Effects of Visual Displays on 3D Interaction in Virtual Environments

Dhruv Manek

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

Master of Science
In
Computer Science and Applications

Dr. Doug Bowman, Chair
Dr. Deborah Hix
Dr. Chris North

23rd June 2004
Blacksburg, VA, USA

Keywords: HMD, CAVE, Virtual Environments,
3D Interaction, Selection, Manipulation

Effects of Visual Displays on 3D Interaction in Virtual Environments

Dhruv Manek

Abstract

In this thesis, we explore some of the effects of virtual environment displays, specifically the head-mounted display (HMD) and the Cave Automatic Virtual Environment (CAVE), on 3D interaction tasks involving selection and manipulation. The motivation for this thesis comes from the lack of previous work that has studied the effects of differences between the HMD and the CAVE on 3D interaction tasks. We conducted three user studies to determine how the differences between these two displays affect selection and manipulation in a 3D environment. Our first study demonstrates that 3D selection and manipulation tasks can be affected by the display type. Our second user study shows that task performance can suffer when a selection and manipulation technique is migrated to a display for which it is not intended. The third user study we conducted suggests that we can modify a selection and manipulation technique and improve its usability in the display to which it is migrated. We conclude with a set of guidelines to ease the migration of selection and manipulation techniques from the HMD to the CAVE while trying to maintain usability.

Acknowledgements

The author would like to thank the following people.

- His advisor, Dr. Doug Bowman, for his continuous help, advice and the critiques to make this thesis better
- The members of his thesis committee: Dr. Deborah Hix and Dr. Chris North
- His parents for their continuous support
- His younger sister, Riddhima, who always had some advice waiting for him
- Chad Wingrave, Dheva Raja, Jian Chen, John Lucas and other members for the 3DI group at Virginia Tech for always helping out
- Patrick Shinpaugh for helping out with the CAVE hardware
- Subjects that participated in all his experiments
- Radhika, for always being there for him

Table of Contents

List of Figures

List of Tables

Chapter 1 - Introduction

In this chapter, we introduce virtual environments (or virtual reality) and 3D interaction. We also discuss the factors that motivated this research, introduce the problem statement and provide an overview of the rest of this thesis.

1.1 Virtual Environments

Howard Rheingold (1991) defined virtual environments (VEs) or virtual reality (VR) as where a user is "surrounded by a three-dimensional computer-generated representation, and is able to move around in the virtual world and see it from different angles, to reach into it, grab it, and reshape it". The first VE display was an HMD which was invented by Ivan Sutherland in 1968 (Sutherland 1968). It consisted of two cathode ray tubes (CRTs) next to the user's ears and tracking hardware to measure head orientation. Two mirrors were placed in front of the user's eyes to reflect the images from the CRT displays. Current HMDs resemble a helmet that attached with two small displays. These displays can be cathode ray tubes (CRTs) or liquid crystal displays (LCDs). The two displays allow for stereoscopic viewing by presenting separate overlapping images to each eye. Some HMDs are also provided with stereo headphones to provide users with audio feedback. Figure 1.1 shows a commercially available HMD.

Figure 1.1 The Cybermind Hi-Res800™ 3D HMD (http://www.hi-res800.com/)

Besides HMDs, various VE displays have emerged onto the scene, ranging from desktop monitors to immersive surround-screen displays like the CAVE (Cruz-Neira et al. 1993). CAVEs are VE displays with generally two or more walls and a floor. The walls are generally rear projected. The floor can be projected from below or from the top. While using the CAVE users are required to wear stereo glasses. These stereo glasses are synchronized with the projectors to allow for stereoscopic viewing by presenting separate images to each eye. Figure 1.2 shows a four-sided CAVE with a top-projected floor.

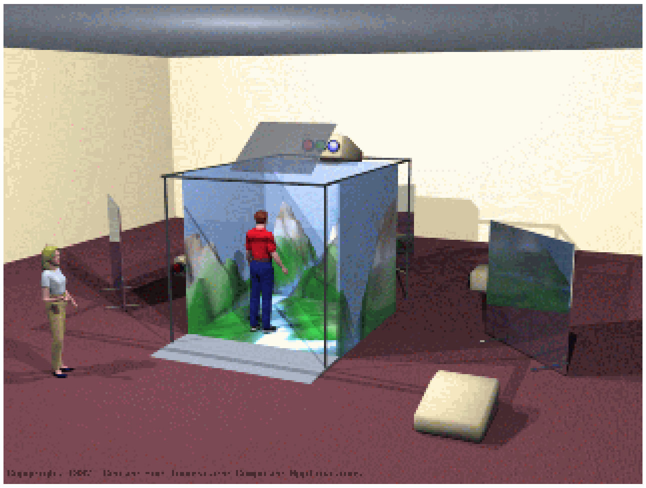

Figure 1.2 Four-sided CAVE (http://www.ncsa.uiuc.edu/VR/cavernus/GALLERY/ CAVEimages/cave_illustration01.gif)

Over the past three decades, computing power has progressed by leaps and bounds. Current tracking technology is able to provide 6 degrees of freedom (DOF) tracking. These advancements in technology have permitted VR to find use in several areas and moved it out of the computer science research laboratory. Some areas where VR technology has found prominent uses include the fields of visualization, training and simulation, architecture, medicine, entertainment, etc. (Brooks 1999, Bricken 1991, Schmitz 1993, Nilan, Silverstein and Lankes 1993)

1.2 3D Interaction

3D interaction refers to the actions of users within a 3D environment. A typical 3D interaction task is manipulating an object in the environment in terms of its location and orientation. 3D interaction is difficult since most people are used to the 2D desktop metaphor. Input devices like the mouse rest on a surface and thus are easier to move and point as opposed to 3D input devices (Bowman 1999). Additionally, 3D interaction is difficult since users typically have no haptic contact with objects in the environment (Flasar 2000). However, VEs can also overcome many limitations of the real world through the use of different interaction techniques. For example, the Go-go technique invented by Poupyrev et al. (1996) maps the motion of the user's physical hand to the virtual hand in a non-linear fashion, allowing a user to extend his virtual hand and select distant objects that would be out of reach in the real world. Most 3D interaction techniques can be classified into one of the following: navigation techniques, selection and manipulation techniques and system control techniques (Bowman et al. 2004). Navigation techniques allow users to travel from one part of the environment to another. A common navigation technique is the World-in-Miniature technique (Pausch et al. 1995). Selection and manipulation techniques allow users to select a particular object of interest and change its location and orientation within the environment. Common selection and manipulation techniques are Go-go (Poupyrev et al. 1996) and Ray-casting (Mine 1995). System control techniques allow users to change the system state. An example of a system control technique is TULIP menus (Bowman and Wingrave 2001). Along with the advancements in technology, the increased acceptance of VR technology in various fields can be attributed to the intuitiveness and ease of use of 3D interaction techniques such as these.

1.3 Motivation

The primary motivation for this thesis comes from the differences between HMDs and CAVEs. Both of them are commonly used immersive displays for different types of VE applications, however, their differences can have astounding effects on the usability of an application. Table 1.1 below illustrates some of the common differences between a commercially available HMD and a four-sided CAVE:

HMD	CAVE
Smaller field of view	Larger field of view
360° field of regard	Smaller field of regard due to the missing back wall
No physical screens that the users might run into	Presence of physical screens that the users might run into
Cannot see physical objects or other users in a multi-user environment	Can see physical objects or other users in a multi-user environment
No occlusion due to body parts	Occlusion by body parts due to see-through stereo glasses
Generally unlimited vertical viewing height	Vertical viewing height is limited by the height of the screens
No floor present	Top projected floors generally have lower visual quality compared to the screens
Lower resolution than CAVEs	Higher resolution compared to HMDs
Most HMDs are stereo capable but stereo not as good as in the CAVE	Better stereo than HMDs
More wires/cables	Fewer wires/cables

Table 1.1: Common differences between HMDs and 4-sided CAVEs

The *field of view* (FOV) of a display is defined as "the maximum number of degrees of visual angle that can be seen instantaneously on a display" (Bowman et al. 2004). HMDs generally have a horizontal FOV between 20 and 120 degrees (Bungert 2004). CAVEs, on the other hand, have a FOV of greater than 180 degrees assuming the user is looking at the front wall. An example of how this difference in FOV might affect interaction is when a user is trying to move objects from one end of the VE to another end of the VE. Due to the wider FOV of CAVEs, users will be able to see more of the VE and use less head rotation allowing them to perform the task faster as compared to HMDs.

The *field of regard* (FOR) of a display is defined as "the amount of the physical space (measured in degrees of visual angle) surrounding the user in which visual images are displayed" (Bowman et al. 2004). Because the screens are attached to the user's eyes, HMDs have a 360-degree horizontal and vertical FOR. A four-sided CAVE, on the other hand, has a 270-degree horizontal FOR. An example of how FOR differences might affect usability is while placing objects behind the user in the VE. In the HMD, users will be able to turn around and place objects behind them. In the CAVE this is not possible due to the absence of the back wall. Users might have to rotate the world (if the application provides the feature) in order to place the object behind them in the VE.

The third difference between HMDs and CAVEs is the presence of physical screens in CAVEs that users might run into. In HMDs, this problem is eliminated since the display screens move with the user. Running into physical screens in CAVEs may lead to loss of immersion and affect the usability of the application.

Owing to the use of see-through stereo glasses, CAVEs do not have the problem of users running into physical objects or other users in a multi-user environment. This is a problem in HMDs that may cause injury, loss of immersion and affect task performance. On the other hand, CAVEs are plagued with the problem of occlusion due to the use of see-through stereo glasses. According to Bowman et al. (2004), occlusion is "the phenomenon in which an object closer to the user partially obstructs the view of an object farther away". An example of occlusion is when a user is trying to touch an object that is behind another object in the VE. Since the screens are a few feet away from the user in the CAVE, the user's physical hand will occlude the object in the front (even though it would still be visible in the real world) causing a break in the believability of the display. HMDs are unaffected by this problem since the display screens are within inches of the user's eyes, and since the user's body is not visible.

HMDs, owing to their design, do not limit the maximum viewable height that users can see. Four-sided CAVEs (with no ceiling display) are limited by the height of their walls. Users will not be able to see beyond a certain maximum height. Selecting objects at a height is a task where the limited height of CAVEs may affect usability. Using the HMD, users may be able to complete the task without changing their position in the VE. In the CAVE users will have to move far away from the object or navigate to the height of the object in order to complete the task.

Another difference among HMDs and CAVEs is the presence of the floor in the CAVE. Top projected floors in the CAVE have lower visual quality than the screens because they are generally made out of wood which has a low reflectivity. Reflectivity is the ratio of light reflected from a surface to the ratio of light incident on that surface (http://emma.la.asu.edu/webdata/glossary.html). Owing to the lower visual quality of the top projected floor in the CAVE, user's might have difficulty in seeing objects that are projected on the floor. This in turn affects interaction, especially selection. In HMDs, the visual quality is the same throughout the space.

HMDs also have a lower resolution compared to CAVEs. The Virtual Research V8 HMD at Virginia Tech has a resolution of 640 x 480 pixels. CAVEs on the other hand offer much higher resolution as compared to HMDs. The CAVE at Virginia Tech has a resolution of 1280 x 1024 pixels on each side.

Most commercially available HMDs have two displays and can display images in stereo. However, most HMD-based applications allow only for monoscopic viewing since the graphics engines running these HMDs are not capable of displaying two different channels at a time (Bowman et al 2004). CAVEs, on the other hand, provide stereo images through the use of stereo glasses. These glasses coordinate with the projectors to display particular images to each eye at considerably high refresh rates (about 60 Hz overall or 30 Hz per eye).

Another difference between HMDs and CAVEs is that HMDs can be cumbersome and have more wires. CAVEs, on the other hand, do not require the user to wear heavy devices and have much fewer wires. This could lead to greater ease of interaction in the CAVE.

Besides the differences between HMDs and CAVEs, the other motivation for this thesis is that most 3D selection and manipulation techniques are HMD-based techniques. HMD-based techniques are interaction techniques that were designed specially for use in HMDs. Hand-centered Object Manipulation Extending Ray-casting (HOMER) (Bowman and Hodges, 1997) and Pen and Tablet (Angus & Sowizral, 1995) are examples of two such techniques. Moreover, a majority of the user studies involving these techniques have been performed using the HMD. Thus, we do not know how a different display like the CAVE will affect the usage of these techniques. Therefore, the purpose of this thesis is to uncover that interaction is indeed affected by the display type. We will also demonstrate that a HMD-based selection and manipulation technique will not perform as well in the CAVE and that it can be modified to provide the same level of usability in a display for which it is not intended. Based on the modifications, the final goal is to provide a set of generalized guidelines that future VE application developers can use to migrate their interaction techniques from HMDs to CAVEs.

The motivation to migrate interaction techniques from the HMDs to CAVEs also arises from the fact that most 3D interaction studies have been conducted using the HMD. Also, most of the common 3D interaction techniques have been developed for the HMD. It would be interesting to study how these HMD-based interaction techniques perform when migrated to the CAVE.

1.4 Problem Statement

The primary goal of this thesis is to develop guidelines to aid in the migration of selection and manipulation techniques from HMDs to CAVEs. However, there are two sides to this problem. First, we need to determine that interaction is indeed affected due to differences in displays. Second, we need to determine how the usability of a particular technique is affected by the migration from the HMD to the CAVE. This leads us to the problem statement:

> *Does the choice of a VE display have an effect on 3D interaction? Can we migrate HMD-based selection and manipulation techniques to the CAVE? If so, will there be any tradeoffs in user performance when this is done? Can we develop generalized guidelines to do this without compromising user performance?*

1.5 Overview of this Thesis

This thesis is divided into six chapters. The second chapter talks about the relevant work that has influenced this thesis. The third chapter presents the first user study that we conducted in order to determine that interaction is indeed affected by the display type. The fourth chapter talks about the second user study that we conducted to determine how the CAVE influences an HMD-based interaction technique. The fifth chapter discusses the third user study that we conducted to show that the changes made to interaction technique in the CAVE do indeed help to improve usability of the HMD-based interaction technique. In the final chapter, we summarize our work, present guidelines to migrate HMD-based selection and manipulation techniques to the CAVE, discuss our contributions to this field and talk about the future work.

Chapter 2 - Related Work

In this chapter, we define a set of terms that we will use throughout this thesis. We then discuss several common VE displays and review some literature that has studied the effects of their differences on 3D interaction. We also review some of the common 3D interaction techniques, specifically those for travel and selection and manipulation.

2.1 Terminology

In this section, we briefly define some of the common terms related to VEs that will be used in this chapter as well as the later chapters of this thesis.

2.1.1 Monocular displays

Monocular displays are one eye displays (http://www.nokia.com/nokia/0,,53716,00.html).

2.1.2 Binocular displays

Binocular displays are two-eye displays (http://www.nokia.com/nokia/0,,53716,00.html). Binocular displays can be biocular or stereoscopic (defined below).

2.1.3 Biocular displays

Biocular displays are two-eye displays where the same image is shown to both eyes (http://www.nokia.com/nokia/0,,53716,00.html).

2.1.4 Stereoscopic displays

Stereoscopic displays are two-eye displays where different images from a slightly different perspective are shown to the left and right eyes (http://www.nokia.com/nokia/0,, 53716,00.html). Showing different images to the left and right eyes makes users feel as if they are watching a 3D image.

2.1.5 Monocular, static depth cues

Monocular, static depth cues provide “depth information that can be inferred from a static image viewed by a single eye” (Bowman et al 2004). They include occlusion, linear perspective, atmospheric perspective, texture gradient, relative height, and relative size.

