task process by enhancing common strategies is one way of increasing collaborator awareness. One example is to incorporate an area-marking feature for map interfaces that enables users to discuss areas within a space. Such a feature would allow participants to know where others have been working, where they are currently working, and which areas have yet to be discussed. Another example is to visually represent each collaborator's viewport and telepointer within both interfaces. This study identifies that the multiple representations afforded different activities, but that collaborators needed greater awareness when working between representations. Providing visual indications in both representations allows the collaborators to better understand what each person is doing.

The work contributes a rich description of a distributed, spatial collaboration activity. Observing an actual group permits a discussion of their communication patterns, strategies, approaches, and dynamic use of subgroups. This work also offers a novel prototype that uses multiple representations to support complex, spatial collaboration tasks. Two-dimensional and three-dimensional interfaces are combined to allow users to explore different aspects of the space. Through the evaluation of this prototype, the research presents the benefits of similar and different representations within a collaborative activity. Lastly, this study is unique in its evaluation setup. Placing the collaborators in the same room with dividers allowed a single

researcher to observe the multiple participants interact. It also eliminated the need for an audio-conferencing tool.

Chapter 7

7 Conclusion

This work has explored the field of distributed, spatial collaboration. It examined how shared, interactive representations can support spatial collaboration activities. These representations reflect the spatial problem the collaborators are working to solve. They portray important features from the physical space and they support synchronous discussions regarding these features. By investigating viewpoint constraint and awareness techniques, this research looked at how the individual collaborators view the representation, and how they stay aware of one another's work efforts. It contrasted different technique combinations as the collaborators used the same techniques, different techniques, and as they chose between techniques. The work considered the collaborators' approaches to problem solving, as well as the product of their work, their efficiency, precision, and solution completeness. It also considered the collaborators' experiences in terms of satisfaction and preferences.

Three main conclusions can be drawn from these evaluations. This work offers contributions that will benefit both the fields of Computer-Supported Cooperative Work and Collaborative Virtual Environments. It advances the knowledge of spatial collaboration, and highlights the need for a continued study of designing spatial collaboration solutions.

7.1 Contributions

This work contributes to the fields of Computer-Supported Cooperative Work and Collaborative Virtual Environments in four ways:

1. Through three unique studies, it offers comparisons between two-dimensional representation techniques and three-dimensional representation techniques. It also explores spatial collaboration in a realistic setting and investigates combining multiple representations.
2. It presents a set of lessons learned about distributed, spatial collaboration activities. These lessons reflect on the commonalities between the studies, and are based on first-hand experience in conducting them.
3. It offers novel prototypes that support realistic, distributed spatial collaboration tasks. These include the spatial discussion tool from the preliminary studies, the fisheye radar map design from the two-dimensional study, and the prototype of the last study that combines two-dimensional and three-dimensional representations.
4. It provides unique ways to evaluate spatial collaboration, offering interesting tasks, metrics, and experimental arrangements that can be used in future studies.

The next sections discuss the first two contributions in detail. Discussions of the prototypes and the evaluations conducted appear in other chapters.

7.1.1 Summary of Findings

In an investigation of viewpoint constraint and awareness techniques, this research is conducted through four phases. Each phase corresponds to a different study with different research interests. These phases explore different aspects of distributed, spatial collaboration.

The first phase is a preliminary study. It presents an initial prototype to support distributed, spatial discussions, and motivates the later phase of this work.

The second two-dimensional map study provides a formal comparison of two viewpoint constraint techniques and two awareness techniques. It looks at the individual and combined effects of the techniques as the pairs of collaborators interact using the same representation. The following summarizes the findings:

- The study contrasts discrete and continuous navigation. Discrete navigation supports easily aligning viewpoints and pointing out features, but it is difficult to navigate. Continuous navigation, on the other hand, is easier to navigate and is preferred for collaboration. In deciding on locations, pointing occurs equally often as lengthy, verbal discussions with continuous navigation.
- The study also contrasts a traditional radar view with a novel, fisheye radar view. The fisheye radar view supports users easily pointing out features, but the current design should be improved so that it is more useful. In comparison, the traditional radar view is intuitive to use, but does not encourage pointing.
- Upon examining the interactions between the navigation and awareness techniques, we discovered that the discrete-fisheye combination limits the collaborators' understanding of the space. The continuous-fisheye combination, in contrast, indicates the potential usefulness of the fisheye design.

The third study uses a three-dimension, collaborative, virtual environment to investigate similar and different viewpoint constraint techniques. It looks at pairs of collaborators playing specific roles in the activity, where one user provides the instructions and the other manipulates an object. The following summarizes the findings:

- The study contrasts assigning the actor to an egocentric and exocentric frame of reference. The egocentric actor can easily perform fine-grain manipulations, but is inefficient in exploring the space and conducting large-scale movements. In reverse, the exocentric actor can easily perform large-scale movements, but struggles with fine manipulations.
- The study also contrasts assigning the director to an egocentric and exocentric frame of reference. The egocentric director provides limited help, while the exocentric director is simply less helpful with a task.
- Examining the frame of reference combinations, two egocentric users take longer to complete a task, but easily understand each other. Exocentric pairs have a good understanding of the layout, but struggle to see the details and use a common reference

system. Egocentric-exocentric combinations are efficient at object manipulations, but have issues with the spatial transformations inherent in the different perspectives.

The fourth study explores distributed, spatial collaboration in a realistic setting and investigates the combined use of two-dimensional and three-dimensional representations. It reuses techniques from the previous studies and validates the findings with a different application. The following summarizes the findings:

- In validating the results, continuous navigation continues to be useful and usable. Egocentric navigation, in contrast, corresponds to different behaviors as the participants with similar egocentric perspectives rarely interacted. Combinations with exocentric and egocentric perspectives again lead to efficient object manipulations.
- Observing a complex task, the collaborators took an ad hoc approach, frequently working on their own and occasionally working closely in pairs. The subgroups were dynamic and changed throughout the sessions. One strategy is to conceive of a space in terms of different areas of functionality. Another strategy is to relate the floor plan design to the collaborator's observations and knowledge of current space uses.
- Combining a two-dimensional and three-dimensional representation allows collaborators to work together to explore different aspects of the space. A two-dimensional map affords reflecting on the overall layout of the space, while a three-dimensional representation affords experiencing how the space feels in actuality.

7.1.2 Lessons Learned

In conducting the later three studies of distributed, spatial collaboration, a number of common results are apparent. These lessons learned provide a better understanding for supporting distributed, spatial collaboration. The future work of this research also directly relates to these findings, as the implications for design are not always clear.

7.1.2.1 Individual Navigation Requirements and Benefits

One of the primary observations from the two-dimensional study is that collaborators need to be able to navigate and explore the representation individually. This allows each participant to contribute to the problem. As the users understand their own view, they are better prepared to understand their collaborators. They do not struggle to comprehend their individual display during the task, but rather can focus on what their collaborators saying. The other evaluation sessions provide further examples of this. If the participants became disoriented in working with the three-dimensional representation, they would simply state "I'm lost," and focus their attention on orienting themselves rather than contributing to the collaboration effort. This demonstrates the need for usable viewpoint constraint techniques to ensure efficient collaboration.

When the collaborators can navigate their independent viewpoints without becoming confused or disoriented, this allows them to collaborate from different perspectives. In the two-dimensional study, this observation was clear as the participants effectively solved spatial puzzles. Each user was given spatial criteria, and working together, they would individually explore the

representation to satisfy both criteria. This is one example of how individual viewpoints enable collaborators to solve complex spatial tasks. Similarly, in the three-dimensional study, the collaborators would work together to take advantage of their unique viewpoints on the representation. When both participants both used an egocentric or exocentric frame of reference, they would frequently navigate to different viewpoints to acquire different perspectives on the space. For instance, in positioning a picture frame against the wall and centered above the bed object, one egocentric user might position himself/herself at the foot of the bed, while another collaborator looked on from a view close to the wall.

The advantages of individual viewpoints are not limited to the cases where the collaborators use the same viewpoint constraint techniques or the same representation, however. In the three-dimensional study, two of the conditions investigated had the collaborators use different frames of reference. One user would use egocentric viewpoints, while the other had exocentric viewpoints. Similarly, in the last phase, it was possible for one collaborator to be looking at the two-dimensional map, while another user was located within the virtual environment. In both of these evaluations, the different viewpoints complemented one another. The results from the three-dimensional study state that the combination supported both fine and coarse grain object manipulations. The user with exocentric viewpoint would provide information relating to the large-scale movements across the environment, while the egocentric user contributed the smaller, positioning details. Likewise, the map user in the last study was observed moving the objects, while a collaborator in the virtual environment offered rotation guidance. These cases illustrate the usefulness of multiple perspectives afforded by individual navigation in solving complex tasks.

In summary:

Collaborators need to be able to navigate and explore the representation individually. This allows them explore the space from different perspectives and contribute different information to the collaboration task.

7.1.2.2 Support For Assistance

Another observation from the two-dimensional study that reoccurs in the other evaluations is the social practice of collaborators assisting one another. In working with the spatially complex, discrete navigation technique, the users with greater spatial abilities frequently helped their partners with less spatial skills in navigating the map. They would direct them to align their viewpoint with their own, occasionally instructing them with the specific interactions required. Similarly, in the studies with the three-dimensional representations, participants would aid one another if someone became disoriented or confused on which way to navigate. With both participants using an egocentric frame of reference, some collaborators would search the environment for their “lost” partner and then direct them on which way to go. With collaborators using different frames of reference, directions from the exocentric viewpoint were also frequent.

The desire of collaborators to help one another goes beyond individual problems with navigation. Collaborators also offer assistance when their colleagues cannot identify features in the

representation and when they struggle with object manipulations. For instance, in the two-dimensional study the participants would remind one another of the legend listing the icons used on the map. When both users were working with the three-dimensional representation, they would frequently confirm the identity of the objects within the environment and assist one another in positioning objects. The collaborators would explicitly say: "I'll help you position it," as they provided directions such as "move it up/down and left/right." These actions were also observed in the last phase of this research with one collaborator looking at the map interface and the other looking at the virtual environment interface. The user within the virtual environment would offer additional information about the object being manipulated by the map user. This collaborator would confirm the map user had the correct object selected and indicate the appropriate rotation desired.