2.1.6 Motion parallax depth cues

Motion parallax depth cues are "depth information [that] is conveyed when objects are moving relative to the viewer (i.e., stationary-viewer motion parallax), when the viewer is moving relative to stationary objects (i.e., moving-viewer motion parallax), or through a combination of the two" (Bowman et al 2004).

2.1.7 Stereopsis

Stereopsis is "the perception of depth based on the differences in images that reach the 2 eyes" (http://www.usabilityfirst.com/glossary/main.cgi).

2.1.8 Refresh Rate

Refresh rate is "the maximum number of frames that can be displayed on a monitor in a second, expressed in Hertz" (http://dict.die.net/refresh%20rate/). Thus, if a display has a refresh rate of 60 Hertz (or Hz), then it can redraw the screen 60 times per second.

2.1.9 Immersion

According to Slater, Usoh and Steed (1995), immersion "includes the extent to which the computer displays are extensive, surrounding, inclusive, vivid and matching. They are surrounding to the extent that information can arrive at the person's sense organs from any (virtual) direction, and the extent to which the individual can turn towards any direction and yet remain in the environment."

2.1.10 Physical Travel Techniques

Physical travel techniques are those "in which the user's body physically translates or rotates in order to translate or rotate the viewpoint" (Bowman et al 2004).

2.1.11 Interocular distance

Interocular distance is defined as "the distance between the centers of rotation of the eyeballs of an individual or between the oculars of optical instruments" (http://www.dtic.mil/doctrine/jel/doddict/data/i/02748.html)

2.2 VE Displays

2.2.1 Monitors

Conventional cathode ray tube (CRT) monitors have commonly been used as displays for 3D applications as they are inexpensive and can provide monocular and motion parallax depth cues. Additionally, CRT monitors can also provide stereopsis through the use of stereo glasses and a stereo-capable graphics card (Bowman et al. 2004). However, these displays are required to have a high refresh rate (generally 100 Hz or greater) since they need to display images for the left and right eye at alternating frames. Thus, liquid crystal display (LCD) monitors cannot typically be used for VE applications as they cannot attain such high refresh rates (Bowman et al 2004). Using CRT monitors in conjunction with a pair of stereo glasses and head-tracking makes moving-viewer motion parallax easier to achieve. These setups are usually known as fish-tank virtual reality (Ware et al., 1993). Figure 2.1 shows a set up of a fish-tank VR.

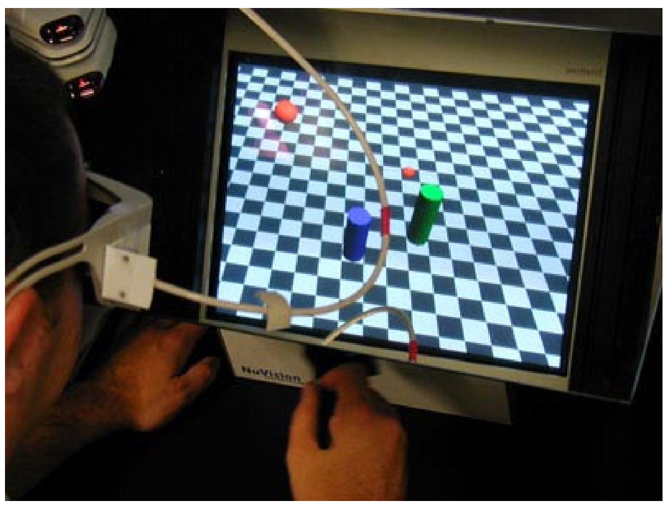

Figure 2.1 Fish-tank VR
(http://www.ccom.unh.edu/vislab/FishTankVR.html)

The advantages of this type of display are that they are inexpensive compared to other types of VE displays and allow the use of any input device, including the keyboard and mouse. On the other hand, CRT monitors are not immersive and limit the user's ranges of movement due to the low FOR they offer. Also, due to the small display size any physical input devices may block the display, breaking any stereo illusion.

2.2.2 Workbench Displays

Another type of VE display is the workbench display. The original workbench was developed by Krüger and Fröhlich in 1994 and was known as the Responsive Workbench (Krüger and Fröhlich 1994). Variants of the Responsive Workbench include the ImmersaDesk, VersaBench and VisionMaker (Bowman et al. 2004). The workbench display is a projection-based, table format display device that uses stereo glasses and head tracking to create a semi-immersive VE. The screen might be oriented vertically, horizontally or at an angle to suit a particular type of task and/or application. Figure 2.2 shows a horizontal workbench display.

Figure 2.2 The Responsive Workbench (http://www.ait.nrl.navy.mil/vrlab/projects/Workbench/Workbench.html)

The advantages of workbench displays are that they provide relatively high resolution and serve as an intuitive display for certain types of applications (Bowman et al. 2004). They also provide visual depth cues like monitors and can be rotated making them pretty flexible. The disadvantages for this kind of display are that they are not completely immersive since they do not enclose the user (Bowman et al. 2004). Additionally, the range of viewpoints is restricted since parts of the screen might not be visible from certain viewpoints (Bowman et al 2004). Another disadvantage of the workbench display is that physical travel techniques are not appropriate since users have limited maneuverability (Bowman et al 2004).

2.2.3 Hemispherical Displays

A third type of VE display device is the hemispherical display. These are projection-based devices that consist of wide-angle lens that is attached to a projector that displays images in a 180° x 180° field of view. Users sit in the front of screen and can interact with the application using a keyboard and mouse or other 3D input devices. The advantages of hemispherical displays are that they are more immersive than monitors and brighter as they are front-projected (Bowman et al. 2004). However, direct selection, manipulation and navigation techniques do not work well in hemispherical displays because moving too close to the display might cast shadows on the screen, breaking the stereo illusion and occluding part of the VE (Bowman et al. 2004). Additionally, due to the curvature of the screen, resolution and image quality are not uniform across the screen (Bowman et al. 2004). Figure 2.3 shows a hemispherical display.

Figure 2.3 Hemispherical Display
(http://www.photonics.com/spectra/applications/XQ/ASP/aoaid.337/QX/read.htm)

2.2.4 Head-Mounted Displays

Head-mounted displays (HMDs) are the most common head-coupled display devices used for VE applications. A typical HMD consists of a helmet with two small CRT or

LCD displays and an adjustable lens system. They can produce stereoscopic viewing by presenting separate overlapping images to each eye. However, since most graphics engines are not able to simultaneously display two different channels most HMD-based applications allow only monoscopic viewing (Bowman et al 2004). Interaction and navigation are achieved through specialized input devices, for example 3D mice, wands or data gloves. The HMD and the input devices are tracked in real-time using a tracking system which is transformed by a computer to update the user's position within the VE. HMD systems are usually used by single participants however it is possible for multiple participants to inhabit the same virtual environment. Figure 2.4 shows a user performing tasks using a HMD. Some HMDs offer see-through options where users can see the real world too. This technology is known as augmented reality.

Figure 2.4 A user wearing a HMD and performing tasks (http://www.view.iao.fhg.de/description.html)

The biggest advantage of HMDs over other VE displays is that users can have complete physical immersion allowing for a 360° FOR (Bowman et al. 2004). Another advantage of HMDs over projection-based displays is that multiple users can each have their own HMD with a tracked view of the virtual world (Bowman et al 2004). Additionally, HMDs are much lighter and portable than most projection-based displays. Disadvantages of HMDs include a limited FOV which may produce distortions in perception of size and distance (Bowman et al 2004). Another disadvantage that plagues stereoscopic HMDs is that they do not provide a way to accommodate different interocular distances of users

(Bowman et al 2004). Lastly, HMDs are heavy and can cause neck strain if used for an extended period of time.

2.2.5 Surround-Screen Displays

A surround-screen display is a VE output device that consists of 3 or more screens that surround the user, completely immersing them in the VE. The screens are typically between 8 to 12 feet in width and height and rear-projected to eliminate the possibility of the user casting shadows on the screen. The first surround-screen VE system was called the CAVE and was developed by Cruz-Neira et al. at the University of Illinois at Chicago in 1993 (Cruz-Neira et al. 1993). It consisted of three walls and a floor. The CAVE projectors use an active stereo system to provide true stereoscopic imagery. Infrared LCD shutter glasses are used in conjunction with the projectors to provide true stereo at a rate of sixty frames per second (thirty frames per second for each eye). Interaction and navigation can be achieved by using 3D input devices like the wand. Users are tracked and the environment is updated when they navigate in the VE or walk a few steps within the enclosed area of the screens. Figure 2.5 shows a CAVE inhabited by several users.

Figure 2.5 Multiple users in a CAVE (http://www.view.iao.fhg.de/description.html)

The advantage of using surround-screen displays like the CAVE is that they provide higher resolutions as opposed to monitors. Additionally, they also have a large FOV and FOR (Bowman et al. 2004). A large FOV permits users to utilize their peripheral vision (Bowman et al. 2004). These displays also provide monocular depth cues and motion parallax and when users are tracked and wearing stereo glasses, they also provide additional motion parallax cues and stereopsis. The disadvantage of surround-screen displays is that they are expensive and require a large amount of space. Additionally, like other projection-based systems users might have difficulty seeing objects in stereo under certain conditions. Also, when more than one user inhabits the surround-screen device, they will see the images from the tracked user's perspective (Bowman et al. 2004). There will be no response when the untracked users move.

The above discussion presents a brief description of the most common VE displays. It is apparent that all these displays have differences pertaining to FOV, FOR, resolution, input devices that can be used, etc. Taking into consideration all these differences, there has been very little work that studies their effect on 3D interaction. Below is a description of some of the work done and their results.

2.3 Effects of VE Displays and Their Differences on 3D Interaction

Since we are studying the effects of different displays, specifically the HMD and the CAVE on 3D interaction, it is necessary to study some of the inherent differences between two displays. For the purpose of this thesis, we used the Virtual Research V8 HMD and a 10 x 10 x 10 feet CAVE with projections on the floor and front, left and right walls. The V8 HMD has a resolution of 640 x 480 pixels and a diagonal FOV of 60 degrees. Figure 2.6 shows the HMD we used for this thesis.

Figure 2.6 The Virtual Research V8 HMD (http://www.virtualresearch.com/)

The CAVE has a resolution of 1280 x 1024 pixels on each screen. The CAVE also has a horizontal FOV of greater than 180 degrees assuming the user is looking at the front wall.

However, in practice this is not achieved due to the use of stereo glasses (Bowman et al. 2002). Additionally, the missing back wall imparts a lower FOR to the CAVE as opposed to the 360-degree FOR of the HMD. Other differences between the HMD and the CAVE include the presence of physical screens, occlusion due to body parts and the lower visual quality of the floor of the CAVE. Additionally, the presence of the physical screens also restricts the maximum viewable height in the CAVE (see section 1.3). Figure 2.7 shows the CAVE at Virginia Tech that we used for this work.

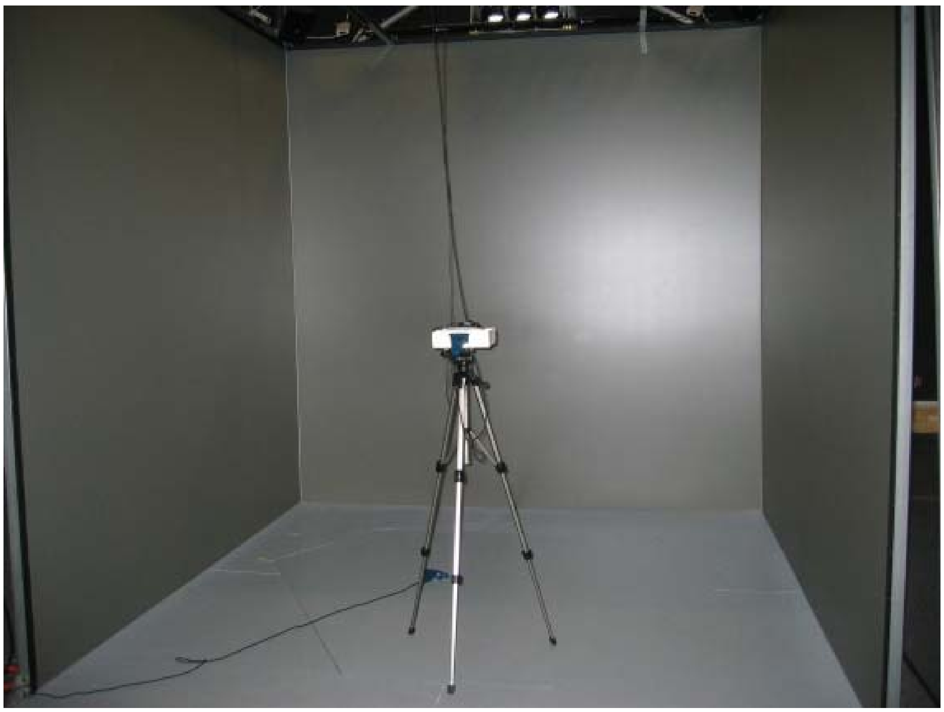

Figure 2.7 The CAVE at Virginia Tech

Considering the differences between the HMD and the CAVE, there is little work done that has studied the effects of differences in VE displays on user performance. Arthur (2000) studied the effects of FOV on performance with HMDs. He compared three HMDs having FOV of 48°, 112° and 176° and found that restricting the FOV degraded performance on searching and walking tasks. Compared to the 176° FOV HMD, user performance degraded by 12% and 24% respectively for the 112° and 48° FOV HMDs for searching tasks. Additionally, for the task of walking through a simple maze-like environment, FOV degraded performance by 23% at 112° and by 31% at 48°. Another study conducted by Dolezal (1982) found that large eye movements were substituted by head movements when the FOV was restricted by vision goggles.

Bowman et al. (2002) compared users' preferences for real and virtual turns in HMDs and CAVEs. They had users walk through corridors in the HMD and the CAVE. The results of their user testing show that natural rotation is the preferred technique of rotation in HMDs. They also found a gender-based preference i.e. females preferred natural rotation and used it as much as possible. Additionally, they also concluded that in environments where turning was frequent, HMD users were more likely to maintain spatial orientation than CAVE users. Based on their experimental results, they posited a guideline: "For VE applications involving navigation through enclosed spaces and frequent turning, choose an HMD with head tracking to provide increased efficiency and spatial orientation to users."

Kjeldskov (2001) conducted a study evaluating the usability of more than 40 interaction techniques for specific combinations of displays and interaction devices. The displays that they used were an 8 x 3 meter panoramic display covering a field of view of 160 degrees and a six-sided CAVE measuring 2.5 meters on each side. The tasks users had to perform primarily addressed orientating and moving in the virtual 3D spaces and selecting and manipulating 3D objects. Their primary conclusion from the experiments was that the same interaction techniques do not work equally well with panoramic displays and six-sided caves. Other specific results include:

1. Untraditional use of headtracking may support orientating in partial immersive displays, though introducing a problematic boundary between interaction in physical and virtual space.
2. Rotating the world in full immersive displays using e.g. joysticks may complement headtracking by letting the user view the VE from odd perspectives but might limit the feeling of immersiveness.
3. Non-tracked 3D interaction devices work fine for orientating and moving in partial immersive displays but can be problematic in full immersive displays if the user turns around physically.

4. The virtual hand approach has different potentials in partial and full immersive displays. The virtual pointer approach works fine for picking 3D objects but is problematic for rotating and moving them.

The above literature talks about the effects of restricting FOV in HMDs and the effect of two different displays on user interaction. However, none of these papers discuss the effects on user performance for selection and manipulation tasks using the HMD and the CAVE. Additionally, none of the papers mention guidelines for migrating applications from one display to the other.

2.4 Related work in 3D Interaction Techniques

2.4.1 Travel Techniques

There are several different travel and navigation techniques for VEs. Pausch et al. (1995) present the World-in-Miniature (WIM) technique for navigation in VEs. Using the WIM, users can move an iconic representation of themselves and they will be flown to that location in the VE. Zeleznik et al. (2002) present the Zoomback and LaserGrab techniques for navigation in VEs. Using these techniques, users can point to an object surface using a virtual laser pointer and they will be translated near that object. Using the ZoomBack technique, users are translated within 2 feet of the selected object while in LaserGrab the relative distance between the user's head and hand is used to determine the user's location relative to the target object. Lastly, Ware et al. (1988, 1990, 1996) present the "flying," "eyeball-in-hand," and "scene-in-hand" metaphors for virtual camera control.

2.4.2 Selection and Manipulation Techniques

Selection and manipulation techniques can be categorized into three categories: reaching techniques, ray-based techniques and multi-modal techniques.

2.4.2.1 Reaching Techniques

Reaching techniques are the most intuitive selection techniques; however, they have the drawback that users cannot select objects beyond their arm distance. This limitation was remedied by Poupyrev et al. (1996) though the invention of the Go-go technique. Go-go is an egocentric interaction technique, meaning that users interact from within the virtual environment (Poupyrev et al., 1998). It works by non-linearly mapping the motion of user's physical hand and the effected motion of the virtual hand in the immersive VE. The result of Poupyrev et. al's initial user evaluation was that Go-go was intuitive and easy to use since it imitated real world human behavior and was modeless (Poupyrev, et al., 1996). Bowman and Hodges (1997) tested Go-go and its variants like fast Go-go and stretch Go-Go with 11 subjects on selection and manipulation tasks using an HMD. A majority of the subjects (greater than 6) commented on the finite range and imprecise grabbing of Go-go.

2.4.2.2 Ray-based Techniques

Ray-based techniques require users to select an object by intersecting it with a ray emanating from their finger. This technique is commonly known as ray-casting (Mine, 1995) and is analogous to selecting a file on the desktop using a mouse pointer. Bowman and Hodges (1997) tested ray-casting and its variant called ray-casting with reeling with 11 subjects using an HMD. Almost all subjects (greater than 9) commented on the inability of ray-casting to move selected objects in or out, the difficulty to rotate the selected objects and ease of grabbing using Ray-casting. Bowman and Hodges (1997) developed a hybrid technique called HOMER (Hand-centered Object Manipulation Extending Ray-casting) that combines ray-casting and object-centered manipulation at-a-distance.

2.4.2.3 Multi-modal Techniques

Multi-modal selection and manipulation techniques include the "put-that-there" technique (Bolt 1980) that allows users to use voice commands as well as gestures to create objects and specify their target location.

There are also some selection and manipulation techniques that do not fit into any of the above categories. These include the Worlds in Miniature (WIM) technique described by Stoakley, Conway and Pausch (1995) and the Voodoo Dolls technique described by Pierce, Steams and Pausch (1999). In the WIM technique, users are provided with a hand-held miniature copy of the world such that they can manipulate the environment using a “God’s eye view”. Thus users can manipulate objects in the environment by two different ways, through the life-sized world or using the WIM. In the Voodoo dolls technique, users create miniature dolls, which are hand-held copies of the objects in VE. Users then use two-handed interaction techniques to manipulate these dolls using one doll as a frame of reference.

The above literature introduces some of the common 3D travel and selection and manipulation techniques. They also mention the results of user studies done using these techniques on the HMD. However, there is no literature that presents the results of user testing done in the CAVE or compares the results of similar tasks performed using one of the several selection and manipulation techniques in the HMD as well as the CAVE.

Chapter 3 – General Effects of Display Type on 3D Interaction

In the previous chapters, we discussed some of the differences between the HMD and the CAVE. This chapter outlines an experiment designed to evaluate the effect of their differences on 3D interaction.

3.1 About the experiment

3.1.1 Purpose

The purpose of this experiment was to show that interaction is affected by the characteristics of the display. Thus, our general hypothesis was that there will be differences in usability and task performance between display platforms for specific types of tasks and environments. These differences will arise due to differences among the display platforms. For example, due to the limited height of the screens in the CAVE, the maximum viewable height in the CAVE is limited. Also, the CAVE has a wider field of view as compared to the HMD. Other differences include the missing back wall and occlusion due to body parts, especially hands and feet, due to the use of see-through stereo glasses in the CAVE. Based on these differences, we laid down a set of specific hypotheses:

1. Due to the limited viewable height in the CAVE, users will find selecting and manipulating objects at a height easier in the HMD than in the CAVE.
2. Due to the wider field of view in the CAVE, users will be able to perform wide field tasks easier in the CAVE versus the HMD.
3. Due to the missing back wall in the CAVE, users will find it easier to place objects behind them in the HMD as opposed to the CAVE.
4. Due to occlusion by body parts in the CAVE, users will find it difficult to place objects near their feet in the CAVE. The same shall not be an issue in the HMD.

3.1.2 Brief Outline of the Experiment

This was an informal experiment as opposed to a formal, controlled experiment since we were more interested in qualitative, large-scale differences between the HMD and the CAVE. Thus, users were asked to perform tasks in the HMD as well as the CAVE. The tasks consisted of selecting and manipulating beams and columns in the environments using the Go-go technique. Tasks were designed specifically to check our hypotheses as to how the differences between the HMD and the CAVE would affect user performance in selection and manipulation tasks.

3.2 Subjects

This user study involved 6 users. All of them were unpaid volunteers and students at Virginia Tech. Their average age was 27. There were five males and one female. The user population was equally split between novice and experienced users of VEs.

3.3 Apparatus and Implementation

The HMD in the experiment was the Virtual Research V8 HMD. It supports a resolution of 640 x 480 pixels with a 60° diagonal FOV. Additionally, the HMD presented biocular images to the users meaning that both the LCD displays presented the user with the same image. The CAVE used in the experiment was 10 x 10 x 10 feet with projections on the floor and front, left and right walls. Each screen had a resolution of 1280 x 1024 pixels and presented stereo images to the users wearing stereo glasses. We used an Intersense IS900 tracker with both displays to track the head and hand of the users. For this study, we used the Virtual Structural Analysis Program (VSAP) application in the HMD and the CAVE. VSAP was developed by Bowman et al. (2003), and allows users to create and modify virtual building structures for the purpose of visualizing the effects of earthquakes. It allows users to navigate to different parts of the world by pointing the wand in the direction they wish to move. Additionally, it uses the Go-go technique invented by Poupyrev et al. (1996) to allow users to select and manipulate objects in the VE. Go-go allows users to select and manipulate objects at a distance by non-linearly mapping the user’s physical hand motion to the motion of the virtual hand in the VE beyond a certain threshold distance. Figure 3.1 shows the mapping function for the Go-go technique.

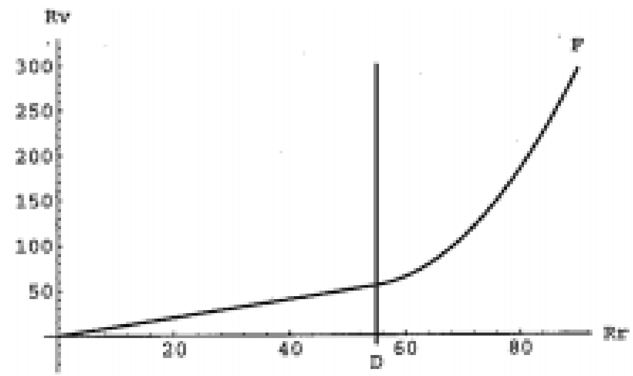

Figure 3.1 Mapping function for the Go-go technique. R_r represent the physical hand distance and R_v represents the virtual hand distance. (Poupyrev et al. 1996)

The VSAP implementation in the HMD allows users to specify commands using a tablet and a stylus. Using the stylus, users can click on the buttons on the tablet and create different types of objects in the VE. Figure 3.2 shows the tablet menu. VSAP for the HMD was developed using the Simple Virtual Environment (SVE) library and runs on a PC running Windows 2000.

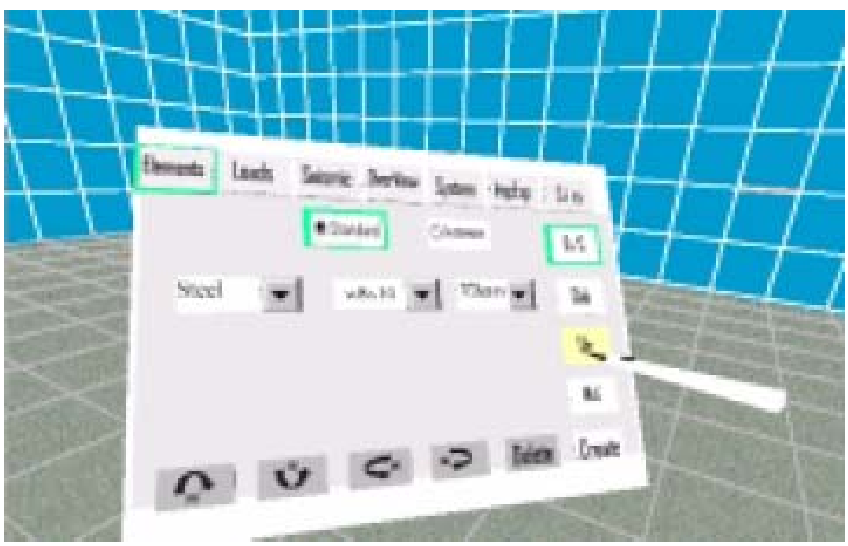

Figure 3.2 The tablet interface for VSAP in the HMD (Bowman et al. 2003)

The VSAP implementation in the CAVE uses a ring menu. Using the upper left and right buttons on the wand, users can rotate the ring menu to bring the specific buttons to the bottom. The active buttons (the two buttons on the bottom) can then be selected using the lower two buttons on the wand. The joystick on the wand can be used to rotate the world around the user. Figure 3.3 shows the ring menu interface for VSAP in the CAVE. The CAVE implementation of VSAP was developed using the Device Independent Virtual Environments - Reconfigurable, Scalable, Extensible (DIVERSE) library (Kelso et al. 2002) and runs on a Silicon Graphics Power Onyx running Irix 6.5.

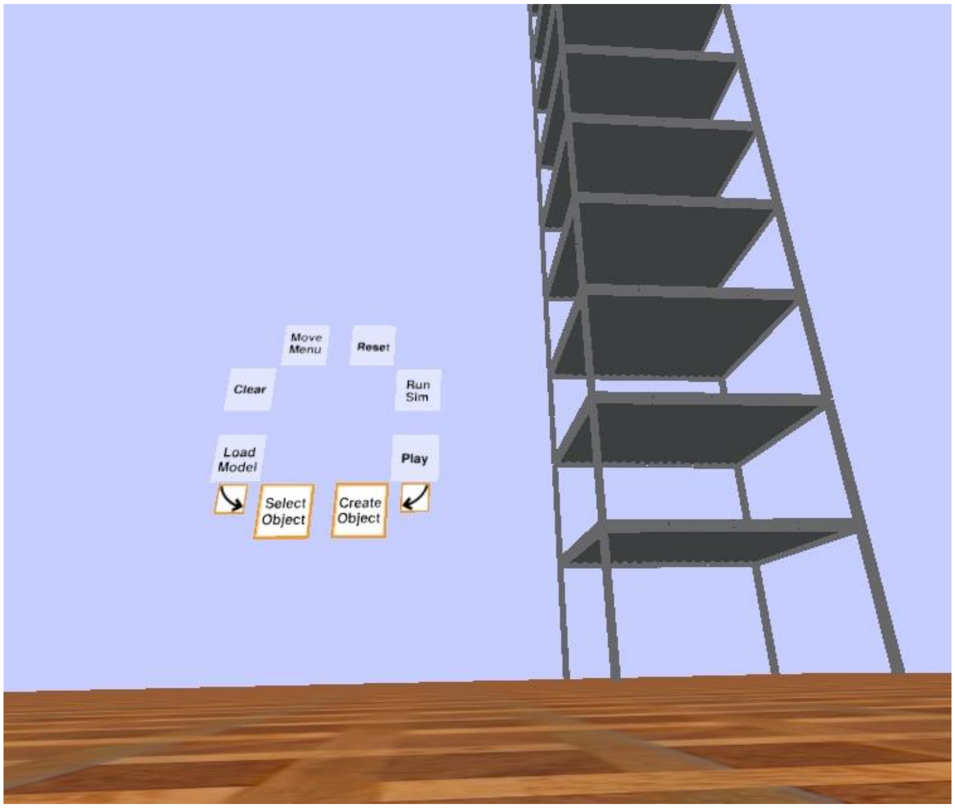

Figure 3.3 The ring menu interface for VSAP in the CAVE

Although both the implementations of VSAP are different in the way they allow users to interact with the environment, these differences were mostly irrelevant for our experiment. Since none of our tasks involved creating new objects in the environment, users did not use the tablet in the HMD. They just used the stylus for navigation and for Go-go. In the CAVE, users used the “Select Object” button on the ring menu so select a particular object after they had touched it using the virtual hand. The only other button that was used was the “Drop” button to drop the object attached to the virtual hand. Thus, functionally these implementations of the VSAP were similar and we could compare them for identical tasks in the HMD and the CAVE.

3.4 Experimental Design

The study was composed of a total of 8 tasks, 4 each in the HMD and the CAVE. The order of the displays was counterbalanced to ensure that users would not always perform better on the second display owing to the practice gained in the first display. We collected four types of data for this experiment: subject comments from the think-aloud protocol, observations of behavior, subjective difficulty ratings for each task and user preferences. Additionally, this was a within-subjects experiment meaning that all users were asked to perform tasks in both display conditions.

3.5 Procedure

Subjects started the experiment by filling out the pre-questionnaire (see Appendix A1) that contained demographic information such as age, gender, handedness and the subject's prior experience with VEs. Next, those unfamiliar with the VE equipment were told about the basic setup and hardware. Also any questions about the hardware and/or the experiment were answered at this stage.

Users were then allowed to spend 5 to 10 minutes each in the HMD and the CAVE to familiarize themselves and to identify ways in which they could interact with the environment. For evaluating task performance, it was absolutely necessary for all users to be introduced to all the functionality offered by the application. Hence, after the users had explored the application by themselves, the experimenter walked the users through the functionalities they had missed. By the end of this phase, the experimenter made sure the users had adequate knowledge of the capabilities of the application. The users were free to ask questions at this stage. They were encouraged to talk aloud about their reactions to the different features.

3.5.1 Tasks

These tasks were not timed. The experimenter read each task aloud and the users were asked to try to think aloud while performing them. The main intention of this set of tasks was to get qualitative data and gain a better understanding of the user's thought process. The users were told to perform the task by themselves. However if they got stuck on a task, they could go on to the next task. Below is a list of tasks that the users were asked to perform (also in Appendix A2):

1. Look at the 10-floor structure in front of you (figure 3.3). Now using your virtual hand, select the column on the 6th floor that is closest to you. You will see that the column is highlight once it is selected.
2. Approach the arch in front of you (figure 3.4). Using your virtual hand and the tablet/ring menu, select each piece and place it behind you to build a new arch. The new arch should be a “mirror reflection” of the old arch.
3. Rotate the world to bring the new arch that you built on the left side/screen. Select each piece again and build a new arch on the right screen such that the new arch is a “mirror reflection” of the old arch.
4. Place the beams in front of you (figure 3.5, left) to create a square around your feet such that in the end you are in the center of the square (figure 3.5, right).