These acts of assistance are common in distributed, spatial collaboration activities, and it is important to consider how collaborative software design can be supportive of them. In some of these scenarios, the help provided is very specific. A collaborator will direct someone else on an interaction to perform, such as clicking a button, dragging a scrollbar or slider widget, or pressing a key. In other cases, the help is more abstract, such as mimicking someone else's viewpoint or providing directional information. One way to support these cases is to enable collaborators to easily align their viewpoints. This allows the collaborator playing the helper role to take on the view of their partner and walk them through the adjustments needed, or it allows the collaborators needing help to quickly return to what one hopes is a less disorienting viewpoint that is the same as their partner's. Another approach is to support the directional assistance cases by enabling one collaborator to gesture directions to another. The prototypes in this work are limited in such gesturing features, which can be beneficial to the assistance.

In summary:

The social practice of collaborators assisting one another is common in distributed, spatial collaboration. Software designs should be supportive of these additional activities.

7.1.2.3 Visual Attention

In observing each of the studies in this work, it is also interesting to note the gaze of the participants' eyes. Frequently, the collaborators are extremely focused on the spatial representation. They continuously stare at the screen and become fixated on the user interface. Often, this leads them to pay less attention to what other people are saying. For instance, in all of the evaluations, participants would respond positively to their collaborators with short phrases like "uh-huh" and "yeah," although it seemed that they did not truly comprehend what was said, or asked. The participants would be consumed by their own thoughts and respond to simply indicate that they registered the other participants comments. On other occasions, the participants would make a conscious effort to stop pursuing their own activities in order to listen and follow their partner's ideas.

In one light, this finding is encouraging. It indicates that the participants were engaged in the spatial collaboration task and that they were honestly interacting with the representation during the evaluations. On the other hand, this finding implies that representations can lead to

inefficient collaboration as the users become engrossed in their own activities. One approach to resolving this issue is to ensure that each collaborator is able to navigate the representations individually. With highly usable viewpoint constraint techniques, the participants will less likely be fixated on the user interface because they do not understand it. This is a primary explanation for the numerous stares at the interface that we observed to be associated with the complex navigation techniques of the two-dimensional study. Another response is to follow the suggestion outlined in the final phase. This suggestion is to support independent activities, but provide ample awareness information so that the collaborators maintain an understanding of what each person is doing and how they are contributing to the problem. This allows the collaborators to focus on their own activities, but it also encourages coordination among the collaborators' work efforts.

In summary:

Users can become engrossed with a spatial representation and their own activities. In addition to supporting independent navigation, it is important to use awareness techniques to encourage coordination.

7.1.2.4 Many Spatial Concepts

Throughout the three evaluations, the collaborators used primarily six different spatial concepts to convey their thoughts involving spatial relationships. These included: location information, direction information, orientation information, area information, grouping information, and distance information. Location information is the most obvious of these. The collaborators point out specific locations to one another and use words such as “here” and “there” as they look at the representation. Direction information involves navigating one's viewpoint, moving and manipulating an object, and referencing other routes through the space. For instance, one collaborator might instruct another on how change his/her viewpoint so that that user is looking at a different area of the space. One collaborator also could offer another directions on how to move an object, or he/she could discuss the common walkways of the space. All of these involve direction information.

Similar to direction information is orientation information. This encompasses rotating one's viewpoint, rotating an object, and referencing other orientations and rotations relationships within the space. For example, a collaborator might describe which way he/she is currently facing, or which way an object should face. One could also consider the orientation between two locations, such as one object being located north of another; or the orientation of one location with respect to the rest of the space, such as an object's location against a wall and facing into the center of a room.

Area and grouping information can also be related. Area information references sections of a large space, or designated areas. For instance, rooms within a floor plan provide distinct areas. Groupings, in contrast, are sets of objects. The objects have some commonality, such as needing to be located near one another spatially. For instance, one collaborator might group a television and a couch together and then converse with others about other objects to add or where to place the group. Often a group of objects is situated with an area, but the area is actually a separate

idea relating to the spatial layout. Distance information refers to the measurements between locations. For instance, one object could be located five paces away from another object.

Collaborators need a way to share these spatial concepts while completing distributed, spatial collaboration tasks. The prototypes in this study supported some of these interactions, but additional techniques are needed. The participants could easily share location information through the telepointers, but indicating directions, orientations, areas, groupings, and distances were less straightforward. The collaborators would frequently gesture with the telepointer and resort to detailed verbal descriptions in the event that such gesturing fails to convey their original intent. The last study also highlighted the importance of collaborators sharing pointers across multiple representations. Looking at a representation that is different from one's fellow collaborators should not preclude one from exchanging ideas.

Three-dimensional representations, such as virtual environments, are more complex than two-dimensional maps with respect to spatial concepts. They encourage spatial information to be specified in three-dimensions. However, these representations benefit from the use of avatars, unlike two-dimensional representations. Avatars are one way facilitate an expression of location, direction, and orientation information. A collaborator can position him/herself in the space so that he/she suggests a location and orientation. Navigating through the space, the avatar can also indicate directional information. Providing spatial information within a virtual environment is also difficult because each collaborator sees the representation from a different orientation. In the three-dimensional study, the collaborator would frequently offer information with respect to his/her own viewpoint, and struggle to perform the mental transformations required to describe the idea from another's perspective. One future aspect of the work of this study is to explore additional techniques that enable the collaborators to share these spatial concepts.

In summary:

There are at least six different types of spatial concepts that can be expressed in spatial collaboration. Interaction techniques are needed that enable collaborators to easily convey these spatial relationships.

7.2 Advances in Understanding

At the beginning of this work, current collaborative software did not address the challenges of distributed, spatial collaboration. Collaborative solutions that included a shared, interactive representation did support spatial activities such as discussing existing spaces, investigating multiple types of spatial data, or exploring distances. Shared whiteboards are one such example. Other collaborative software that enabled a spatial representation of a physical place lacked the support for exchanging ideas. This technology allows someone to present a spatial representation, such as a map, to other collaborators by sending an email with an image, or posting the representation on a Web page. Some of the preliminary work initiated the question of how to design solutions for distributed, spatial collaboration. It presented a survey of interaction techniques for map software, reviewing techniques used in common Web-based map applications and commercial map software. These techniques suggested possibilities for distributed, spatial collaboration, but their applicability and usefulness were unexplored.

The later three phases of this research take a large step toward understanding the requirements imposed by distributed, spatial collaboration. In particular, they show that shared, interactive representations can indeed support spatial collaboration activities. These phases present prototypes and evaluate them with realistic users. The prototypes are unique, as each addresses a different spatial collaboration activity, and as each explores different features to support spatial collaboration. The evaluations investigate these activities using different participants with each prototype. In each evaluation session, the users successfully complete the task. Although some sessions last longer than others and some collaborators encounter more communication problems, all of the groups finish the task. This demonstrates the utility of the shared, interactive representations.

These evaluations are beneficial also since they provide valuable feedback about the prototypes and their features. Some designs prove to be more supportive of distributed, spatial collaboration than others. The evaluations reveal the advantages and disadvantages of different representation techniques. They describe differences in two-dimensional navigation techniques. They also contrast approaches to providing awareness information through radar view displays. Examining three-dimensional representations, the evaluation findings discuss different frames of reference and different collaborator roles. They compare egocentric and exocentric frames of reference and consider a collaborator who performs object manipulations, as well as a collaborator who plays a director role. The evaluations also report the effects of using similar and different exocentric and egocentric perspectives among different collaborators. In one case, both egocentric and exocentric perspectives correspond to a three-dimensional representation, and in the other a two-dimensional map provides the exocentric view. Lastly, the evaluations reveal uses for different dimensional representations and how these unique representations can support complex, spatial collaboration tasks.

In addition to these findings about specific techniques, the evaluations expose the numerous possibilities for viewpoint constraint and awareness techniques. As each study examines a specific set of techniques, the results suggest prototype modifications. They describe improvements to current features as well as additional features to include. These suggestions are ways to enhance the support for spatial collaboration. Often, there are many approaches to implementing these suggestions. For instance, the study of two-dimensional navigation encourages combinations of the two techniques evaluated. It reports positive effects of discrete and continuous navigation and offers integrated approaches to capitalize on the positives. Similarly, there are many options in designing exocentric and egocentric frame of reference techniques for a three-dimensional representation. This work constrains the egocentric viewpoints, so that the user views the representation from a specific height. Other implementations might support viewpoints at floor level, at ceiling level, and in between. The exocentric viewpoints are also limited to the views available along a half sphere above the space, while other techniques could allow zooming. The last phase of this work also suggests additional awareness techniques to provide information across different representations. Such techniques should allow someone working with a two-dimensional map to know where a collaborator viewing a corresponding three-dimensional representation is positioned and vice versa. This research points out these numerous possibilities for viewpoint constraint and awareness techniques. It also initiates an exploration of additional techniques.

Considering the beginning of this research, there was also little known about how people solve spatial collaboration problems and how to support this work from different locations. The issues of spatial collaboration stated that people had different spatial abilities and different understandings of the same location. It was also known that many of the ideas conveyed in spatial collaboration would require an understanding of the space. The implications of these issues within a distributed setting, however, had yet to be investigated. With respect to distributed software, the lack of the nonverbal, face-to-face cues was a concern, as was the need to provide additional awareness information to recreate the social work setting. It was unclear how these distributed collaboration issues related to spatial collaboration. It was presumed that distributed, spatial collaboration would have greater awareness needs as people frequently use gestures to convey location and direction information in face-to-face settings.

This research informs our understanding of spatial collaboration activities and distributed, spatial collaboration activities, in particular. Numerous evaluation sessions were conducted to complete this work and each one describes a distributed, spatial collaboration scenario. Using spatial collaboration tasks, the sessions provide a description of collaborators solving a spatial problem together.

These descriptions reference how people work together to understand spatial representations. The representations represent features found in the real physical space and people will collaborate to understand these features. For example, a pair of collaborators might follow each other's navigation movements as they explore the representation together. The descriptions also reveal the techniques people use to establish a common ground. With a common ground, the collaborators have a shared understanding of the space and use a common set of names to reference features. To establish this level of communication, collaborators might navigate to similar or identical viewpoints, use unique terms, or reference objects in real physical space.

The evaluation sessions also illustrate common behavior patterns of distributed, spatial collaboration. They reveal how collaborators divide up the work and how the participants play certain roles in the activity. For example, some participants are natural leaders that take charge, while others are more team players who choose to work closely with their collaborators. The sessions also provide realistic examples of spatial concepts collaborators desire to express. Working on a spatial problem, the observations illustrate collaborators pointing out locations, describing directions, gesturing rotations, measuring distances, etc.