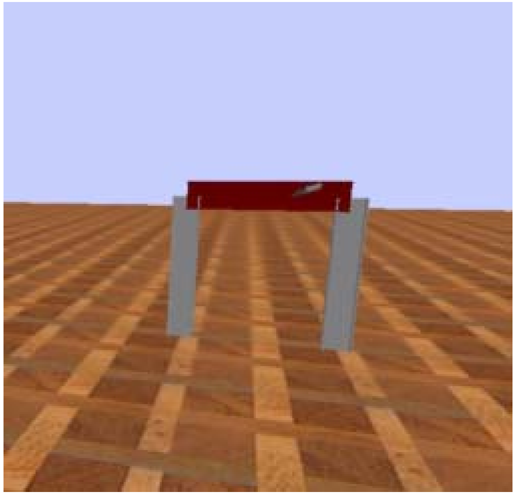

Figure 3.4 The arch that users were to manipulate in tasks 2 and 3

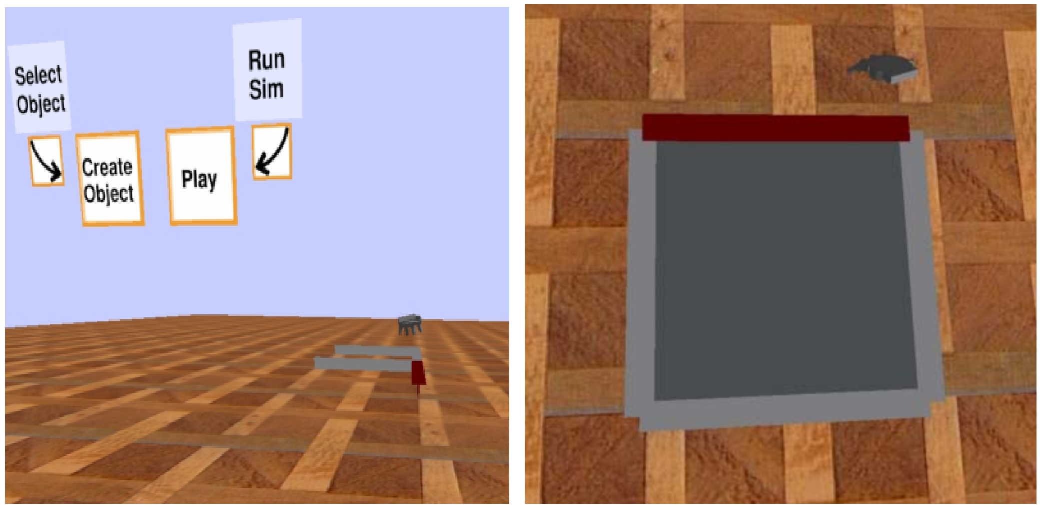

Figure 3.5 The start and end conditions for task 4

At the end of each task, users were asked to rate the difficulty level of the task on a scale of 1 to 5, 1 being least difficult to 5 being most difficult. This phase allowed the experimenter to collect qualitative data about the users' perception and also some quantitative data on the difficulty level of each task. The tasks that users performed were a good mix of tasks that users will perform in any selection and manipulation application and also serve to check our hypotheses.

At the end of the experiment, users were asked to fill out a post-questionnaire (see Appendix A3) asking their preference of display for each type of selection and manipulation tasks. Users were also asked about specific difficulties they had using the Go-go technique.

3.6 Observations

Several observations were made and noted down while users performed the tasks listed above. The observations include comments by users while "thinking aloud" and other notes taken by the experimenter while observing the user performing the tasks.

For the first task in the HMD, a majority of the users (5 out of 6) looked up at the column on the 6^{th} floor, navigated to that level by flying up and completed the selection task. In the CAVE, all users complained that they could not see the column on the 6^{th} floor. A couple of users tried to select the column by moving away from the structure. The other four users navigated to the 6^{th} floor and then completed the selection task. About 3 users commented on losing the virtual hand while trying to select distant objects.

For the second task in the HMD, all users did not have any difficulty placing objects behind them in the environment. In the CAVE, however, one user was completely lost as to how the task was to be accomplished. It struck him after a while that the world had to be rotated in order to see environment behind him. Another problem that users had was that sometimes they rotated the world too much or too little. They realized they had rotated the world too much when they saw the original arch back on one of the screens. While rotating the world, one experienced user realized that this task could be accomplished in a left to right screen fashion and thus had rotate to world only by 90° and then completed the task with ease.

The third task involved users rebuilding the arch from the left to the right. In the HMD, most users did not have any trouble finishing this task. One user complained of neck strain due to the weight of the HMD and excessive head rotation from the left to the right. In the CAVE, the same user mentioned that the task was easier than in the HMD since it involved less head rotation. Other users mentioned that this task was a duplicate of the previous task in terms of selecting and placing the pieces of the arch.

The last task involved users placing beams around their feet to form a square in such a way that they would be inside the square. This task was used to validate our hypothesis that due to occlusion due to feet in the CAVE, users would prefer the HMD. In the HMD, one user commented that it was easy to place the objects since he could not see his feet. In the CAVE, one user flew up to a considerable height, placed the beams to form the square and flew back down inside the square. During an interview after completing the tasks, one other user commented that she found the CAVE easier for this task since the sight of her feet provided her a sense of location and orientation.

3.7 Results

Figures 3.6 through 3.9 below show the difficulty levels for each task by all 6 users:

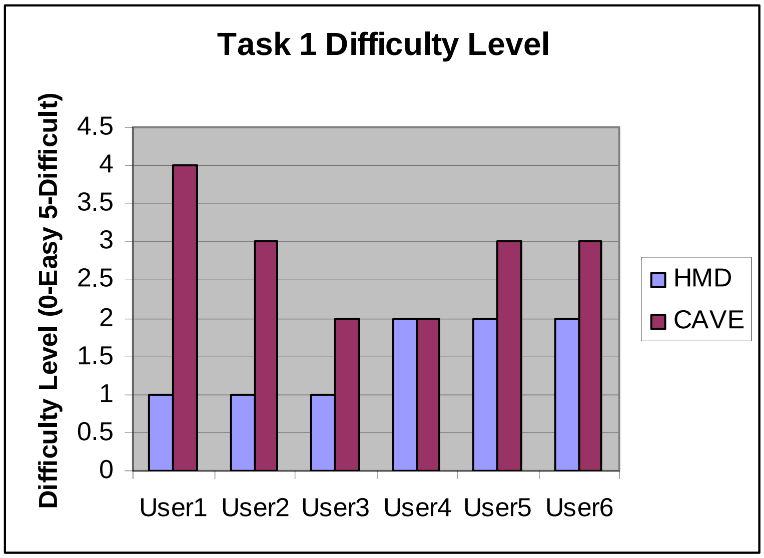

Figure 3.6 Graph showing the average difficulty levels for Task 1

From figure 3.5, the average difficulty level of task 1 in the HMD was 1.50 (standard deviation = 0.500) while the average difficulty level in the CAVE was 2.83 (standard deviation = 0.687). The lower average difficulty level for this task in the HMD can be explained by most users' comments in the think-aloud protocol that they found it harder to see the column on the 6^{th} floor in the CAVE compared to the HMD.

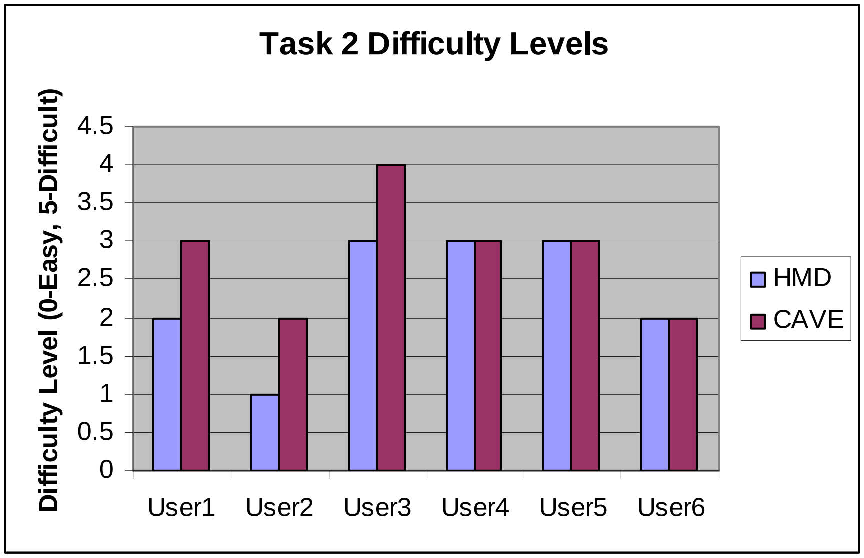

Figure 3.7 Graph showing the average difficulty levels for Task 2

From figure 3.6, the average difficulty level of task 2 in the HMD was 2.33 (standard deviation = 0.746) while the average difficulty level in the CAVE was 2.83 (standard deviation = 0.687). Most users found this task more difficult to accomplish in the CAVE because they had to rotate the world as opposed to the HMD where they could just rotate their head to see the VE behind them.

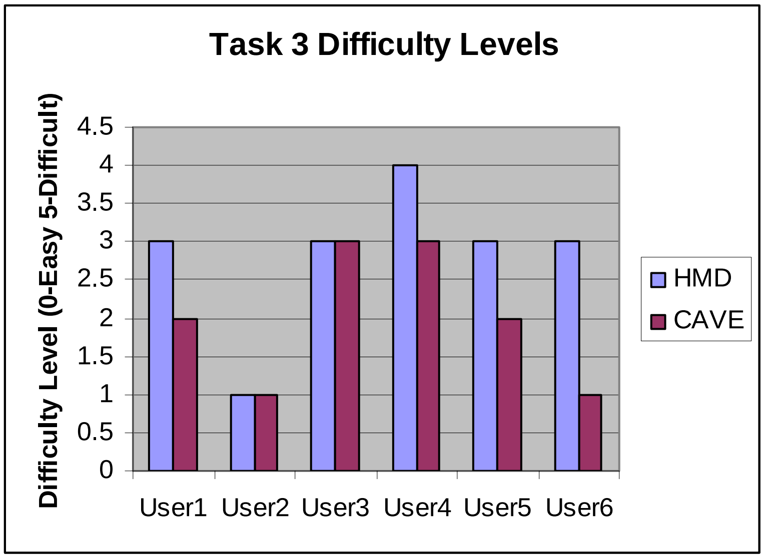

Figure 3.8 Graph showing the average difficulty levels for Task 3

From figure 3.7, the average difficulty level of task 3 in the HMD was 2.83 (standard deviation = 0.898) while the average difficulty level in the CAVE was 2.00 (standard deviation = 0.816). Users felt that this task was easier to complete in the CAVE because it involved less head rotation compared to the HMD. Additionally, the weight of the HMD also made head rotation difficult.

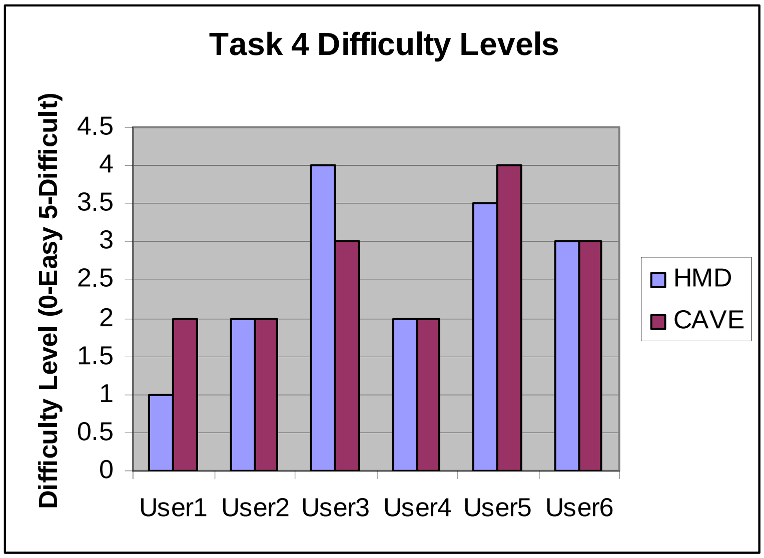

Figure 3.9 Graph showing the average difficulty levels for Task 4

From figure 3.8, the average difficulty level of task 4 in the HMD was 2.58 (standard deviation = 1.017) while the average difficulty level in the CAVE was 2.66 (standard deviation = 0.745). Some users in the HMD felt that since there was no occlusion due to their feet, they felt that this task was easier to accomplish in the HMD. On the other hand, other users in the CAVE commented that since they could see their feet, it gave them a sense of orientation and made it easier to place the beams around them, thus they preferred the CAVE.

Table 3.1 presents users' responses for the preference of one display type over the other from the post-questionnaire.

Task Type	HMD	CAVE	Both	No pref.
Selecting objects at a height	6	0	0	0
Placing objects behind the user	6	0	0	0
Placing objects from left to right	1	5	0	0
Placing objects close to the user	2	2	1	1

Table 3.1 Table showing user's preferences from the post-questionnaire for the HMD or the CAVE for particular types of tasks

3.8 Analysis

The difficulty levels of Task 1, users' "think aloud" comments and responses from the post-questionnaire support our first specific hypothesis that users find it easier to select objects at a height in the HMD compared to the CAVE.

Based on the difficulty levels of task 2 and from the user's responses, our second specific hypothesis is also supported that users prefer the HMD for placing objects behind them in the environment.

Similarly, the difficulty levels of Task 3, users' comments and their preferences from the post questionnaire also support our third specific hypothesis that the CAVE is the preferred display type for wide field tasks like placing objects from left to right.

For the last task of placing objects close the user, the averages of the difficulty level are too close to draw a decision in any favor. Additionally, the users' responses are also conflicting and ambiguous. From user interviews, some users felt that the HMD was a better fit for such tasks since they could not see their hands or feet which allowed them to place the beam around their feet. However, users who preferred the CAVE felt that since they could see their feet it gave them a sense of their location which made it easier for them to place the beams around their feet.

3.9 Conclusions

From the analysis above we conclude that the HMD is preferred for selecting objects at a height and placing them behind the user. The CAVE is preferred for wide field tasks like selecting and placing objects from one side of the environment to the other owing to its wide FOV. These results validate the first three of our specific hypotheses. The fourth hypothesis could not be validated since users preferred the HMD and the CAVE for placing objects close to them.

More importantly, we demonstrated that display type can have an effect on several aspects of usability during the performance of 3D interaction tasks. The properties of the displays can be used to predict the usability of 3D interaction tasks using the same 3D interaction technique in different displays.

Chapter 4 – The Effect of Display Migration on the Usability of a 3D Interaction Technique

In the previous chapter, we discussed the experiment whose purpose was to show that interaction is affected by the characteristics of the display. This chapter outlines the experiment we designed to examine the migration of an HMD-based interaction technique, the World-in-Miniature (WIM) technique, which we tested in the HMD and the CAVE.

4.1 About the experiment

4.1.1 Purpose

The purpose of this experiment was to demonstrate that when an interaction technique designed for the HMD is migrated to the CAVE, usability may decrease. Thus, our hypothesis was that the usability of the WIM technique will decrease when migrated to the CAVE. Based on the results of this experiment, we can make changes to the WIM technique such that its usability in the CAVE is compromised minimally (see chapter 5).

4.1.2 Brief Outline of the Experiment

This was a formal, controlled experiment where users were timed for each task they performed in order to gather some quantitative data. Users were asked to perform selection and manipulation tasks using the WIM technique in the HMD and the CAVE. The WIM scale in our experiment was $1/20^{th}$ the scale of the world. The tasks consisted of selecting a cube from a set of nine cubes in the VE and then placing it between two target indicators which were essentially wooden cylinders in the WIM (see figure 4.3). These tasks were chosen as they required users to perform precise selection and manipulation.

4.2 The World-in-Miniature Technique

The WIM technique was invented by Stoakley, Conway and Pausch (1995). It works by providing users with a handheld copy of the virtual world. Users can select and manipulate objects in the environment by using the miniature copies of the objects in the WIM. An advantage of the WIM technique is that users can get a "God's eye view" of the world in addition to the immersive environment (Pausch et al. 1995). Pausch et al. (1995) have also demonstrated the use of the WIM technique as a navigation and locomotion technique for VEs. Figure 4.1 shows a WIM and a part of the VE that is represents in the background.

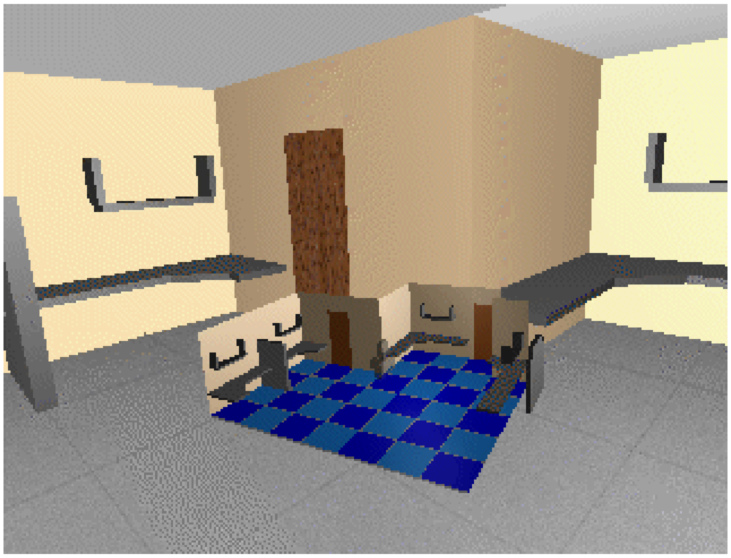

Figure 4.1 WIM against a part of the VE that it represents (Pausch et al. 1995).

4.3 Subjects

This experiment involved 8 users. All of them were unpaid volunteers and students at Virginia Tech. Their average age was 25. There were six males and two females. The user population consisted of three experienced users and five novice users of VEs.

4.4 Apparatus and Implementation

We used the Virtual Research V8 HMD for this experiment. It supports a resolution of 640 x 480 pixels with a 60° diagonal FOV and presented biocular images to the users. The CAVE used in the experiment was 10 x 10 x 10 feet with projections on the floor and front, left and right walls. Each screen had a resolution of 1280 x 1024 pixels. In order to ensure consistency among our displays, the CAVE was set up to display biocular images rather than stereo. We used an Intersense IS900 tracker to track the head and hands of the users. Additionally, we used a testbed implementation designed by Bowman, Johnson and Hodges (2001). The testbed was developed using the SVE library. For the HMD, we ran the testbed on a PC running Windows 2000. The CAVE used a cluster of five PCs each running Redhat Linux 9.0. Moreover, we used 2 wands in the HMD and the CAVE. The wand in the user's dominant hand was used to control the virtual hand while the wand in the non-dominant hand was used as a tracker to which the WIM was attached. The top left button on the wand in the user's dominant hand was used to select and release objects in the WIM and the bottom left button was used to start the next trial once they had completed the current trial. The environment consisted of two rooms, the practice room and the trials room. Users started off the practice room containing pieces of furniture that they could select and manipulate. Figure 4.2 shows the practice room. After spending some time in the practice room, users were taken into the trials room that contained a set of nine cubes and two cylindrical targets. Each trial consisted of selecting a particular cube from the nine cubes and placing it between the target cylinders.

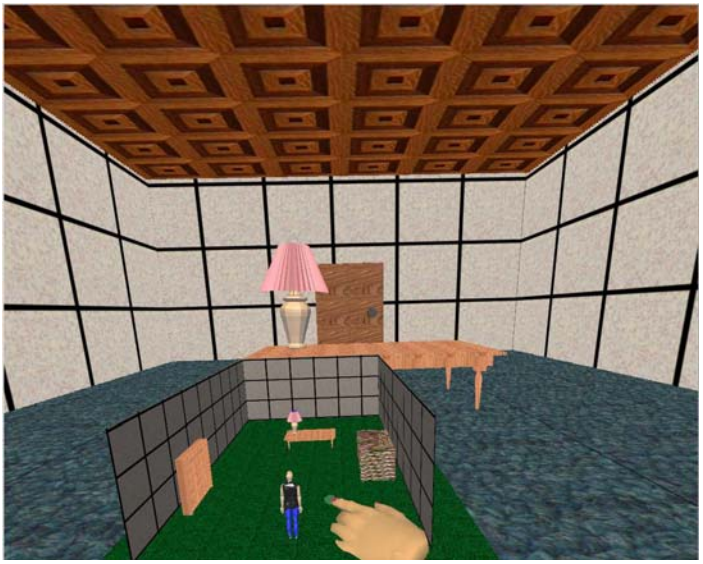

Figure 4.2 The practice room used in this experiment (also showing the WIM)

4.5 Experimental Design

The study consisted of 8 different types of tasks. Each type of task had 4 trials, totaling 32 trials each in the HMD and the CAVE. Like the previous experiment, the order of the displays was counterbalanced. The main independent variable for this experiment was the display type. The dependent variables included selection time (the time it took users to select the blue cube starting from the beginning of the trial), manipulation time (the time it took users to place the cube in between the targets after selecting the cube) and selection errors (errors made while trying to select the cube like selecting the wrong cube). Other dependent variables were the difficulty level ratings that users gave after each task and the comfort level ratings that users gave after the completion of approximately every 8 tasks. These ratings were on a scale of 1 to 10, 1 being least difficult to 10 being most difficult. Also, this was a within-subjects design as all users were made to perform tasks in the HMD as well as the CAVE.

4.6 Procedure

At the start of the experiment, subjects were asked to fill out a pre-questionnaire (see Appendix B1). The pre-questionnaire consisted of questions related to demographic information such as age, gender, handedness and the subject's prior experience with VEs. After that they were asked to go through an instruction sheet explaining the tasks that were to be performed and the WIM technique (see Appendix B2). Next, those unfamiliar with the VE equipment were told about the basic setup and hardware. Also any questions about the hardware and/or the experiment were answered at this stage. Users were then allowed to practice for about 5 to 10 minutes before the start of the trials in the HMD and the CAVE. The testbed includes a practice room with different pieces of furniture that users could select and manipulate. This allowed users to practice the typical selection and manipulation tasks that they were going to perform later in the experiment. Also, at this stage users were free to ask any questions and encouraged to mention any comments they had.

4.6.1 Tasks

The tasks that users were asked to perform consisted of selecting a cube from a set of 9 cubes in the environment and placing it between two target cylinders. The cubes were arranged in a 3 x 3 pattern and the users were asked to select the center cube which was blue in color. The surrounding cubes were grey. The targets cylinders were vertically on top of each other and users were supposed to place the blue cube in the space between them. The set of tasks for this experiment involved 2 variables for the selection part and 1 variable for the manipulation part. The selection variables were size and density of the cubes. The two sizes of cubes were 0.4 meters and 1.0 meter. Figure 4.3 shows the cubes with the two different sizes. The different densities were 0.4 meters and 1.0 meter between the cubes. Figure 4.4 shows the two densities between the cubes. The variable for the manipulation part was ratio of target size to the size of the cube. The two ratios for this variable were 1.5 and 2.5 times the size of the cube. Figure 4.5 shows the two target sizes. Thus, from the combination of the 3 variables, there were 8 different types of tasks and users performed 4 trials of each type. Each trial consisted of selecting the blue cube and placing it between the target cylinders. Users were given feedback by highlighting the virtual hand every time a cube was "selectable" when the index finger of virtual hand coincided with it. Additionally, feedback was also provided in the same manner when the blue cube was in the correct location between the target cylinders. The trials were randomized so that users did not get used to performing one type of task. Additionally, all the task parameters were read from an input file and the system kept track of the selection and manipulation times, selection errors and position errors. This ensured that the recorded times were accurate and the experimenter could devote his attention to taking down notes about the critical observations that he made. Qualitative data like difficulty and comfort level were recorded by the experimenter by asking users after they had completed the trials.

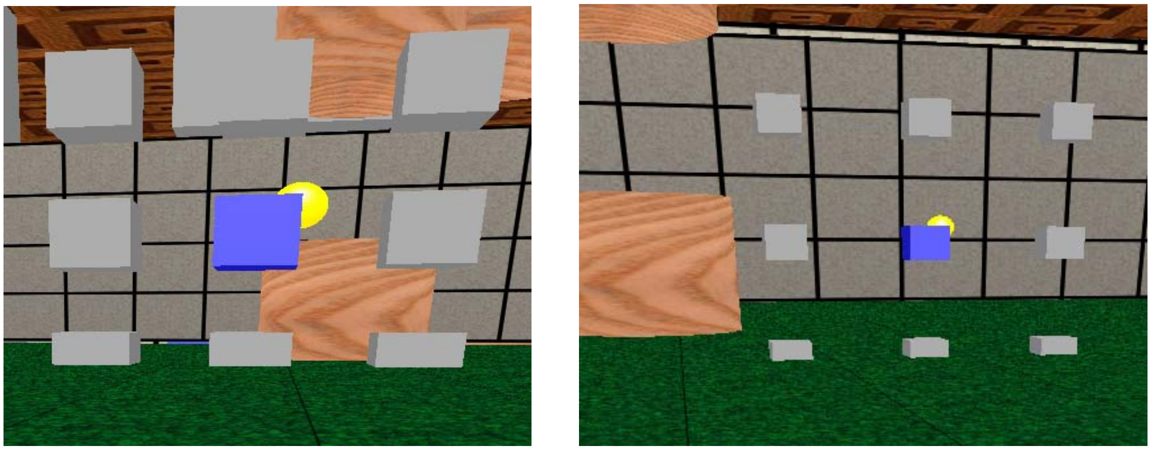

Figure 4.3 The two size conditions used in our experiment, the 1.0 m cubes on the left and the 0.4 m cubes on the right

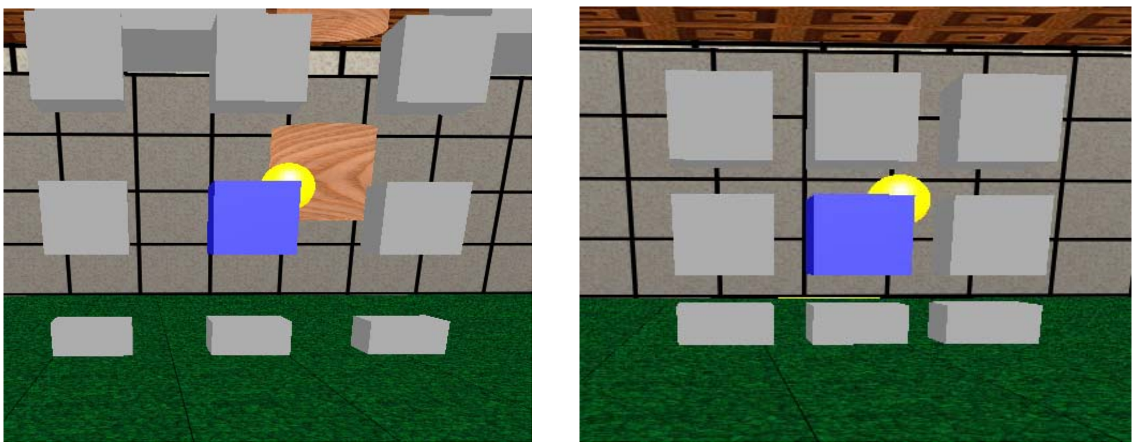

Figure 4.4 The two density conditions used in our experiment, 1.0 m distance on the left and the 0.4 m distance on the right

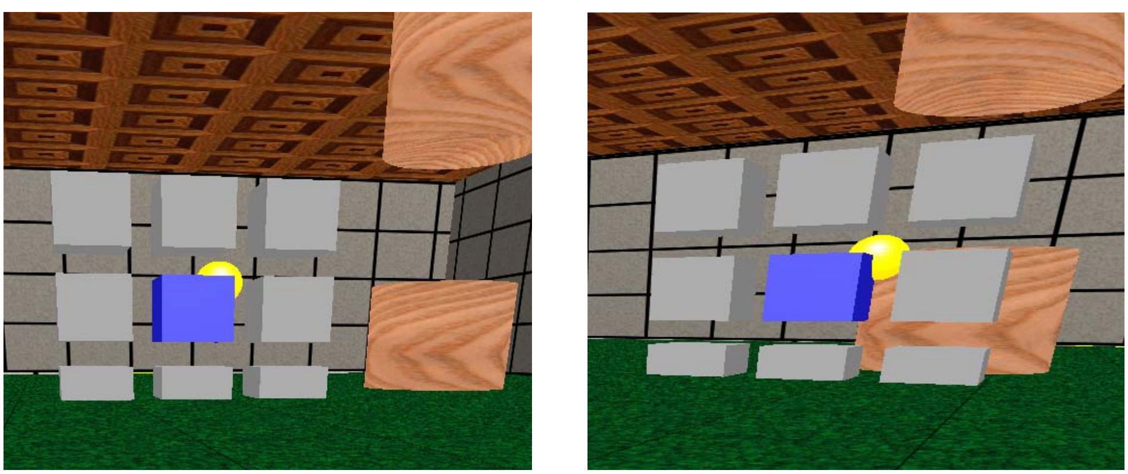

Figure 4.5 The two target size conditions used in our experiment, 1.5x on the left and 2.5x on the right

At the end of the experiment, users were asked to fill out a post-questionnaire (see Appendix B3) asking them about the difficulties they had while selecting and manipulating the blue cube. They were also asked for suggestions as to what changes could be incorporated into the WIM technique in the CAVE in order to make it more usable.

Figure 4.6 A user performing the tasks using the HMD

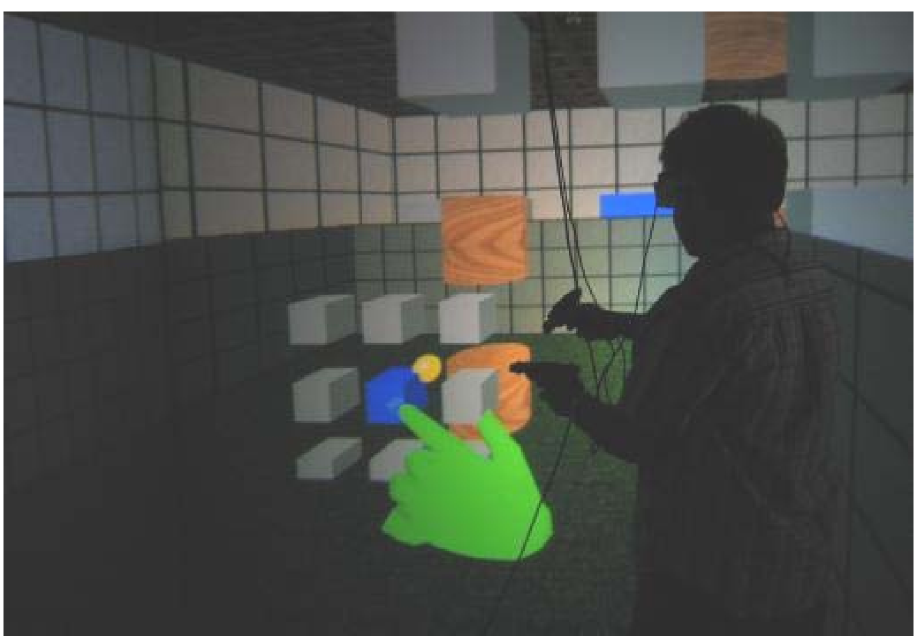

Figure 4.7 A user performing the tasks in the CAVE

4.7 Observations

Since users were timed while performing tasks, they were not asked to provide comments. Thus, most of these observations were made by the experimenter while users performed the tasks.