All of these records of common activities, working together to understand spatial representations, establishing a common ground, behavior patterns, and common spatial concepts, provide a better understanding of how people solve spatial collaboration problems in a distributed setting. They document some of the requirements of distributed, spatial collaboration and encourage designs to support these social processes.

7.3 Future Work

This focus of this work examines a small, yet important, aspect of using spatial representations to support distributed, spatial collaboration. This prompts a further investigation of representation and interaction techniques. There are likely other useful ways for representing the physical

space and more creative interactions that enable collaborators to share ideas. A continued investigation will lead to the design of novel solutions for spatial collaboration that address the unresolved issues.

One way to continue this work is to evaluate additional viewpoint constraint and awareness techniques. The findings report on the numerous possibilities for these techniques. For instance, they suggest other approaches for navigating a two-dimensional map. The two-dimensional study also encourages investigating various design improvements for the fisheye radar view. In the last phase, incorporating similar awareness information across representations is also strongly suggested. Making these improvements and investigating different design options further explores the design space for spatial collaboration solutions. It reveals the usefulness and usability of different techniques.

This research also stimulates a future study of interaction techniques for distributed, spatial collaboration. In its discussion of common spatial concepts used, this work encourages the design and evaluation of techniques that support sharing spatial information. Activities such as labeling locations, providing route-based directions, and indicating orientations lacked corresponding interactions in the prototypes evaluated. Similarly, collaborators required features that enabled them to mark areas, group objects, and measure distances during the sessions. The awareness techniques implemented in prototypes provided some support for conveying spatial concepts, but additional techniques are required. Such techniques should be simple and easy to use, similar to the non-interactive awareness information, so as to not interrupt the communication between the collaborators. The techniques also should not distract the users from their primary task.

Another avenue for future research is to investigate additional complex tasks. The last phase of this research took an initial step at examining realistic distributed, spatial collaboration. It involved an active group discussing furniture layout designs. This is just one application of spatial collaboration, however. The introduction of spatial collaboration mentioned many other tasks. It discussed members of an outdoor club working together to maintain a set of hiking trails and a collaborative effort to preserve and restore a local watershed. Future research should work to support applications such as these.

Research into additional spatial collaboration applications should first consider the current set of findings. Possibly there are other spatial concepts expressed in these different settings that were not mentioned here. Also, these applications may require different spatial features. For instance, the watershed collaborators may need to look at layered data sets, where each layer can be toggled on and off and the users can change the color scheme of the data. How is this handled with collaborative software? Another possibility is that the collaborators need a way to exchange textual information in addition to spatial concepts. Perhaps they want to associate map locations with specific calendar dates. What novel designs can support an integrated information exchange?

In order to conduct this future research, one approach is to focus on local community groups with environmental interests, paying particular attention to their spatial collaboration needs. Such groups are a natural extension to this work as they typically involve non-technical people

working from different locations. These groups also present complex, spatial problems as natural areas present multifaceted spaces with unique terrains, vegetation, and water sources. The outdoor club and the watershed efforts are good examples. These groups encourage this research to consider both synchronous and asynchronous collaboration. This work reports on only synchronous activities, but asynchronous work is also common as people are not always available to work together from different locations synchronously.

8 References

- Bartlett, Joel F. (1994). Ramonamap: an example of graphical groupware. In proceedings of the *ACM Symposium on User Interface Software and Technology*, 83-84.
- Begole, J.M.A., (1998). Flexible collaboration transparency: supporting worker independence in replicated application-sharing systems. PhD Dissertation. Virginia Polytechnic Institute & State University. Retrieved from <http://scholar.lib.vt.edu/theses/available/etd-011999-152714/unrestricted/BEGOLEDISSERTATION1998.PDF>.
- Benford, S., Reynard, G., Greenhalgh, C., Snowdon, D., & Bullock, A.(2000). A poetry performance in a collaborative virtual environment. *IEEE Computer Graphics and Applications*, 20(3), 66-75.
- Brewer, I., MacEachren, A.M., Abdo, H., Gundrum, J. & Otto, G. (2000). Collaborative geographic visualization: enabling shared understanding of environmental processes. In *Proceedings of IEEE Information Visualization Symposium*, 137-141.
- Billinghurst, M., Kato, H., Kiyokawa, K., Belcher, D., & Poupyrev, I. (2002). Experiments with face-to-face collaborative AR interfaces. *Virtual Reality Journal*, 6(3), 107-121.
- Card, S.K., Mckinlay, J.D., & Robertson, G.G. (1990). The design space of input devices. In proceedings of *Computer-Human Interaction*, 117- 124.
- Carpendale, M.S.T., Cowperthwaite, D.J., & Fracchia, F.D. (1997). Making distortions comprehensible. *IEEE Symposium on Visual Languages*, 36- 45.
- Carpendale, M.S.T., Tigges, M., Cowperthwaite, D.J., & Fracchia, F.D. (1998). Bringing the advantages of 3D viewing into focus. *IEEE Conference on Information Visualization*, 17-20.
- Carroll, J. M., Rosson, M. B., Isenhour, P. L., Van Metre, C.A., Schafer, W.A. & Ganoe, C.H. (2001). MOOsburg: Multi-user domain support for a community network. *Internet Research*, 11(1), 65-73.
- Churcher, N.I., & Churcher, C.D. (1996). GroupARC – A collaborative approach to GIS. In *Proceedings of the 8th Annual Colloquium of the Spatial Information Research Centre*. New Zealand: University of Otago, 156-163.
- Churcher, N.I., & Churcher, C.D. (1999). Realtime conferencing in GIS. *Transactions in GIS*, 3(1), 23-30.
- Churchill, E.F., & Snowdon, D. (1998). Collaborative virtual environments: An introductory review of issues and systems. *Virtual Reality: Research, Development and Applications*, 3(1), 3-15.

- Dourish, P., & Bellotti, V. (1992). Awareness and coordination in shared workspaces. In proceedings of *ACM Computer Supported Cooperative Work*, 107-114.
- Dyck, J., & Gutwin, C. (2001). Where are you and what can you? Maintaining location awareness in collaborative 3D workspaces. (Tech. Rep. HCI-TR-2001-02). University of Saskatchewan.
- Dyck, J., & Gutwin, C. (2002). Groupspace: A 3D workspace supporting user awareness. In proceedings of *Conference on Human Factors in Computing Systems, Extended Abstracts*, 502-503.
- Egido, C. (1990). Teleconferencing as a technology to support cooperative work: Its possibilities and limitations. In Galegher, J., Kraut, R.E. & Egido, C. (Eds.), *Intellectual teamwork: Social and technological foundations of cooperative work*, (pp. 351-371). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Ellis, C., Gibbs, S.J., & Rein, G.L. (1991). Groupware: Some issues and experiences. *Communications of the ACM*, 34(1), 38-58.
- Evans, A., Kingston, R., & Turton, I. (1999). Web-based GIS to enhance public democratic involvement. *Geocomp99 Conference Proceedings*.
- Fraser, M., Benford, S., Hindmarsh, J., & Heath, C. (1999). Supporting awareness and interaction through collaborative virtual interfaces. In *Proceedings of the 12th Annual ACM Symposium on User Interface Software and Technology*, 27-36.
- Furnas, G. (1982). The fisheye view: A new look at structured files (Technical Memorandum 82-11221-22). Bell Laboratories. Murray Hill, New Jersey.
- Fussell, S.R., Kraut, R.E., & Siegel, J. (2000). Coordination of communication: Effects of shared visual context on collaborative work. In *Proceedings of Computer Supported Cooperative Work*, 21-30.
- Granum, E., Moeslund, T.B., Stoerring, M., Broll, W. & Wittkaemper, M. (2003). Facilitating the presence of users and 3D models by the augmented round table. In *Proceedings of The 6th Annual International Workshop on Presence*, Presence 2003, 19.
- Greenberg, S., Gutwin, C., & Cockburn, A. (1998). Awareness through fisheye views in Relaxed-WYSIWIS groupware. *Graphics Interface*, 28-38.
- Greenberg, S. & Roseman, M. (1996). GroupWeb: A WWW browser as real time groupware. In *Proceedings of Conference on Human Factors in Computing Systems, Companion Proceedings*, 271-272.

- Gutwin, C., & Greenberg, S. (1998). Design for individuals, design for groups: Tradeoffs between power and workspace awareness. In proceedings of *Computer Supported Cooperative Work*, 207-216.
- Gutwin, C., & Greenberg S. (1999). The effects of workspace awareness support on the usability of real-time distributed groupware. *ACM Transactions on Computer-Human Interaction*, 6(3), 243-281.
- Gutwin, C., & Greenberg, S. (2002). A descriptive framework of workspace awareness for real-time groupware. *Computer Supported Cooperative Work, The Journal of Collaborative Computing*, 11(3-4), 411-446.
- Gutwin, C., Roseman, M., & Greenberg S. (1996). A usability study of awareness widgets in a shared workspace groupware system. In proceedings of *Computer Supported Cooperative Work*, 258-257.
- Guzdial, M., Rick, J. & Kerimbaev, B. (2000). Recognizing and supporting roles in CSCW. In *Proceedings of the Conference on Computer Supported Cooperative Work*, 261-268.
- Hayne, S., Pendergast, M., & Greenberg, S. (1994). Implementing gesturing with cursors in group support systems. *Journal of Management Information Systems*, 10(3), 43-61.
- Hindmarsh, J., Fraser, M., Heath, C., Benford, S., & Greenhalgh, C. (2000). Object-focused interaction in collaborative virtual environments. *ACM Transactions on Computer-Human Interaction*, 7(4), 477-509.
- Isenhour, P.L., Rosson, M.B. & Carroll, J.M. (2001). Supporting interactive collaboration on the Web with CORK. *Interacting with Computers (13)*, Special Issue on "Interfaces for the Active Web," 655-676.
- Jankowski, P., Nyerges, T.L., Smith, A., Moore, T.J., & Horvath, E. (1997). Spatial group choice: A SDSS tool for collaborative spatial decision making. *International Journal of Geographical Information Science*, 11(6), 577-602.
- Keahey, T.A. (1997). Nonlinear magnification. PhD Dissertation. Indiana University. Retrieved from <http://www.cs.indiana.edu/~tkeahey/research/papers/thesis.html>
- Keahey, T.A., & Roberson, E.L. (1996). Techniques for non-linear magnification transformations. In *Proceedings of the IEEE Symposium on Information Visualization, IEEE Visualization*, 38-45.
- Kessler, G., Bowman, D., & Hodges, L. (2000). The simple virtual environment library: An extensible framework for building VE applications. *Presence: Teleoperators and Virtual Environments*, 9(2), 187-208.