The first observation was that left and right wands clashed several times while users were trying to select objects in both the HMD and the CAVE. Users generally found this frustrating but found a way to work around it in order to complete the task, however, at the cost of time.

The second observation was that users found it difficult to see the WIM when it was being projected on the floor of the CAVE. This was because the floor of the WIM in the testbed was dark green with some texture. Due to the dark color and poor visual quality of the floor in the CAVE, users had to lift the WIM such that it was projected on one of the screens in order to complete the task.

The third observation was made on a couple of user's comments while they were practicing in the VE. Both users felt that the size of the WIM in the CAVE was larger than in the HMD, although in both cases the WIM was $1/20^{th}$ scale of the virtual room.

4.8 Results and Analysis

We made some interesting observations while users performed tasks in the HMD and the CAVE. In order to extract interesting results, we performed a two-factor with replication analysis of variance (ANOVA) on the selection times, manipulation times and selection errors that the system recorded. For the full results, see Appendix B4.

We found that display type had a significant effect on manipulation time ($p = 0.048$). This supports our hypothesis that the usability of the WIM technique is affected by the display type. The average manipulation time for all users was 3.607 seconds in the HMD and 4.254 seconds in the CAVE.

We did not find any significant results of the display type on selection time or selection errors; however, the average selection time for all users in the HMD (4.580 seconds) was lower than the average selection time for all users in the CAVE (4.826 seconds).

Other results that we found were that size of the object is significant for selection time ($p < 0.001$). Moreover, selection errors were significantly affected by size ($p < 0.001$). This is because larger objects are easier to select. These results conform to previous results by Bowman, Hodges and Johnson (1999, 2001).

We also conducted a correlation analysis on the demographic data from the pre-questionnaire and the average selection times, manipulation times and selection errors for the HMD and the CAVE for all subjects. We found a correlation between previous 3D gaming experience (with games like Doom, Quake, etc.) and the manipulation time in the HMD ($r = -0.712$) and the CAVE ($r = -0.725$). Additionally, we also found a correlation between previous VE experience and selection times in the HMD ($r = -0.891$) and the CAVE ($r = -0.514$). We also found correlations between the selection time in the HMD with manipulation time in the HMD ($r = 0.521$), selection errors in the HMD ($r = 0.775$) and selection time in the CAVE ($r = 0.564$). There were also correlations between manipulation time in the HMD and the CAVE ($r = 0.757$). We did not find any gender based correlations.

4.9 Discussion

From the manipulation time results for all users in the HMD and CAVE, it is evident that the user task performance with the WIM technique was hindered in the CAVE. Combining this result with the user comments about the perceived size of the WIM, we have developed the following explanation for this: in a 3D scene people use a range of depth cues, including pictorial cues, stereopsis, motion parallax and oculomotor cues. In our experimental environment, pictorial cues were not sufficient to determine the depth of the WIM. Stereopsis was not present since we did not use stereo in the CAVE in order to ensure consistency among both displays. In the HMD, due to the limited FOV users were forced to move their heads which gave them better depth perception due to motion parallax. In the CAVE, motion parallax was ineffective because users did not move their heads or the WIM very much during the experiment. This means that users must have relied on oculomotor cues to a provide depth estimate for the WIM in the CAVE. Oculomotor cues are "depth cues derived from muscular tension in the viewer's visual system, called accommodation and convergence. *Accommodation* is the physical stretching and relaxing of the eye lens caused by the eye muscles when focusing on an image…*Convergence* is the rotation of the viewer's eyes so images can be fused together at varying distances" (Bowman et al. 2004). Again, since we did not use stereo, users did not use convergence, meaning most of the depth perception was through accommodation. In the HMD, since the screens are within a few inches off the users eyes, accommodation cues probably made users feel that they had a local WIM that was close to them (in their hand). On the other hand, due to the screens being a few feet away from the users in CAVE, accommodation cues made users perceive that they were working indirectly, on a remote copy of the WIM. This may explain the significant difference in the average manipulation time for all users between the HMD and the CAVE. Note that this effect would have likely not been present had we used stereo in the CAVE. Additionally, the lower visual quality of the floor of the CAVE was also a factor for the lower user task performance in the CAVE. Due to the dark color of the floor of the WIM, users had to lift it to project it on one of the screens every time the WIM was being projected on the floor which increased selection and manipulation time.

4.10 Conclusions

Our primary conclusion is that the WIM technique is somewhat more usable in the HMD than in the CAVE. Thus, in order to make the WIM technique more usable in the CAVE, we propose the following changes based on our observations and experimental results:

1. Scale down the WIM, the virtual hand and its motion such that it appears to be the same size as in the HMD.
2. Eliminate any dark colors and textures from the WIM such that users can see and select and manipulate objects with ease even when the WIM is being projected on the floor of the CAVE.

3. Train users to use motion parallax by moving their head and/or the WIM or walking a couple of steps to view the object to be selected from multiple viewpoints in order to get better depth perception.

We believe that the above changes will make the WIM technique in the CAVE more usable. Our next step would be to make these changes to the current testbed implementation and conduct another round of usability evaluation in order to determine that the changes we proposed do have a positive effect on the usability of the WIM technique in the CAVE.

Chapter 5 – Verification of the Usability of the CAVE-specific Interaction Technique

In the previous chapter, we discussed the experiment whose purpose was to show that the WIM technique, which was designed specifically for the HMD, needs to be modified in order to provide the same level of usability in the CAVE. This chapter describes the experiment that we conducted with users with the modified WIM technique in the CAVE.

5.1 Our CAVE-specific WIM Technique

Based on the results from the previous experiment, we made four modifications to the WIM technique to increase its usability in the CAVE. The first modification involved adding different vertical offsets to the virtual hand and the WIM from their tracked position. This was done to minimize clashing between the two wands when users performed tasks on the WIM. Note that this was also a problem in the HMD, and therefore this change should also be implemented in HMD-based WIMs.

The second modification involved changing the brightness of the floor of the (practice and trial) rooms and the WIM. The colors of the rooms and the WIM were changed from dark blue to grey and from dark green to light green respectively.

The third modification involved changing the scale of the WIM and virtual hand. The scale of the original WIM was $1/20^{th}$ the size of the virtual room. The improved WIM and virtual hand were $3/4^{th}$ of their original size (the improved WIM was approximately $1/27^{th}$ the size of the virtual room). We also scaled down the hand motion by the same amount so that intricate selection and manipulation is not affected with the smaller WIM.

The last modification we made was to train users to use motion parallax to get better depth perception while selecting and placing objects. Users were taught to move their heads and/or the WIM to get more than one viewpoint on the object of interest. They were also taught to walk a couple of steps in different directions, again to obtain more than one viewpoint.

5.2 About the experiment

5.2.1 Purpose

The purpose of this experiment was to evaluate that the changes we made to the WIM technique in the CAVE have a positive effect on usability. Our hypothesis for this experiment was that the modifications to the WIM technique do indeed have a positive effect on user task performance. Based on the results of this experiment, we can posit guidelines to migrate selection and manipulation techniques from the HMD to the CAVE while maintaining usability.

5.2.2 Brief Outline of the Experiment

Like the previous experiment, this too was a formal, controlled experiment where users were timed for each task they performed in order to gather some quantitative data. Users were asked to perform the same tasks using the WIM technique as the previous experiment: selecting the blue cube and placing it between the target cylinders. Users were asked to perform the same tasks as the previous experiment because it would help us compare the data collected from both the experiments and determine if the modifications we made indeed affected user task performance.

5.3 Subjects

This experiment involved 12 users. All of them were unpaid volunteers and students at Virginia Tech. Their average age was 24. There were ten males and two females. The user population consisted of four experienced users and eight novice users of VEs.

5.4 Apparatus and Implementation

The purpose of this experiment was to evaluate that the changes we made to the original WIM technique in the CAVE did have a positive effect on its usability. Thus, we used the same apparatus as the previous experiment. The CAVE used in the experiment was 10 x

10 x 10 feet with projections on the floor and front, left and right walls. Each screen had a resolution of 1280 x 1024 pixels and displayed biocular images. The wand in the user's dominant hand was used to control the virtual hand while the wand in the other hand was used as a tracker to which the WIM was attached. Again, we used the Intersense IS900 tracker to track the head and hands of the users. We also used the testbed implementation designed by Bowman, Johnson and Hodges (2001). However, the testbed implementation was modified to change the color and scale of the WIM on different keystrokes. This would enable us to make modifications on the fly while conducting the experiment. One other modification was that we added different vertical offsets for the virtual hand and the WIM in order to minimize clashing of the two wands.

5.5 Experimental Design

As opposed to the previous two experiments, this was a between subjects experiment. We had 4 groups of users, namely Groups A, B, C and D. All groups consisted of 3 users, 1 experienced and 2 novice users of VEs. All subjects performed 32 tasks in the CAVE. The tasks were identical to the previous experiment. However, there was a slight modification. All subjects in groups A, B, and C were asked to perform 16 tasks with the original WIM and an additional 16 with the improved WIM. The improvements varied by groups. Group A subjects had the improvement in terms of brightness. Figure 5.1 shows the original WIM and the improved WIM. Group B subjects had the improvement in terms of the scale of the WIM, virtual hand and its motion. Figure 5.2 shows the original WIM and the improved, scaled down WIM. Group C subjects had the improvement in terms of training. All these subjects were trained to use motion parallax for better depth perception. Group D subjects had all three improvements, brightness, scale and training, and performed tasks with 8 different versions of the WIM (see table 5.2). The main independent variables were the display type and the improvement we made to the WIM. The dependent variables included the difficulty level and comfort level ratings, selection times and manipulation times. We did not test the offsets to the trackers of the WIM and the virtual hand independently because it was not a CAVE-specific improvement.

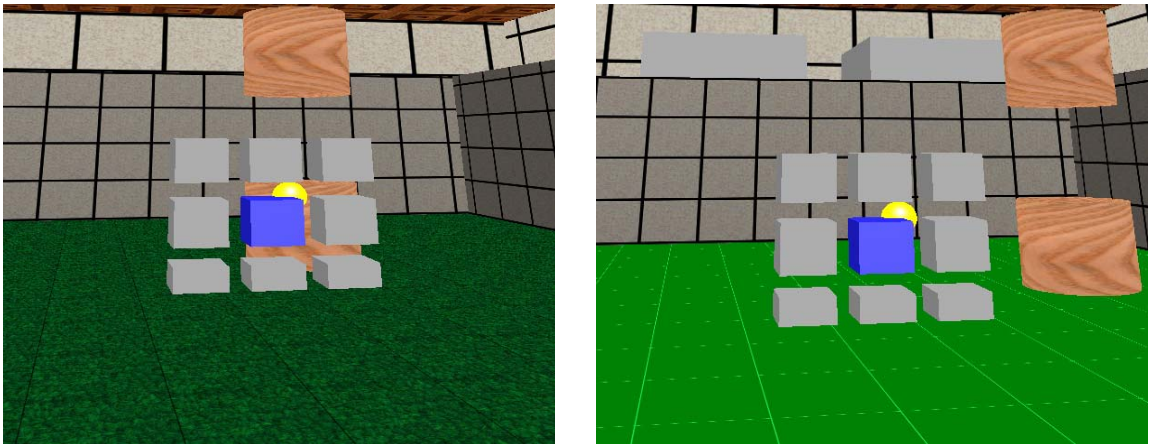

Figure 5.1 The original WIM (left) and the improved WIM with brightness change (right)

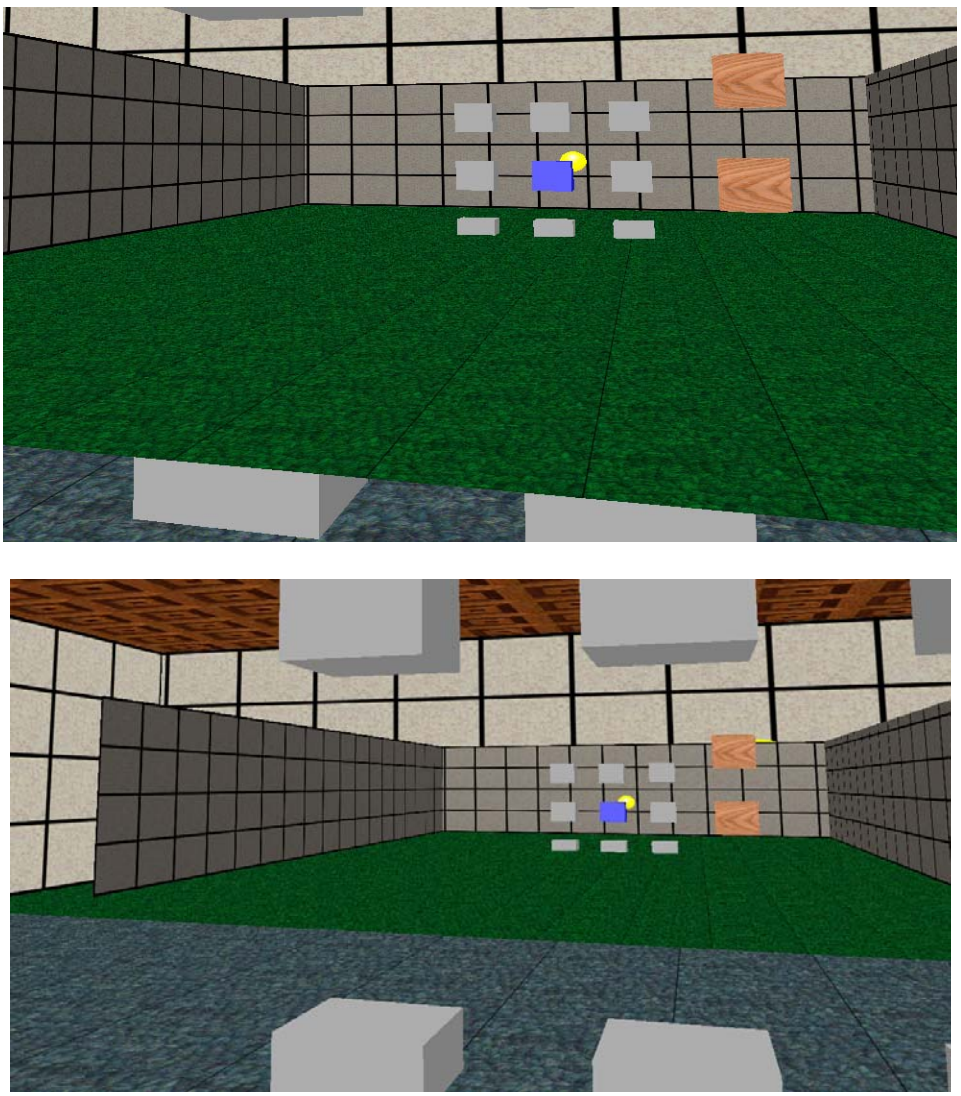

Figure 5.2 The original WIM (above) and the improved scaled down WIM (below)

5.6 Procedure

At the start of the experiment, subjects were asked to fill in a pre-questionnaire (see Appendix C1) that contained demographic information such as age, gender, right handedness or left handedness and the subject's prior experience with VEs. We did not test subjects for color blindness. After that they were asked to read an instructions sheet (see Appendix C2) explaining the tasks that were to be performed and the WIM technique. Next, those unfamiliar with the VE equipment were told about the basic setup and hardware. Also any questions about the hardware and/or the experiment were answered at this stage. Users were then allowed to spend about 5 to 10 minutes in the practice room before the start of the trials. This allowed them to familiarize themselves and to identify ways they could interact with the environment. Users were free to ask questions at this stage. They were encouraged to mention any comments they had at this stage.