- Lamping, J., Rao, R., & Pirolli, P. (1995). A focus+context technique based on hyperbolic geometry for visualizing large hierarchies. *Human factors in Computing Systems*, 401-408.
- Lau, L.M.S., Curson, J., Drew, R., Drew, P.M., & Leigh, C. (1999). Use of virtual science park resource rooms to support group work in a learning environment. In *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work*, 209–218.
- Leigh, J. & Johnson, A.E. (1996). Supporting transcontinental collaborative work in persistent virtual environments. *IEEE Computer Graphics and Applications*, 16(4), 47-51.
- Leigh, J., Johnson, A.E. & DeFanti, T.A. (1996). CALVIN: An Immersimedia design environment utilizing heterogeneous perspectives. In *Proceedings of IEEE International Conference on Multimedia Computing and Systems*, 20-23.
- Leigh, J., Johnson, A.E., Vasilakis, C.A. & DeFanti, T.A. (1996). Multiple-perspective collaborative design in persistent networked virtual environments. In *Proceedings of IEEE Virtual Reality Annual International Symposium*, 253-260.
- Leland, M.D.P., Fish, R.S., & Kraut, R.E. (1988). Collaborative document production using Quilt. In *Proceedings of the Conference on Computer Supported Cooperative Work*, 206-215.
- Leung, Y.K., & Apperley, M.D. (1994). A review and taxonomy of distortion-oriented presentation techniques. *ACM Transactions on Computer-Human Interaction* 1(2), 126-160.
- Lynch, K. (1960). *The image of the city*. Massachusetts: MIT Press.
- MacEachren, A.M., Brewer, I. & Steiner, E. (2001). Geovisualization to mediate collaborative work: Tools to support different-place knowledge construction and decision-making. In *Proceedings of 20th International Cartographic Conference*.
- Mastaglio, T.W. & Williamson, J. (1995). User-centered development of a large-scale complex networked virtual environment. In *Proceedings of the Conference on Human Factors and Computing Systems*, 546-552.
- McCarthy, J. F., Meidel, E.S. (1999). ActiveMap: A visualization tool for location awareness to support informal interactions. In proceedings of *Handheld and Ubiquitous Computing*, 158-170.
- Nyerges, Timothy L. & Jankowski, P. (1997). Enhanced adaptive structuration theory: A theory of GIS-supported collaborative decision making. *Geographical Systems*, 4(3), 225-259.
- Palen, L. (1999). Social, individual and technological issues for groupware calendar systems. In *Proceedings of Conference on Human Factors in Computing Systems*, 17-24.

- Pekkola, S. (2002). Critical approach to 3D virtual realities for group work. In *Proceedings of the Second Nordic Conference on Human-Computer Interaction*, 129-138.
- Piolat, A., Roussey, J., & Thunin, O. (1997). Effects of screen presentation on text reading and revising. *International Journal of Human-Computer Studies* 47(4), 565-589.
- Poupyrev, I., Tan, D.S., Billingham, M., Kato, H., Regenbrecht, H., & Tetsutani, N. (2001). Tiles: a mixed reality authoring interface. In *Proceedings of Interact 2001*.
- Rinner, C., (1999). Argumentation maps – GIS-based discussion support for online planning. PhD Dissertation. University of Bonn. Retrieved from <http://www.gmd.de/publications/research/1999/022/>.
- Roussos, M., Johnson, A.E., Moher, T., Leigh, J., Vasilakis, C.A., & Barnes, C. (1999). Learning and building together in an immersive virtual world. *Presence*, 8(3), 27-263.
- Sarkar, M., & Brown, M.H. (1992). Graphical fisheye views of graphs. *Human Factors in Computing System*, 83-91.
- Schafer, W.A. (2001). *A classification of interactive map software* (Technical Report TR-01-26). Blacksburg, Virginia: Computer Science, Virginia Tech. Received from <http://eprints.cs.vt.edu:8000/archive/00000669/01/TechReport.pdf>
- Schafer, W., & Bowman, D. (2002). Enhancing radar views using fisheye techniques. In conference supplement of *Computer Supported Cooperative Work*, 147-148.
- Schafer, W.A., Bowman, D.A., & Carroll, J.M. (2002). Map-based navigation in a graphical MOO. *ACM Crossroads*, 9(1), 8-15.
- Schafer, W.A., & Bowman, D.A. (2003). A comparison of traditional and fisheye radar view techniques for spatial collaboration. In proceedings of *Graphics Interface*, 39-46.
- Schafer, W.A., & Bowman, D.A. (2004). Evaluating the effects of frame of reference on spatial collaboration using desktop collaborative virtual environments. To appear in *Virtual Reality*.
- Schwarz, E., Beldie, I.P., & Pastoor, S. (1983). A comparison of paging and scrolling for hanging screen contents by inexperienced users. *Human Factors* 25(3), 279-292.
- Smith, R. B., O'Shea, T., O'Malley, C., Scanlon, E., & Taylor, J. (1989). Preliminary experiments with a distributed, multi-media, problem solving environment. *First European Conference on Computer Supported Cooperative Work*, 19-34.
- Stefik, M., Bobrow, D.G., Foster, G., Lanning, S., & Tatar, D., (1987). WYSIWIS revised: Early Experiences with Multiuser Interfaces. *ACM Transactions on Office Information System*, 5(2), 147-167.

- Stoakley, R., Conway, M.J. & Pausch, R. (1995). Virtual reality on a WIM: Interactive worlds in miniature. In *Proceedings of Human Factors and Computer Systems*, 265-272.
- Tang, J.C. (1991). Findings from observational studies of collaborative work. *International Journal on Man-Machine Studies*, 34(2), 143-160.
- Thorndyke, P., & Hayes-Roth, B. (1982). Differences in spatial knowledge obtained from maps and navigation. *Cognitive Psychology*, 14, 560-589.
- Tolsby, H. (1993). Navigating in a process landscape. In proceedings of *East-West International Conference on Human-Computer Interaction*, 141-151.
- Wickens, C.D., & Prevett, T.T. (1995). Exploring the dimensions of egocentricity in aircraft navigation displays. *Journal of Experimental Psychology: Applied*, 1(2), 110-135.
- Wickens, C.D. (1999). Frames of reference for navigation. In Gopher, D. & Koriat, A. (Eds.), *Attention and Performance XVII* (pp. 113-144). Cambridge, MA: MIT Press.
- Yang, H. (2002). Multiple perspectives for collaborative navigation in CVE. In *Extended Abstracts for the ACM Conference on Human Factors in Computer Systems*, 560-561.
- Yang, H., & Olson, G.M. (2002). Exploring collaborative navigation: The effect of perspectives on group performance. In *Proceedings of The 4th International Conference on Collaborative Virtual Environments*, 135-142.

Appendix A Map software applications surveyed

Name	Location	Navigation Technique	Support for Collaboration	Data Sources Available
ActiveMap (McCarthy & Meidel, 1999)	Andersen Consulting	continuous panning and discrete zooming	synchronous using an external resource	one data source
The map serves as an at-a-glance awareness tool that shows where people are within an office space.				
ArcData Online	www.esri.com/data/online/	discrete zooming and discrete panning	no support	non-interactive
A web application that allows users to create maps using data for sale at ArcData Online				
ArcView GIS 3.2	ESRI product	discrete zooming and discrete panning	no support	layers
A GIS produce that provides data visualization, querying, analysis, editing, and integration capabilities.				
Autodesk MapGuide LiteView Extension	www.mapguide.com	discrete zooming and discrete panning	no support	layers
An on-line application that serves maps without having users download and install a special viewer.				
Autodesk MapGuide Viewer	www.mapguide.com	discrete zooming and discrete panning	asynchronous using the map interface	layers
Software that delivers custom-specific maps and design data easily across the Web.				
Bay Area Rapid Transit (BART) Schedule Animation	www-itg.lbl.gov/vbart/homepage.html	no support	no support	multiple data sources viewed separately
A simulation that provides an infrastructure visualization of San Francisco's transit system.				
City of Tucson ArcIMS demo	www.ci.tucson.az.us/ed/ed.htm	discrete zooming and discrete panning	no support	layers
An on-line application to find and compare vacant commercial properties in Tucson, Arizona.				

Collaborative Map Annotator (Jedrysik, Moore, Stedman, & Sw	www.rl.af.mil/tech/programs/ADII/adii_cma.html	no support	synchronous using the map interface	layers
---	--	------------	-------------------------------------	--------

A demonstration application for the Interactive Data Wall research project.

EPIC Web Browser	www.epic.noaa.gov/epic/ewb	discrete zooming	no support	multiple data sources viewed separately
------------------	--	------------------	------------	---

An on-line tool to provide access to EPIC hydrographic data sets.

EtakGuide Map	www.etakguide.com	discrete zooming and discrete panning	no support	non-interactive
---------------	--	---------------------------------------	------------	-----------------

A demonstration of geocoding, a procedure to find and display any address on a map.

Geophysical Fluid Dynamics Lab Data	www.cdc.noaa.gov/cgi-bin/DataMenus.pl?dataset=gfdl	no support	no support	multiple data sources viewed separately
-------------------------------------	--	------------	------------	---

An on-line tool to visualize global climate variable data from the Geophysical Fluid Dynamics Lab.

GeoMedia Web Map	www.intergraph.com/gis/gmwm	discrete zooming and discrete panning	no support	layers
------------------	--	---------------------------------------	------------	--------

Commercial software produce that places GIS data on the Web.

GRASSLinks	www.regis.berkeley.edu/gldev/regis.html	discrete zooming and discrete panning	no support	layers
------------	--	---------------------------------------	------------	--------

A Web interface to a GIS system that facilitates data sharing between environmental planning agencies, public action groups, citizens, and private entities.

Hourly/Daily Rain Data	precip.fsl.noaa.gov/hourly_precip.html	discrete zooming and continuous panning	no support	layers
------------------------	--	---	------------	--------

An on-line tool to display daily and hourly precipitation totals from various locations all over the US.

Houston Real-Time Traffic Map	traffic.tamu.edu/incmap/incmap.asp	no support	no support	multiple data sources viewed separately
-------------------------------	--	------------	------------	---

A web-based application to convey real-time traffic conditions in Houston, Texas.

Interactive California Environmental Management, Assessment, and Planning System (ICEMAPS)	icemaps.des.ucdavis.edu/icemaps2/ICEMapInit.html	discrete zooming and discrete panning	no support	layers
--	--	---------------------------------------	------------	--------

A Web interface to an environmental, California GIS.

Interactive Web-Mapping Tool	atlas.geo.cornell.edu/ima.html	discrete zooming and discrete panning	no support	layers
------------------------------	--------------------------------	---------------------------------------	------------	--------

Cornell's Geoscience Information System tool on the Web.

Image Web Server	www.earthetc.com/ecwcounty/ecw_county_frame.asp	discrete zooming, continuous zooming, and continuous panning	no support	one data source
------------------	---	--	------------	-----------------

A demonstration of Image Web Server technology – county assessor data is provided on-line.

JShape Software	www.jshape.com/index0.html	discrete zooming and discrete panning	no support	layers
-----------------	----------------------------	---------------------------------------	------------	--------

A Java software package that supports fast and easy creation of interactive, web-based, GIS applications.

Macromedia Shockwave FreeHand maps	www.stmaartenstmartin.com/freehand	discrete zooming and discrete panning	no support	one data source
------------------------------------	------------------------------------	---------------------------------------	------------	-----------------

This software displays a FreeHand file on the web.

MAPBLAST!	www.mapblast.com	discrete zooming and discrete panning	asynchronous using an external resource	layers
-----------	------------------	---------------------------------------	---	--------

An on-line application to drawn maps of anywhere in the US and provide routing information.

Map-It	crust.er.usgs.gov:80/mapit	no support	no support	one data source
--------	----------------------------	------------	------------	-----------------

A demonstration of the free, open source software package GMT. This on-line tool draws a map of user-supplied latitude and longitude pairs.

MapQuest	www.mapquest.com	discrete zooming and discrete panning	asynchronous using an external resource	layers
----------	------------------	---------------------------------------	---	--------

An on-line application to draw maps of anywhere in the US and provide routing information.

Maps.com	www.maps.com/learn	discrete zooming and continuous panning	no support	layers
----------	--------------------	---	------------	--------

Education resources to make learning fun. Maps are animated and interactive.

Maps On Us	www.mapsonus.com	discrete zooming and discrete panning	asynchronous using an external resource and synchronous using the map interface	layers
------------	------------------	---------------------------------------	---	--------

An on-line application to draw maps of cities and yellow page listings in the US and to provide routing information.

MOOsburg map widget	moosburg.cs.vt.edu	continuous zooming and continuous panning	asynchronous and synchronous using an external resource	non-interactive
---------------------	--------------------	---	---	-----------------

A navigation tool, awareness tool, and end-user developer tool for MOOsburg.

MyWay.com's javamap	www.zip2.com	discrete zooming and discrete panning	asynchronous using an external resource	non-interactive
---------------------	--------------	---------------------------------------	---	-----------------

An on-line application to draw maps of anywhere in the US.

NOAA Real-time TAO Buoy Data Display	www.pmel.noaa.gov/toga-tao/realtime.html	no support	no support	multiple data sources viewed separately
--------------------------------------	--	------------	------------	---

An on-line tool to display real-time data from ocean buoys in order to analyze El Niño and La Niña.

Pacific Century Systems Limited's GIS Map	lps.pcgsys.com/wml/elpas/html	discrete zooming and discrete panning	no support	layers
---	-------------------------------	---------------------------------------	------------	--------

A password-protected website to track co-workers' locations at PCS.

Portland State University Child Welfare Partnership Thematic Mapping Program	www.ncn.com/~rilg/Mapper/ore.htm	no support	no support	multiple data sources viewed separately
--	--	------------	------------	---

A web-based tool that provides child welfare data for the state of Oregon.

Puget Sound Traffic Conditions	www.wsdot.wa.gov/PugetSoundTraffic	multiple maps	no support	one data source
--------------------------------	--	---------------	------------	-----------------

A webpage that displays the current traffic conditions for the Puget Sound area in Washington state.

Race Track Locator	chasinracin.net/track-locator	multiple maps	no support	non-interactive
--------------------	--	---------------	------------	-----------------

A website to locate race tracks around the US.

Ramonamap (Bartlett, 1994)	Western Research Lab	no support	asynchronous using the map interface	one data source
----------------------------	----------------------	------------	--------------------------------------	-----------------

An application to share office information spatially.

Rand McNally TripMaker Deluxe 1999 Edition	Rand McNally	discrete zooming and discrete panning	no support	layers
--	--------------	---------------------------------------	------------	--------

A commercial software product to plan road travel.

Town of Blacksburg WebGIS	www.webgis.net/blacksburg	discrete zooming and discrete panning	no support	layers
---------------------------	--	---------------------------------------	------------	--------

An on-line tool that provides information regarding the town's street system, properties, parks, bus stop locations, zoning, etc.

Taiwan Map	peacock.tnrc.edu.tw/ADD/maps/taiwanmap.html	multiple maps	no support	non-interactive
------------	--	---------------	------------	-----------------

A primitive interactive map website.

Triscape Map Explorer	www.triscape.com	continuous zooming and continuous panning	asynchronous using the map interface	non-interactive
-----------------------	------------------	---	--------------------------------------	-----------------

An on-line application that provides an alternative approach to on-line mapping.

Virtual Slaithwaite (Evans, Kingston, & Turton, 1999)	www.ccg.leeds.ac.uk/slaithwaite	discrete zooming and discrete panning	asynchronous using the map interface	one data source
---	---------------------------------	---------------------------------------	--------------------------------------	-----------------

A demonstration of a local, environmental decision-making, public participation tool.

Virtual Tourist	www.virtualtourist.com	multiple maps	asynchronous and synchronous using an external resource	only one data source
-----------------	------------------------	---------------	---	----------------------

A web-based application that allows people from all over the world to share travel plans.

Appendix B Map-based Navigation Prototype Evaluation

A formative evaluation of the map-based navigation tool described in Section 3.2 was conducted to explore designing an interactive map of a familiar place. The sessions began with a background questionnaire and involved three sections that required working with paper maps and the prototype. The researcher used an evaluation data sheet to record observations for each participant.

B.1 Background Questionnaire

Demographic Information

Name: _____

Age: _____

Gender: male female

Occupation: student faculty staff other: _____

Computer Experience

How many years have you used a computer?

0-2 years 2-4 years 4-6 years 6+ years

How often do you use a computer?

many times a day once a day a couple times a week less frequently more frequently

What computer applications do you regularly use? (Circle all that apply)

Word Processors Spreadsheets Internet Browsers Email

Others: _____

Do you use map programs on the Internet? Yes No

If yes, how often do use them?

twice a month once a month once every 6 months once a year less frequently

Circle all the online map programs you have used:

Mapquest

MapBlast!

Micosoft Expedia

MapsOnUs

Rand McNally

Yahoo! Maps

Lycos RoadMaps

Others: _____

Do you use any other type of map software programs? Yes No

If yes, how often do you use them?

twice a month once a month once every 6 months once a year less frequently

Circle all the map software you have used:

Rand McNally TripMaker
Microsoft AutoMap Streets Plus

Rand McNally StreetFinder
Microsoft Trip Planner

Others: _____

Map Activities

Do you own an atlas? Yes No

If yes, how often to do you use it?

twice a month once a month once every 6 months once a year less frequently

Do you own other maps? Yes No

If yes, how often to do you use them?

twice a month once a month once every 6 months once a year less frequently

B.2 Evaluation Data Sheet

Evaluation Checklist

Group: 1 2 3 4

Prompt: I II

Section 1: Think aloud testing with paper prototypes

Flat Map Test with just roads:

Prompt: Do you recognize this map? What is it a map of?

Follow up: How did you know that?

Could the user say that it was Blacksburg: yes, quickly yes, after some time no, not really no, not at all

Landmarks mentioned: Drillfield Main Street 460 Prices Fork Rd Tom's Creek Rd Patrick Henry Dr

Other comments:

Prompt: Knowing that it this a map of Blacksburg, can you point out the approximate location of where we are now?

Follow up: How did you know that?

User could point out current location: yes, quickly yes, after some time no, not really no, not at all

Knowledge used: route survey landmark

Rotate map: yes no

Draw on map: yes no, but could have no

Refer to road names: yes no

Landmarks mentioned: Drillfield Main Street 460 Prices Fork Rd Tom's Creek Rd Patrick Henry Dr

Particular problems:

Other ways to support navigation:

Flat Map Test with roads and road labels:

Prompt: Can you point out where your grocery store is located?

Follow up: How did you know that?

Did the road names help you in your task? yes no

Refer to road names: yes no

Knowledge used: route survey landmark

Rotate map: yes no

Draw on map: yes no, but could have no

Landmarks mentioned: Drillfield Main Street 460 Prices Fork Rd Tom's Creek Rd Patrick Henry Dr

Particular problems:

Other ways to support navigation:

Flat Map Test with roads and a collection of labeled landmarks:

Prompt: Can you point where your bank is located?

Follow up: How did you know that?

Did the landmarks help you in your task? yes no

Are you familiar with all of the landmarks on this map? yes no

Do you consider these to be major landmarks in Blacksburg? yes no

Are there others that should be included on the map?

User referenced labeled landmarks: yes no

Knowledge used: route survey landmark

Rotate map: yes no

Draw on map: yes no, but could have no

Landmarks mentioned: Drillfield Main Street 460 Prices Fork Rd Tom's Creek Rd Patrick Henry Dr
 Airport VPI Golf Course Town Golf Course BMS BHS Kipps Elementary
 Duck Pond Linkous Elementary Harding Avenue Elementary Lane Stadium
 Margaret Beeks Elementary

Particular problems:

Other ways to support navigation:

Section follow-up questions to learn more about the user population:

How long have you lived in Blacksburg?

When have you used a map of Blacksburg before? upon first moving looking for a place or address finding best route

In cases where you need to use a map, do you usually read the map or do you generally let someone else do it?
I read it Someone else reads it

Section 2: Think aloud testing with the flat view of the mapwidget

Flat Map Test with roads continuously displayed and buildings displayed at certain zoom levels:

Prompt: This is a new type of map program I'm developing. I would like you to experiment with it and see if you center your home.

Follow up: Did you figure out that you could zoom, click, and drag the map? yes, all 3 yes, 2 yes, 1 none

Interactions used: zoom click drag

Mapwidget responded as the user expected: definitely probably not really user was confused

Was the user ever baffled: yes no

Other user reaction comments:

Flat Map Test with roads continuously displayed and buildings displayed at certain zoom levels:

Prompt: From your home, now I would like you to center the downtown Main Street area.