5.6.1 Tasks

The set of tasks for this experiment involved 1 variable each for the selection and manipulation part. We did not include the density variable for this experiment because it did not have any significant effects on selection and manipulation times in the previous experiment. The selection variable was size of the cubes while the manipulation variable was the relative target size. The two sizes of cubes were 0.4 meters and 1.0 meter. The two ratios for the relative target size were 1.5 and 2.5 times the size of the cube (see figures 4.3 and 4.5). Density was kept constant with 1 meter between the cubes. Thus, the combination of these two variables combined with the third variable (brightness, scale or training) leads to 8 different combinations. Users in Groups A, B and C performed 4 trials for each combination of the variables, leading to a total of 32 trials. Users in Group D performed tasks for all improvements to the WIM with only the selection variable. The target size was kept constant at 1.5 times the size of the cube. Thus they had 4 variables in all (size of cube, brightness, scale and training) and performed 2 trials for each combination of the variables. Each trial consisted of selecting the blue cube and placing it between the target cylinders. Again, users received feedback through the highlighting of the virtual hand when it coincided with a cube and when the cube was in the right location between the target cylinders. Table 5.1 shows the number of tasks for each condition for subjects in Groups A, B and C. Table 5.2 shows the number of tasks for each condition for subjects in Group D.

	Original WIM	Improved WIM
Size1 (0.4 m cubes)		
Target1 (1.5x)	4	4
Target2 (2.5x)	4	4
Size2 (1.0 m cubes)		
Target1 (1.5x)	4	4
Target2 (2.5x)	4	4

Table 5.1 Number of trials for each condition for subjects of Groups A, B and C

	None	Brightness	Scale	Training	Brightness + Scale	Brightness + Training	Scale + Training	All
Size1 (0.4 m cubes)	2	2	2	2	2	2	2	2
Size2 (1.0 m cubes)	2	2	2	2	2	2	2	2

Table 5.2 Number of trials for each condition for subjects of Group D

Users of Groups A, B and C randomly started with the original WIM or the improved WIM to ensure that they did not always perform better on improved WIM due to practice from tasks on the original WIM. Additionally, the trials for all users were randomized so that users did not get used to performing one type of task. All the task parameters were read from an input file and the system kept track of the selection and manipulation times, selection errors and position errors. This ensured that the recorded times were accurate. Subjective data like difficulty and comfort level were recorded by the experimenter by asking users after they had completed the trials.

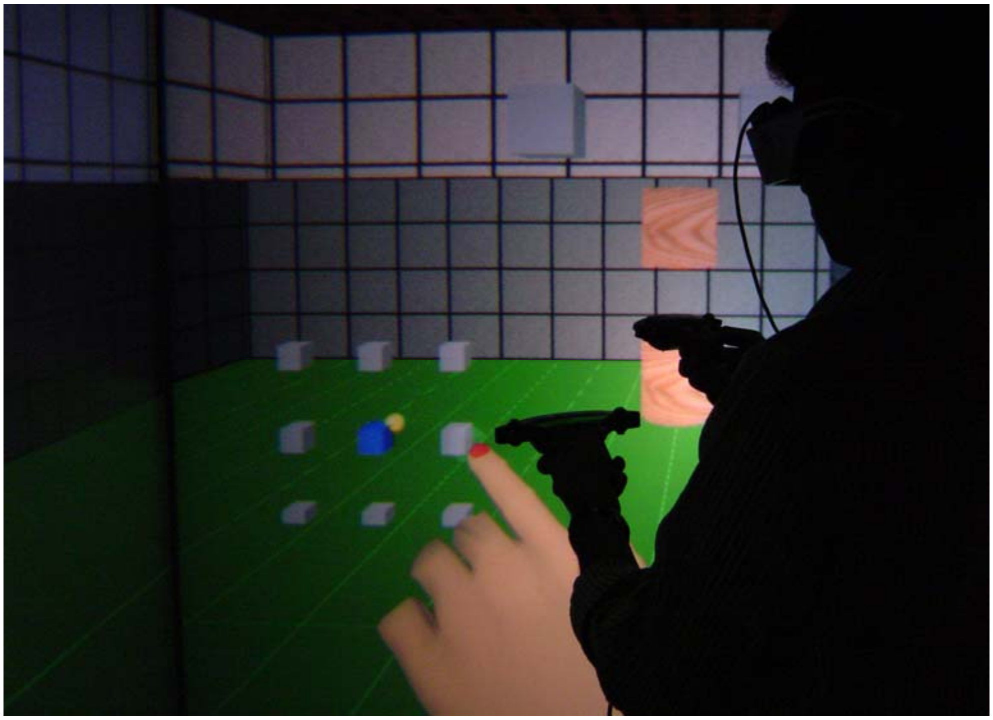

Figure 5.3 A subject of Group A performing a trial with the improved color of the WIM in the CAVE

At the end of the experiment, all users were asked to fill out a post-questionnaire (see Appendix C3) asking them about the difficulties they had while selecting and manipulating the blue cube before and after the improvement. Users of Groups A, B and C were asked specific questions related to the improvement that they got. Group D users were required to answer all questions related to Groups A, B and C since they got all improvements.

5.7 Observations

Like the previous experiment, users were not asked to provide comments while performing tasks. However, while observing users performing task, we made an important observation. None of the users complained of the wands clashing while performing the trials. One user did complain of the wands clashing while practicing however the same was not the case while performing trials.

5.8 Results and Analysis

In order to determine whether the modifications we made to the WIM technique in the CAVE had effects on user task performance, we conducted ANOVAs on data like selection times, manipulation times and selection errors that we had collected. For Groups A, B and C we conducted a two-factor ANOVA with replication while for Group D we conducted a two-factor ANOVA for selection times and errors and a single-factor ANOVA for manipulation times. For full results, see Appendix C4. We found some interesting results which are discussed below.

We found a non-significant trend for an interaction between improvement and size ($p = 0.095$) for group A, and that the brightness change helps with the smaller cube size. Examining the average selection times in this group, we found that all three subjects improved their average selection time for the smaller cubes when using the improved brightness WIM. Figure 5.4 illustrates this fact.

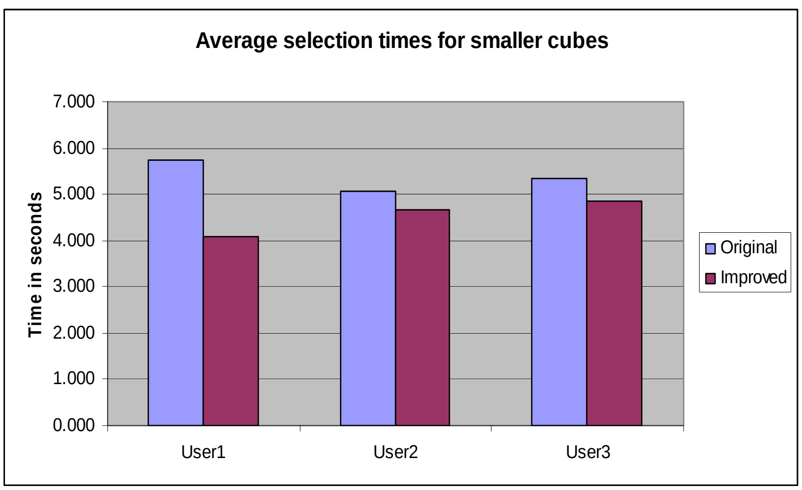

Figure 5.4 Graph showing the average selection times for all users for the smaller cubes with the improved brightness WIM

The brightness change did not have any significant effect on the manipulation times or selection errors.

The scale change appears to have no effect on the selection times, manipulation times and selection errors. We presume this is the case since the hand motion was scaled down from the actual motion of the user's physical hand. Thus, most users tended to undershoot the target when attempting to touch or place an object in the WIM. It was later that they realized that they had to move their physical hand even further to touch or place the object, which took additional time.

The training change does not show any statistical significance on selection and manipulation times or selection errors. However, all subjects in this group improved on the selection time with the training change. The average selection time for all three subjects went down from 4.476 seconds to 3.750 seconds with the improvement. Ironically, manipulation times generally increased with the training change. The average manipulation time for all users was 4.347 seconds with the original WIM and 5.449 seconds with the training change. This can be explained by the response of one of the subjects in the post-questionnaire. He felt that he generally took more time to complete the trials while using the motion parallax training because he had to move his head and/or the WIM to get better depth perception. Additionally, he also felt that sometimes the opposite motions of his head and the WIM confused him. A couple of users also commented that that they did not necessarily need to use the motion parallax training because the feedback through the hand highlighting was sufficient for them to judge that the cube was in the right location. One of these users also mentioned that if there was no such feedback then the training would have been really helpful to place the cube precisely.

We did not find any significant differences between the various combinations of improvements on selection and manipulation times for the subjects of Group D, who had all the improvements.

5.9 Conclusions

Based on the results from this experiment, we can tentatively conclude that some of the improvements to the WIM technique have positive effects on user task performance in the CAVE. We did not find any statistically significant results, so more research is needed. Our opinion is that statistical significance would be achieved with a greater number of subjects (greater power in the experimental design), but of course this is only speculation. We did find a non-significant trend for an interaction between the brightness improvement and size on selection time ($p = 0.095$). Also, the brightness improvement helped reduce the selection times especially for the smaller cubes. This suggests that the brightness improvement did help improve the usability of the WIM technique in the CAVE. The training change did not have any significant effects on selection and manipulation times. However, selection times for all users decreased with the training change. Additionally, from users' comments in the post-questionnaire it was evident that training would be really helpful when there is no feedback while placing objects in the target location. This result suggests that the motion parallax training helped improve the usability of the WIM technique in the CAVE.

Chapter 6 - Conclusions and Future Work

This thesis is part of an initiative to explore how differences between HMDs and CAVEs affect user interaction in a 3D environment.

In this chapter, we summarize our preliminary answers to the questions we pondered at the beginning of the thesis. We also recapitulate some of the results of the previous experiments. Additionally, we present our contributions to this field and suggest directions for future work that came out as a result of our research.

6.1 Conclusions

In the first chapter, we mentioned the problem statement:

> *Does the choice of a VE display have an effect on 3D interaction? Can we migrate HMD-based selection and manipulation techniques to the CAVE? If so, will there be any tradeoffs in user performance when this is done? Can we develop generalized guidelines to do this without compromising user performance?*

Based on the results of the three rounds of experiments, we have gathered evidence in favor of particular answers. The results of our first experiment strongly suggest that the choice of display does have an important effect on 3D interaction. We had users perform different selection and manipulation tasks and saw that the differences between the HMD and the CAVE had varying effects on different types of tasks. Users preferred the HMD to select objects at height and place objects behind them in the environment and the CAVE for wide field tasks like placing objects from left to right. Observations and difficulty ratings also supported these choices.

In the second and third experiment, we migrated the WIM technique, an HMD-based selection and manipulation technique, to the CAVE. Thus, we can say that it is possible to migrate HMD-based selection and manipulation techniques to the CAVE.

The results of the second experiment showed that display type significantly affected manipulation time with the WIM technique. Also, the average selection time for all users in the HMD was 4.580 seconds compared to the average selection time in the CAVE of 4.826 seconds. These results demonstrate that there is a tradeoff in user performance when a display-specific interaction technique is migrated to a display for which it is not intended.

In order to answer the last question of our problem statement, we need to briefly review the second and third experiments. The second experiment compared the WIM technique in the HMD and CAVE for selection and manipulation tasks. We kept the tasks and all other variables identical so we could evaluate how different display types affect interaction. Based on user comments and results from the second experiment, we made three improvements to the WIM technique and tested it, this time only in the CAVE. Users were tested with one improvement individually or all of the improvements so we could determine if each improvement and also all improvements in conjunction had positive effects on user interaction. The results of the third experiment suggest that the brightness and training improvement might have a positive effect on the usability of the WIM technique in the CAVE.

Recapitulating the results of previous experiments, we posit four guidelines:
In the first experiment, we demonstrated that users preferred the HMD over the CAVE for selecting objects at a height. Based on this result, we posit the following guideline:

> *In the CAVE, either avoid selection and manipulation high above the user or provide simple navigation techniques to allow the user to fly up to the higher location.*

In the first experiment we also demonstrated that users preferred the HMD for placing objects behind them in the environment. This is not possible in the CAVE unless the application provides this feature, which leads us to our second guideline:

> *Provide a simple or automatic way for users to rotate the world when performing manipulation tasks in a CAVE.*

The results of our third experiment suggest that the brightness improvement helps improve the selection time in the CAVE due to the lower visual quality of the floor. This is especially true for smaller objects in the VE. We posit our third guideline based on this result:

> *In the CAVE, due to the lower visual quality of top projected floors (compared to the screens), use bright colors (like grey or light green) without any texture for the floor of the VE.*

The third experiment results also suggested that motion parallax training helped improve the selection times in the CAVE. Additionally, in the post-questionnaire one of the subjects commented that the motion parallax training would be really helpful in environments where there is no feedback for selection and manipulation. Thus, we posit our fourth guideline on this result:

> *For tasks requiring accurate selection and manipulation, train users to use motion parallax in the CAVE so that they view the object from multiple viewpoints and get better depth perception.*

However, keep in mind that this result may be solely due to our use of a non-stereo CAVE in our experiments. More research is needed to determine whether users' depth perception for such tasks would be remedied through the use of stereo.

6.2 Contributions

In addition to the above conclusions, we have the following contributions:

1. **Demonstrated that display type does affect 3D interaction.**

   In the first experiment, we had users perform different types of selection and manipulation tasks in the HMD and the CAVE. We determined that users preferred a particular display type over the other in terms of usability for three of these tasks. Thus, we demonstrated that display type does affect 3D interaction.

2. **Demonstrated that an interaction technique designed for one display is less usable in another display.**

   In the second experiment, we had users perform identical selection and manipulation tasks using the WIM technique in the HMD and the CAVE. The WIM technique is a HMD-based technique. From the user's responses from the post-questionnaire and the results from the ANOVA, we established that the WIM technique was less usable in the CAVE compared to the HMD.

3. **Demonstrated that redesigning that technique based on the properties of the second display may result in higher levels of usability.**

   In the third experiment, we made improvements to the WIM technique in the CAVE-based on results from our second experiment. The results of this experiment suggested that some of the improvements resulted in higher levels of usability of the WIM technique in the CAVE.

4. **Presented guidelines to migrate HMD-based selection and manipulation techniques to CAVE.**

   In the earlier section of this chapter, we presented a set of guidelines some of which can be used by future VE application developers to easily migrate their applications involving selection and manipulation from HMDs to CAVEs with minimal loss of usability.

6.3 Future Work

During the course of this thesis, we noticed interesting areas that are relevant to this work, but beyond the scope of this thesis.

1. **Compare the improved WIM technique in the CAVE with the original WIM technique in the HMD.**

   In this thesis, we only compared the improvement to the WIM technique with the original WIM technique in the CAVE that we migrated from the HMD. It would be interesting to study how the improved WIM technique in the CAVE compares to the original WIM technique in the HMD.

2. **Study the effects of the differences between the HMD and the CAVE with another HMD-based selection and manipulation technique.**

   In experiment two, we had users perform tasks using the WIM technique. However, it would be interesting to learn how the usability of different selection and manipulation techniques like HOMER would be affected by the differences between HMDs and the CAVEs.