Follow up: Have you noticed that buildings appear once you zoom in? yes no

Did you find this confusing? yes no

Have you figured out that clicking centers the map to the place where you clicked? yes no

Did you find this confusing? yes no

And, have you learned that clicking and dragging drags the whole map? yes no

Did you find this confusing? yes no

Was the user ever baffled: yes no

Mapwidget responded as the user expected: definitely probably not really user was confused

Interactions used: zoom click drag

Other user reaction comments:

Section 3: Think-aloud testing using various views of the mapwidget

Paper prototype test that shows views similar to what would be seen on the screen:

Prompt: I've been experimenting with different ways to display maps.

Can you point out the intersection of Prices Fork Rd and Main Street on this map?

Can you point out the post office?

Can you point out the library on this one?

Where is the Blacksburg Middle School on this one?

Follow up: Which view did you find the easiest to understand?

A C B D
Parabolic x^2 ArcTan Hyperbolic Parabolic $x^{1.5}$

Which one was the most confusing? Parabolic x^2 ArcTan Hyperbolic Parabolic $x^{1.5}$

Did the user struggle with one view more? Parabolic x^2 ArcTan Hyperbolic Parabolic $x^{1.5}$

Is one view more intuitive? Parabolic x^2 ArcTan Hyperbolic Parabolic $x^{1.5}$

Did the user express a view preference? Parabolic x^2 ArcTan Hyperbolic Parabolic $x^{1.5}$

Interactive tests with the mapwidget:

Prompt: Given that you are currently located at the center of this flat map, which side of Bypass 460 are you on? Are you closer to town or further away from town? Now, given this fisheye view with the same zoom, which side are you on?

Users should be broken up into four groups so that we can test the four different fisheye views.

Group 1 - Parabolic x^2 - A

Group 2 - Hyperbolic - B

Group 3 - Parabolic $x^{1.5}$ - D

Group 4 - ArcTan - C

Response for flat: correct incorrect

Response for fisheye view: correct incorrect

Did the user struggle to understand the fisheye view: yes, definitely yes, a little no, not really no, not at all

Did the user express a view preference? Flat Fisheye

Other comments:

Prompt I: Can you find the post office with this Arctan/Hyperbolic view? Now can you find it again using this Parabolic view?

Prompt II: Can you find the post office with this Parabolic view? Now can you find it again using this ArcTan/Hyperbolic view?

Follow up: Which view was easier to work with? ArcTan/Hyperbolic Parabolic

- Group 1 - Parabolic $x^{1.5}$, ArcTan – D, C
- Group 2 - ArcTan, Parabolic x^2 – C, A
- Group 3 - Parabolic x^2 , Hyperbolic – A, B
- Group 4 - Hyperbolic, Parabolic $x^{1.5}$ – B, D

Time to navigate with Parabolic view:

Zoom level with Parabolic:

Did the user take advantage of the greater display area with the Parabolic view: yes no

Time to navigate with ArcTan/Hyperbolic view:

Zoom level with ArcTan/Hyperbolic view:

Did the user struggle to understand either fisheye view: yes, definitely yes, a little no, not really no, not at all

Which fisheye view did the user struggle with more: Parabolic ArcTan/Hyperbolic

Other comments:

Prompt I: Using this flat view, can you find the library? Now, I'd like you to find the library with this fisheye view.

Prompt II: Using this fisheye view, can you find the library? Now, I'd like you to find the library with this flat map.

Follow up: Which view was easier to work with? Flat Fisheye

- Group 1 - Parabolic x^2 - A
- Group 2 – Hyperbolic - B
- Group 3 - Parabolic $x^{1.5}$ - D
- Group 4 – ArcTan - C

Time to navigate with Flat view:

Zoom level with Flat view:

Time to navigate with fisheye view:

Zoom level with fisheye view:

Did the user take advantage of the greater display area with the fisheye view: yes no

Did the user struggle to understand the fisheye view: yes, definitely yes, a little no, not really no, not at all

Other comments:

Prompt I: Using this flat view and not changing the zoom, can you find the intersection of Prices Fork Rd and Main Street? Now, I'd like you to find that intersection with this fisheye view, again not changing the zoom.

Prompt II: Using this fisheye view and not changing the zoom, can you find the intersection of Prices Fork Rd and Main Street? Now, I'd like you to find that intersection with this flat map, again not changing the zoom.

Follow up: Which view was easier to work with? Flat Fisheye

- Group 1 – Hyperbolic - B
- Group 2 - Parabolic $x^{1.5}$ - D
- Group 3 – ArcTan - C
- Group 4 - Parabolic x^2 - A

Time to navigate with Flat view:

Zoom level with Flat view:

Time to navigate with fisheye view:

Zoom level with fisheye view:

Did the user take advantage of the greater display area with the fisheye view: yes no

Did the user struggle to understand the fisheye view: yes, definitely yes, a little no, not really no, not at all

Other comments:

Section follow-up questions to learn more the user's preference:

Which view did you find the easiest to navigate with?

A C B D
Parabolic x^2 ArcTan Hyperbolic Parabolic $x^{1.5}$ Flat

If given another task, which view would you choose to use?

Parabolic x^2 ArcTan Hyperbolic Parabolic $x^{1.5}$ Flat

Why?

Did you notice any advantages with using the fisheye views?

Appendix C Spatial Discussion Prototype Evaluation

An initial evaluation of the spatial discussion prototype described in Section 3.3 was conducted with two members of the Stroubles Creek Focus Group. They were asked to perform a realistic distributed, spatial collaboration task and answer a short questionnaire. Their typed chat session was recorded.

C.1 Evaluation Task

The Virginia Museum of Natural History at Virginia Tech is holding a workshop to kick off the start of the new school year. The event will occur in the fall on a Saturday when there is not a Tech football game and will last for a half day. The goal of the workshop is to educate a variety of people about Stroubles Creek. Blacksburg residents, some grade school children, and college students thinking about majoring in Environmental Science have all signed up. To make the event engaging the entire time will be spent in the field.

You are to work with a partner to help the Museum decide where to visit. You should identify and agree on four specific locations that enable the workshop to give an overview about the creek. These locations could include different land uses, historical locations, unique features, problematic points, etc. Assume that the people know very little about the creek, if anything at all.

This task will be completed in MOOsburg. The collaborative map tool will be your primary means of communication. It allows you to share typed text and drawings on the map

C.2 Questionnaire

Follow-up Questions for Collaborative Map Tool Evaluation

Think about how you interacted with the tool:

Was the tool easy to learn and use? Explain.

What did you like about the user interface?

What did you not like about the interface?

Other suggestions for improvement?

Think about your overall experience:

Did the session go as you had expected? What was different?

What aspects of the session went well? What went wrong?

Were you able to communicate your ideas effectively? Describe.

Think about using this tool in the future:

What other situations might benefit from this tool?

What might have to be different in order to support these situations?

C.3 Session Log

User	Timestamp	Link	Text
lin	Tue Dec 04 13:02:44 EST 2001		hi
tamim	Tue Dec 04 12:56:45 EST 2001		hi
tamim	Tue Dec 04 13:03:53 EST 2001		Hi Lynn, Good to see you gain!
lin	Tue Dec 04 13:12:23 EST 2001		Hi Tamim-- Ever notice how everybody spells my name different? We are spozed to pick 4 sites for a Stroubles Field Day--can you help me?
tamim	Tue Dec 04 13:05:11 EST 2001		Lynn, We need to select four sites. Each should have unique characteristics.
lin	Tue Dec 04 13:14:05 EST 2001	yes	ignore this line
tamim	Tue Dec 04 13:08:21 EST 2001		One site in urban area One site in agricultural area One site in forested area One site where some type of maximum impact is speculated What do you think?
lin	Tue Dec 04 13:15:43 EST 2001	yes	One site we need to go to is just below the duck pond where people can see lots of impacts in a short distance and we can document with past data--Llyn
lin	Tue Dec 04 13:16:20 EST 2001		Right! That's what's so good about this little watershed. L~~
lin	Tue Dec 04 13:19:29 EST 2001		Where is an ag site?
tamim	Tue Dec 04 13:12:06 EST 2001		Llyn, What does your blue line indicate?
lin	Tue Dec 04 13:20:28 EST 2001		cyan line is supposed to be pointing to the site downstream of duckpond, but it didn't lie up?

tamim	Tue Dec 04 13:16:28 EST 2001	yes	The ag site is on plantation road about a mile below the Duck Pond.
tamim	Tue Dec 04 13:18:01 EST 2001		For the urban site we can use one of SOS sites in downtown area. Do you have any suggestions?
lin	Tue Dec 04 13:26:26 EST 2001		How is that ag? The site after you go under 460 might be more clearly ag, unless I don't know what the landuses are--seems like that area is still mostly urban--commuter parking, etc.
lin	Tue Dec 04 13:28:23 EST 2001		The urban site could be that one downstream of DP? OR we could go to the one by the fire station?
tamim	Tue Dec 04 13:21:15 EST 2001		By the plantation Rd. there are VT barns, sheep and cattle fields as you walk toward Foxridge. That area must have been influenced by animal waste.
tamim	Tue Dec 04 13:21:42 EST 2001		Llyn, What is DP?
lin	Tue Dec 04 13:29:25 EST 2001		duck pond
tamim	Tue Dec 04 13:23:00 EST 2001		The urban site should be, in my opinion, above the DP. The FS site can be appropriate.
lin	Tue Dec 04 13:32:16 EST 2001		Maybe the DuckPond site can be the "maximum impact" site--it's nice because it is so accessible.
tamim	Tue Dec 04 13:25:27 EST 2001		For forested site we need to look up Ray's database. You're right. The site could be near confluence of Slate Branch.
tamim	Tue Dec 04 13:27:00 EST 2001		About the DP site, do you suggest a site immediately below the pond? how about the bridge by the computer parking lot below the Duck Pond?
lin	Tue Dec 04 13:35:33 EST 2001	yes	This could be the urban site by the fire station it is all gabion basketed
tamim	Tue Dec 04 13:28:44 EST 2001	yes	Llyn, See attached for DP site.
tamim	Tue Dec 04 13:29:59 EST 2001		Llyn, I agree with the urban site by the fire station. So site one is selected.
tamim	Tue Dec 04 13:34:01 EST 2001	yes	Llyn, For max ag. impact we can do sampling at the bridge by Foxridge.
lin	Tue Dec 04 13:44:46 EST 2001	yes	I think you're a little further downstream than foxridge, I think it's about here
tamim	Tue Dec 04 13:38:20 EST 2001		I guess you're right. Let's go for a visit sometime.
lin	Tue Dec 04 13:46:19 EST 2001		It's a beautiful day, let's ditch these computers and get out in the field!
lin	Tue Dec 04 13:47:50 EST 2001		Let's use your downstream site as the forested site, but we need to be upstream of the sewer treatment plant'

tamim	Tue Dec 04 13:41:01 EST 2001		Lly, Then perhaps site 2 should be immediately below DP and site 3 by the Foxridge bridge. I think we need more info for site 4 (forested area). What do you think?
lin	Tue Dec 04 13:49:30 EST 2001		
tamim	Tue Dec 04 13:43:14 EST 2001	yes	
lin	Tue Dec 04 13:51:01 EST 2001		
lin	Tue Dec 04 13:52:52 EST 2001		Urban site is at fire station Impacted site is below duck pond Ag site is a t Foxige Forest site is at edge of mmap for want of mor map
tamim	Tue Dec 04 13:48:01 EST 2001		Ok. Let's agree about 3 sites and then follow up later about site 4.
tamim	Tue Dec 04 13:49:10 EST 2001		How about coffee?
lin	Tue Dec 04 13:56:46 EST 2001		OK we'll need to get some more stream on the map, and maybe go out in the field for the forest stie, see you in a minit--L

Appendix D Two-Dimensional Map Study Materials

For the two-dimensional map study described in Chapter 4, each participant filled out a background questionnaire and reviewed a set of instructions beforehand. After each map task participants completed a rating questionnaire and, at the end, they completed a final comparison questionnaire. The researcher used an evaluation data sheet to record observations for each collaboration session.

D.1 Background Questionnaire

Demographics

- a. Age:
- b. Gender: male female
- c. Occupation:
- d. Length of time in Blacksburg, VA:
- e. Highest level of school you completed?
high school 2-year college/certificate program 4-year college program graduate program

Maine Experience

- a. Are you from Maine? yes no
- b. Circle all the cities in Maine you have visited: Belfast Biddeford Rockland Augusta

Map Experience

- a. How often do you use paper maps and atlases?
more than twice a month twice a month once a month every two months twice a year or less
- b. Think about the last time you used a paper map or atlas.
What were you doing? _____
How long ago was it? _____
- c. How often do you look up maps online (Mapquest, Yahoo! Maps, Microsoft Expedia)?
more than twice a month twice a month once a month every two months twice a year or less
- d. Think about the last time you used an online map.
What were you doing? _____
How long ago was it? _____

The purpose of this study is to evaluate different representation techniques for collaborative map software. We are interested in different ways to display maps and the awareness features that help people collaborate.

Procedure:

Your role in the study is to work with an assigned partner to complete four tasks. You and your partner will be sitting at separate computers and you are not allowed to look at the other's screen. You will be able to talk with one another.

The four tasks you will complete are similar. Each has the same objective, while the maps and the user interface are different. Before each task, you will have a chance to explore the upcoming interface using a map of Blacksburg. When you are comfortable with the interface we will give you the task and the map for the town in Maine. At the end of each task, you will be asked to rate the experience. When you have completed all four tasks, there will be one last questionnaire asking about your preference. The procedure for the study is:

- Explore interface #1 with a map of Blacksburg
- Use interface #1 with a map of Belfast, Maine to position traffic lights and road signs
- Rate interface #1

- Explore interface #2 with a map of Blacksburg
- Use interface #2 with a map of Biddeford, Maine to position traffic lights and road signs
- Rate interface #2

- Explore interface #3 with a map of Blacksburg
- Use interface #3 with a map of Rockland, Maine to position traffic lights and road signs
- Rate interface #3

- Explore interface #4 with a map of Blacksburg
- Use interface #4 with a map of Augusta, Maine to position traffic lights and road signs
- Rate interface #4

- Final questionnaire

Objective:

For the next couple of minutes, you and your partner are going to pretend to be Traffic Engineers. Your job is to work together and decide where to position new traffic lights and road signs in four towns. You will be working with maps from small towns in Maine.

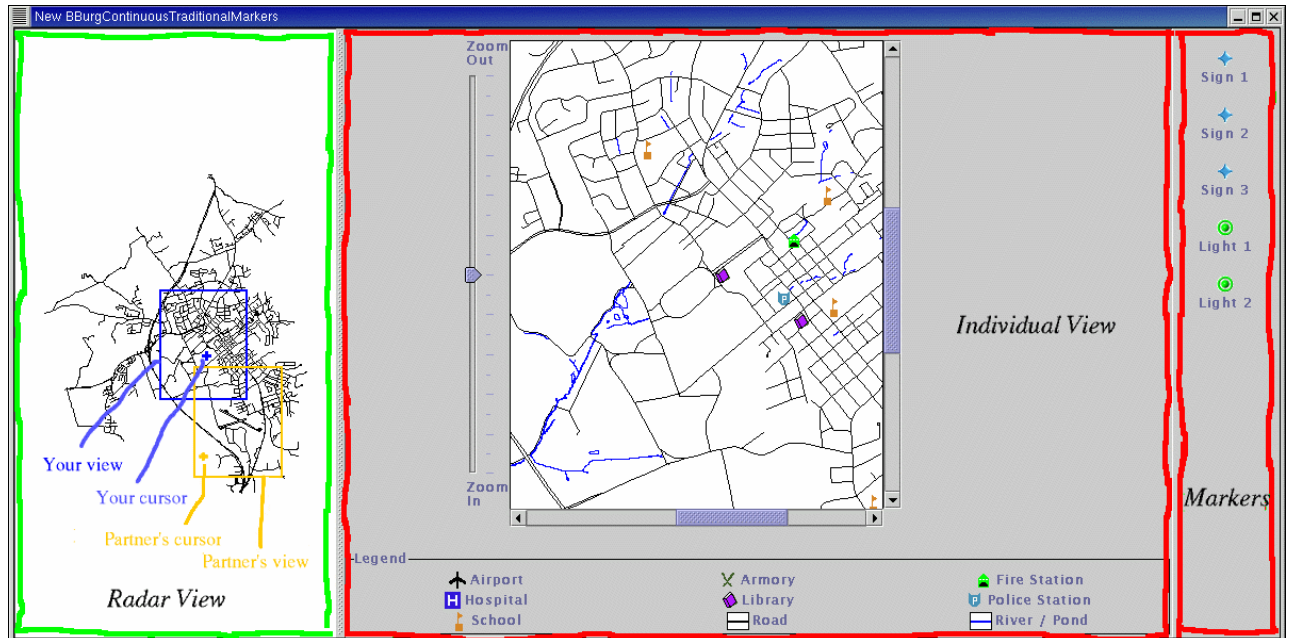
Each of you will have a different set of criteria that specifies where to place the lights and signs. **Both sets of the criteria need to be satisfied** in your solution so it is important that you work together. Both you and your partner will position these items on your individual computers.

For each town, you will be positioning three new road signs and two new traffic lights. The criterion for the road signs is open-ended. You and your partner will have to discuss where to position them, making sure you agree on the locations. The traffic lights, on the other hand, need to be placed at a specific intersection.

It is important that you position the items as **precise as possible** and that you and your partner mark the **exact same location**. We will be measuring the distance between you and your partner's markers. We will also measure the distance between each traffic light marker and the exact center of the intersection. One way to be precise and keep these differences to a minimum is to **zoom in** before marking a location.

Zooming in enlarges an area, making you more likely to mark the same location as your partner and be closer to the exact center of an intersection.

Basic Interface:



On the left-hand side there is a radar view that indicates what you are viewing and what your partner is viewing. The radar view also shows the locations of your cursor and your partner's cursor. This view will be the same on both monitors.

In the center of the display is your view of the town. You can navigate the map through the zoom slider on the left and the scrollbars to the right and below. An easy way to zoom the map is by clicking on the "Zoom In" and "Zoom Out" labels. Below the map is a legend that explains the symbols found on the map.

On the right-hand side are the markers you will position on the map. Your partner has his/her own set of markers that you cannot see. Both of you will position three sign markers and two light markers for each task. You can move a marker by clicking and dragging it. Markers that you drop onto the map will move with the map when you zoom or scroll.

1 2 3 4 5 6 7
much effort little effort

c. How hard did you have to concentrate to do the task?

1 2 3 4 5 6 7
very hard not hard

Section 2.

Think about your experience with the collaboration software in the most recent task.

a. Overall, how easily could you work with the other person?

1 2 3 4 5 6 7
not very easily very easily

b. In general, how easily could you understand your partner?

1 2 3 4 5 6 7
not very easily very easily

c. How easily could you tell what was displayed on your partner's screen?

1 2 3 4 5 6 7
not very easily very easily

d. How easily could you understand where your partner wanted to place a marker?

1 2 3 4 5 6 7
not very easily very easily

e. In general, how easily could you communicate your ideas?

1 2 3 4 5 6 7
not very easily very easily

f. How easily could you describe where you wanted to place a marker?

1 2 3 4 5 6 7
not very easily very easily

g. How well did your partner understand what was displayed on your screen?

1 2 3 4 5 6 7
he/she had a poor understanding he/she had a good understanding

h. How helpful was the radar view on the left-hand side of the interface for the task?

a. In general, how well were the two of you able to work together?

1 2 3 4 5 6 7
not very well very well

b. How easily could you agree on where to position the markers?

1 2 3 4 5 6 7
not very easily very easily

Section 5. (Not used with first task)

Think about the most recent task you completed in comparison to the task before.

a. How did the most recent interface differ in terms of how easily you could collaborate with your partner?

1 2 3 4 5 6 7
harder to collaborate easier to collaborate

b. How did the most recent interface differ in terms of how easily you could navigate your personal map?

1 2 3 4 5 6 7
more difficult to navigate easier to navigate

c. Which interface did you prefer?

most recent interface the interface before

D.4 Final Questionnaire

1. Think about the four trials you just completed. Using the four screenshots, rank the conditions based on how easy/hard it was to collaborate with a partner. (A-Belfast, B-Biddeford, C-Rockland, D-Augusta)

Easiest to collaborate with

1.

2.

3.

4.
Hardest to collaborate with

2. Think about your personal use of the map software. Using the four screenshots, rank the conditions based on how easy/hard it was to navigate the map. (A-Belfast, B-Biddeford, C-Rockland, D-Augusta)

Easiest to navigate with

1.

2.

3.

4.
Hardest to navigate with

3. Think about the conditions you liked the most and the least. Using the screenshots, rank the conditions based on which ones you preferred. (A-Belfast, B-Biddeford, C-Rockland, D-Augusta)

Most desirable

1.

2.

3.

4.
Least desirable

4. Two of the trials limited your navigation options (discrete navigation) while the other two offered more views of the map (continuous navigation).

- | | | | |
|--|----------|------------|----------------|
| a. Which was better for collaboration? | Discrete | Continuous | About the same |
| b. Which was easier to use? | Discrete | Continuous | About the same |
| c. Which did you prefer? | Discrete | Continuous | About the same |

5. Two of the trials had you use a traditional radar view, while the other two had a bubbled radar view.

- | | | | |
|--|------------------|-------------|----------------|
| a. Which was better for collaboration? | Traditional view | Bubble view | About the same |
|--|------------------|-------------|----------------|

- | | | | |
|-----------------------------|------------------|-------------|----------------|
| b. Which was easier to use? | Traditional view | Bubble view | About the same |
| c. Which did you prefer? | Traditional view | Bubble view | About the same |

6. How much effort did your partner put into the tasks, in comparison to your own efforts?

1	2	3	4	5	6	7
more effort						less effort

D.5 Evaluation Data Sheet

Trial 1

Treatment: Traditional/Bubble Discrete/Continuous

Town: Augusta/Belfast/Biddeford/Rockland

Completion Time:

a. How well were the subjects able to work together?

1	2	3	4	5	6	7
very compatible						not compatible at all

b. How different were the subjects' motivation levels?

1	2	3	4	5	6	7
equally motivated						extremely different levels

c. Explain any noticeable differences in the subjects' spatial abilities.

d. Explain any positive/negative critical incidents that occurred while using the radar view.

e. Explain any positive/negative critical incidents that occurred while navigating.

f. What basic strategy did the subjects take in solving the puzzle?

g. Describe any usability problems with:

Mouse operations -

h. Note any interesting/insightful/quotable comments made.

Trial 4

Treatment: Traditional/Bubble Discrete/Continuous

Town: Augusta/Belfast/Biddeford/Rockland

Completion Time:

a. How well were the subjects able to work together?

1 2 3 4 5 6 7
very compatible not compatible at all

b. How different were the subjects' motivation levels?

1 2 3 4 5 6 7
equally motivated extremely different levels

c. Explain any noticeable differences in the subjects' spatial abilities.

d. Explain any positive/negative critical incidents that occurred while using the radar view.

e. Explain any positive/negative critical incidents that occurred while navigating.

f. What basic strategy did the subjects take in solving the puzzle?

g. Describe any usability problems with:

- Mouse operations -
- Scrollbars or panning -
- Zooming -
- Moving the markers -
- Understanding the legend -
- Hearing one another -
- Slow computer response –

h. Note any interesting/insightful/quotable comments made.

Appendix E Three-Dimensional Virtual Environment Study Materials

For the three-dimensional virtual environment study described in Chapter 5, each participant filled out a background questionnaire and reviewed a set of instructions beforehand. After each task participants completed a rating questionnaire.

E.1 Background Questionnaire

Demographics

- f. Age:
- g. Gender: male female
- h. Occupation:
- i. Highest level of school you completed?
high school 2-year college/certificate program 4-year college program graduate program

Virtual Environments Experience (including 3D games)

Have you ever had an experience with virtual reality or a virtual environment?

Think about the last experience you had.

What were you doing? _____

How long ago was it? _____

What kind of display did you use? (desktop, CAVE, ImmersaDesk, etc.) _____

Do you use virtual environments on a regular basis?

How often do you use a virtual environment?

- more than twice a month
- twice a month
- once a month
- every two months
- twice a year
- less than twice a year

Spatial Abilities

In general, how would you rate your spatial abilities?

1 2 3 4 5 6 7

Have you ever experimented with a collaborative virtual environment? Yes No

If yes, which one? (DIVE, MASSIVE, multi-player DOOM, multi-player Quake) _____

E.2 Instructions

The procedure for this experiment involves three stages:

- You will be asked to fill out a background questionnaire first.
- Then, there will be a training session that lasts less than an hour. During this session you will practice working with a virtual environment.
- After the training, you will work with a partner to complete four tasks and four surveys. The session should also last less than an hour and will be concluded with an informal discussion.

The task for this experiment is to work with a partner to position furniture in four different houses. Each of you will have your own computer and you will be able to talk with another across the room. You will not be able to see what is on your partner's screen. There will be an avatar face in the world so you know your partner's position.

In 3D there are two different perspectives. Egocentric views show the world from first person point of view similar to a game like DOOM or Quake. Exocentric views have a bird's eye or top down view. During this experiment, you will try out different exocentric-egocentric combinations with your partner. Sometimes you will have the same perspective and other times they will differ.

There are also two roles in the experiment. The director knows what object to move and where to move it. The other person, the actor, performs the actual movement. During the experiment, you be the director and the actor twice.

We will be recording the time it takes to complete each task and the audio conversation you have with your partner. The task is complete when the object turns yellow **AND** it is unselected by pressing the "Done" button on the pop-up window.

Egocentric Navigation:

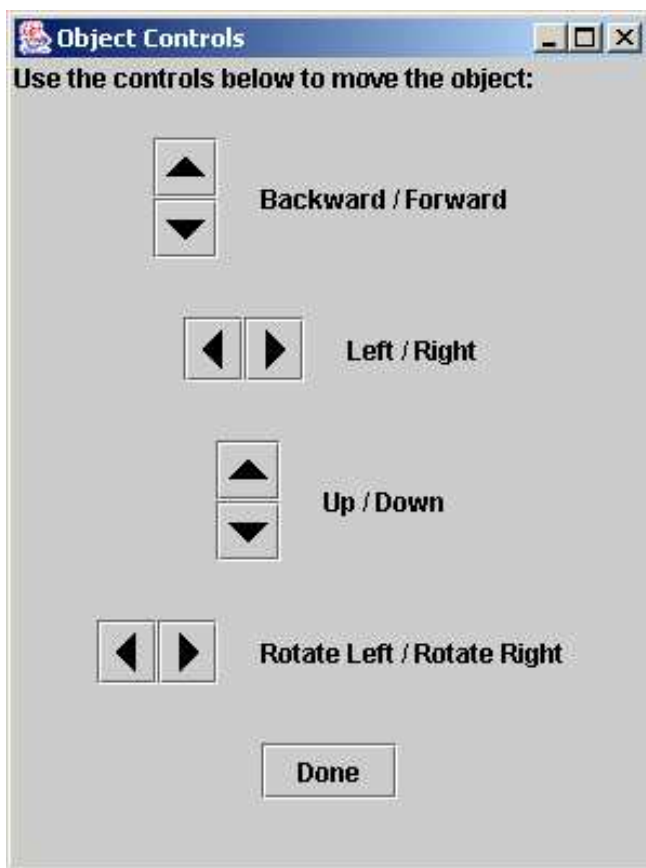
To move around in the egocentric perspective, use the arrow keys to move forward/backward and left/right. You can also look around by clicking and dragging the view left, right, up, and down. You can not change the height of your view as you will always see the world as a 6 foot person would, although you can walk through walls and other objects in the world.

Exocentric Navigation:

Navigation with the exocentric perspective is similar to moving around a half sphere that surrounds the virtual world. Use the left/right arrow keys to move around the world and the up/down arrow keys to change the viewing angle and height. You can not zoom in on the world as you will always see the world from outside.

Selection and Manipulation:

To select an object right click on it. After it is selected, a red box will be drawn around the object on your screen and your partner's screen. The object can be moved with the buttons in the pop-up window (see below). The object moves relative to the way you see the world. If you are looking up, moving the object backward will actually move it up and away from you. The object can be moved through walls and other objects in the world. Selecting the object does not prevent you from moving, continue to use the keyboard and mouse to navigate. When the object is in the correct position it will turn yellow on both screens. Clicking "Done" completes the task.



Pop-up window for moving an object.

E.3 Task Survey

Instructions: Please answer the following questions about the task you just completed. Each question asks you to use a rating scale from one to seven with four being a neutral response.

Section 1 – Collaboration Rating

Think about the task you just completed and your experience with the collaboration software.

Appendix F Real-World Application and Validation Materials

For the real world study described in Chapter 6, the researcher used an evaluation data sheet to record observations during the collaboration sessions.

F.1 Evaluation Data Sheet

Usability Issues:

Critical Incidents:

Group Observations:

Approach - Use of goals? What goals were formed? Other approach?

Strategies used? How often?

Communication and subgroups – people breaking off to work with some? Everyone talking? One person working alone? Dynamics from start to finish?

How do subgroups use 2D/3D representations? Same or different representations? Counts of each.

Critical incidents related to providing both forms of awareness?

Map Observations:

Struggle to navigate alone? Frustrated?

With two people looking at map, does one guide the other or go separate ways? Does one inform another on how to adjust?

Do users look at separate areas without issues?

Can they pursue different subgoals? And then meet up to discuss?

How much do they struggle to align viewpoints? Point out features to one another? Establish common frame of reference?

Do they communicate through pointing, discussions, or both?

Awareness critical incidents:

Radar map Observations:

Problems pointing out features due to lack of precision?

Intuitive to use?

Egocentric Observations:

Individuals can identify objects and areas of the space without problems due to lack of detail?

Do they see each other's avatars and make use of the avatar during the task?

Can they provide useful instructions not relative to their own viewpoint? Relative to their own viewpoint?

Are there spatial transformation problems? Descriptions are relative to viewpoint or no?

Can they share spatial ideas easily?

Do they struggle to explore the space?

Do they struggle with long-distance object movements?

Awareness critical incidents (+/-) in know where others were located or what they could see?

Exocentric-Egocentric Observations

Are there spatial transformation problems?

Object manipulation is efficient or problematic?

Is 2D used to perform large movements and 3D to finalize position?

What role does the user with the map play? The user with CVE?

Examples of actor-director roles? How were the roles assigned?

How many people look at map together and how many look at CVE together for subgoal?

Awareness critical incidents (+/-) in know where others were located or what they could see?

Appendix G Vita

Wendy Schafer
wschafer@vt.edu
<http://csgrad.cs.vt.edu/~wschafer>

Master of Science, Computer Science, December 1999
Virginia Polytechnic Institute & State University, Blacksburg, VA

Bachelor of Science, Computer Science and Mathematics, August 1998
Wofford College, Spartanburg, SC