3. **Study the effects of the difference between the HMD and the CAVE with travel techniques.**

   This thesis concentrated only on the selection and manipulation aspect of VEs. It would be interesting to know the usability of travel techniques would be affected by the differences between HMDs and CAVEs. As a subset, it would also be interesting to learn how the usability of the WIM technique used as a locomotion technique would be affected when migrated to the CAVE.

4. **Study effects of other differences between the HMD and the CAVE on the usability of 3D interaction techniques.**

   In this thesis, we only studied the effects of a few differences between the HMD and CAVE on selection and manipulation. It would be interesting to study the effects of other differences like brightness, resolution, etc. between these displays on other 3D interaction techniques.

5. **Study the effects of differences between other VE displays on 3D interaction.**
   In this thesis, we only studied the effects of the differences between the HMD and the CAVE on 3D interaction. There is very little work done comparing other VE displays like workbench displays, hemispherical displays, monitors, etc. and the effects of their differences on 3D interaction.

6. **Development of a methodology for migration of interaction techniques from one display to another.**
   Earlier in this chapter, we presented a few guidelines to migrate HMD-based selection and manipulation techniques to the CAVE. It would be interesting to develop a structured methodology for the smooth migration of any 3D interaction technique from one display type to another.

7. **Development of a set of guidelines for choice of display for particular interaction techniques.**
   In the second chapter, we discussed several VE displays. It would be interesting to learn how different 3D interaction techniques map to these displays and develop a set of guidelines for the choice of display for each of these interaction techniques.

References

Arthur, Kevin, *Effects of Field of View on Performance with Head-Mounted Displays,* doctoral dissertation, Dept. of Computer Science, University of North Carolina, Chapel Hill, 2000.

Benokraitis, A., Kriz,R., and Hix, D., “VT CAVE Information Center”, May 1998; http://www.cave.vt.edu

Bolt, R., “Put-That-There": Voice and Gesture at the Graphics Interface. *Computer Graphics, 14(3), July 1980, Proc. ACM SIGGRAPH*, 1980, pp. 262-270.

Bowman, D. and Hodges, L., “An Evaluation of Techniques for Grabbing and Manipulating Remote Objects in Immersive Virtual Environments”, *Proc. 1997 Symposium on Interactive 3D Graphics*, 1997, pp. 35-38.

Bowman, D., Johnson, D., and Hodges, L., “Testbed Evaluation of VE Interaction Techniques,” *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, 1999, pp. 26-33.

Bowman, D., *Interaction Techniques for Common Tasks in Immersive Virtual Environments: Design, Evaluation, and Application,* doctoral dissertation, Dept. of Computer Science, Georgia Tech, Atlanta, 1999.

Bowman, D. and Wingrave, C., “Design and Evaluation of Menu Systems for Immersive Virtual Environments,” *Proceedings of IEEE Virtual Reality*, 2001, pp. 149-156.

Bowman, D., Johnson, D., and Hodges, L., “Testbed Evaluation of Virtual Environment Interaction Techniques,” *Presence: Teleoperators and Virtual Environments*, vol. 10, no. 1, 2001, pp. 75-95.

Bowman, D. and Wingrave, C., “Design and Evaluation of Menu Systems for Immersive Virtual Environments,” *Proceedings of IEEE Virtual Reality*, 2001, pp. 149-156.

Bowman, D., Datey, A., Ryu, Y., Farooq, U., and Vasnaik, O., “Empirical Comparison of Human Behavior and Performance with Different Display Devices for Virtual Environments,” *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 2002, pp. 2134-2138.

Bowman, D., Setareh, M., Pinho, M., Ali, N., Kalita, A., Lee, Y., Lucas, J., Gracey, M., Kothapalli, M., Zhu, Q., Datey, A., and Tumati, P., “Virtual-SAP: An Immersive Tool for Visualizing the Response of Building Structures to Environmental Conditions,” *Proceedings of IEEE Virtual Reality*, 2003, pp. 243-250.

Bowman, D., Kruijff, E., LaViola, J., and Poupyrev, I., *3D User Interfaces: Theory and Practice*. Addison-Wesley, Boston, to appear July 2004.

Bricken, M., "Virtual Reality Learning Environments: Potentials and Challenges," Human Interface Technology Laboratory, University of Washington, Seattle, WA, 1991.

Brooks, Frederick, Jr., "What's Real About Virtual Reality?" Computer Graphics and Applications, IEEE, Volume: 19, Issue: 6, Nov-Dec 1999 pp.16-27.

Bungert, C., "HMD/Headset/VR-helmet Comparison Chart", Feb. 2004; http://www.stereo3d.com/hmd.htm

Chung, J., "A Comparison of Head-Tracked and Non-head-tracked Steering Modes in the Targeting of Radiotherapy Treatment Beams," *Proceedings of the Symposium on Interactive 3D Graphics*, 1992, pp. 193-196.

Cruz-Neira, C., Sandin, D., and DeFanti, T., "Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE," *Proceedings of SIGGRAPH*, in *Computer Graphics*, 1993.

Dolezal, H., *Living in a world transformed: Perceptual and performatory adaptation to visual distortion*. New York, NY: Academic Press, Inc. 1982

Flasar, Jan. *3D Interaction in Virtual Environment*,. Botanická, Czech Republic, 2000. http://www.cg.tuwien.ac.at/studentwork/CESCG/CESCG-2000/JFlasar/

Iwata, H. and Fujii, T., "Virtual Perambulator: A Novel Interface Device for Locomotion in Virtual Environment," *Proceedings of the IEEE Virtual Reality Annual International Symposium*, 1996, pp. 60-65.

Kjeldskov, J., "Combining Interaction Techniques and Display Types for Virtual Reality," In Proceedings of OzCHI 2001, Perth, Australia. Churchlands, Edith Cowan University Press, 2001.

Krüger, W., Fröhlich, B., "The Responsive Workbench," IEEE Computer Graphics and Applications, pp 1215, May 1994.

Mercurio, P., Erickson, T., Diaper, D., Gilmore, D., Cockton, G., and Shackel, B., "Interactive Scientific Visualization: An Assessment of a Virtual Reality System," *Proceedings of INTERACT*, 741-745, 1990.

Mine, M., *Virtual Environment Interaction Technique*, Tech. report TR95-018, Department of Computer Science, University of North Carolina, Chapel Hill, 1995.

Neale, D. *Head-Mounted Displays: Product Reviews and Related Design Considerations*, Hypermedia Technical Report HCIL-98-02, Human-Computer Interaction Laboratory, Department of Industrial and Systems Engineering, Virginia Tech, Blacksburg, VA, 1998.

Nilan, M. S., Silverstein, J. L., and Lankes, R. D., “The VR Technology Agenda in Medicine.” Virtual Reality 93: Special Report, 1993, pp.33-37.

Paush, R., Burnetter, T., Brockway, D., and Weiblen, M., “Navigation and Locomotion in Virtual Worlds via Flight into Hand-Held Miniatures,” *ACM SIGGRAPH `95 Conference Proceedings, Computer Graphics,* July 1995.

Pierce, J., Stearns, B., and Pausch, R., “Voodoo Dolls: Seamless Interaction at Multiple Scales in Virtual Environments,” *Proceedings of the 1999 Symposium on Interactive 3D Graphics*, 1999, pp.141-145.

Poupyrev, I., Billinghurst, M., Weghorst, S., Ichikawa, T., “Go-Go Interaction Technique: Non-Linear Mapping for Direct Manipulation in VR,” *Proceedings of UIST'96,* 1996 pp. 79-80.

Poupyrev, I., Weghorst, S., Billinghurst, M., and Ichikawa, T., “A Framework and Testbed for Studying Manipulation Techniques for Immersive VR,” *Proceedings of the ACM Symposium on Virtual Reality Software and Technology*, 1997, pp. 21-28.

Rheingold, H. *Virtual Reality*. Summit, New York, 1991.

Robinett, W. and Holloway, R., “Implementation of Flying, Scaling, and Grabbing in Virtual Worlds.,” *Proceedings of the ACM Symposium on Interactive 3D Graphics,* 1992, pp. 197-208.

Schmitz,B., "Virtual Reality: On the Brink of Greatness." Computer-Aided Engineering, April 1993, pp.26-32.

Setareh, M., Bowman, D., and Tumati, P., “Development of a Collaborative Design Tool for Structural Analysis in an Immersive Virtual Environment,” *Proceedings of the International Building Performance Simulation Association Conference*, 2001.

Slater, M., Usoh, M., and Steed, A., “Taking Steps: The Influence of a Walking Technique on Presence in Virtual Reality,” *ACM Transactions on Computer-Human Interaction, 2*(3), 1995, pp. 201-219.

Stoakley, R., Conway, M., and Pausch, R., “Virtual Reality on a WIM: Interactive Worlds in Miniature,” *Proceedings of CHI*, 1995, pp. 265-272.

Sutherland, I., “A head-mounted three dimensional display,” *Proceedings of Fall Joint Computer Conference*, *33,* 1968, pp. 757-764.

Ware, C. and Jessome, D., “Using the Bat: a Six-Dimensional Mouse for Object

Placement," *IEEE Computer Graphics and Applications*, *8*(6), 1988, pp. 65-70.
Ware, C. and Osborne, S., "Exploration and Virtual Camera Control in Virtual Three Dimensional Environments," *Proceedings of the ACM Symposium on Interactive 3DGraphics*, in *Computer Graphics*, *24*(2), 1990, pp. 175-183.

Ware, C., Arthur, K., & Booth, K. "Fish tank virtual reality" Proceedings of the SIGCHI conference on Human factors in computing systems, 1993, pp. 37-42.

Zeleznik, R., LaViola, J., Acevedo, D., and Keefe, D. "Pop Through Buttons for Virtual Environment Navigation and Interaction," In the Proceedings of Virtual Reality 2002, 127-134, March 2002.

Appendix A

A1: Pre-questionnaire for Experiment 1

1. Age:
2. Gender: Male Female
3. Left-handed or right-handed? Left handed Right handed
4. Do you have perfectly uncorrected vision or wear glasses/contacts?

   Uncorrected vision Wear glasses/contacts
5. Do you have a background in architecture or construction?

   Yes No
6. How often do you play 3D games like Doom, Quake, etc.?

   Never Occasionally Often
7. Would you consider yourself a beginner or an experienced user with Virtual Reality devices like the CAVE and HMD?

   Beginner Experienced
8. Which of the following interaction technique(s) have you used previously? Circle **all** that apply.

   Go-go Ray-casting HOMER None

A2: User Tasks for Experiment 1

HMD Tasks:

1. Look at the 10 floor structure in front of you. Now using your virtual hand, select the column on the 6^{th} floor that is closest to you. You will see that the column is highlighted once it is selected
2. Approach the arch in front of you. Using your virtual hand, select each piece and place it behind you to build a new arch. The new arch should be a "mirror reflection" of the old arch.
3. Rotate the world to such that the new arch you built is on your left. Select each piece again and build a new arch on the right such that the new arch is a "mirror reflection" of the old arch.
4. Place the beams placed in front of you to create a square around your feet such that in the end you are in the center of the square.

CAVE Tasks:

1. Look at the 10 floor structure in front of you. Now using your virtual hand, select the column on the 6^{th} floor that is closest to you. You will see that the column is highlighted once it is selected
2. Approach the arch in front of you. Using your virtual hand and the ring menu, select each piece and place it behind you to build a new arch. The new arch should be a "mirror reflection" of the old arch.
3. Rotate the world to bring the new arch that you built on the left screen. Select each piece again and build a new arch on the right screen again such that the new arch is a "mirror reflection" of the old arch.
4. Place the beams placed in front of you to create a square around your feet such that in the end you are in the center of the square.

A3: Post-questionnaire for Experiment 1

1. For each task below, circle whether you preferred the HMD or the CAVE in tersm of ease of use.

   a. Selecting objects at a height. HMD CAVE
   Why?

   b. Placing objects behind you in the environment. HMD CAVE
   Why?

   c. Placing objects from left to right in the environment. HMD CAVE
   Why?

   d. Placing objects close to you in the environment. HMD CAVE
   Why?

2. What were other difficulties you had while selecting and placing objects using the Go-go technique? Give specific instances.

Appendix B

B1: Instructions for Experiment 2

Introduction & Instructions
Selection / Manipulation Experiment:

Welcome, and thank you for agreeing to participate in our study. We are studying the effect of two different displays, namely the Cave Automatic Virtual Environment (CAVE) and the head-mounted display (HMD), for selecting and manipulating objects in a three-dimensional virtual environment. Through this experiment, you will be helping us to better understand the effect of these displays on a selection and manipulation technique. You will be asked to perform some tasks in a 3D virtual environment, using the HMD and the CAVE and a button device called the wand for input in each. We may also ask you some questions about how you liked the technique. It is very important that you remember that <u>we are testing these techniques - we are not testing you.</u>

This experiment will consist of some practice time at the beginning to familiarize you with the selection and manipulation technique, and then 32 experimental trials using the technique. You should be able to finish the entire experiment in 30 to 45 minutes.

Your task will be to pick up a specified object and place it in a target location. You will see an array of nine cubes. The eight cubes on the outside of the array will be gray, while the center cube will be blue. The blue cube is the one you are to pick up. After you have grabbed this blue cube, you will try to place it between two target indicators. The target indicators are cylinders, you just need to place the object between them and do not have to match the orientation. In all cases, the cube you have picked up will change colors when you have it in the correct position. A trial will not end until you have selected the center cube and placed it in the target.

Select the blue cube

Place it between the target indicators

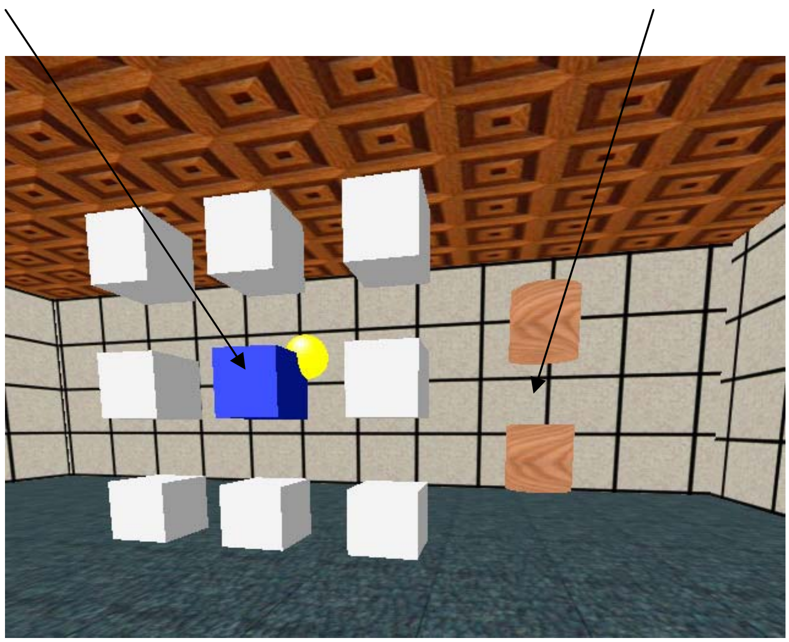

The overall score for each trial will depend on several factors:

- the time you take to grab the blue cube
- the number of incorrect cubes you grab before the blue cube
- the time you take to position and drop the blue cube

Please ask any questions you have before you begin the experiment. If you have questions during a trial, please wait until the trial is over before asking the experimenter. However, if at any time you feel dizzy or nauseous, feel free to take off the HMD or stereo glasses immediately.

At the end of each trial, we will ask you some question regarding the difficulty level of the task. You will be asked to give a rating between 1 and 10 to assess the level of difficult of the task. Additionally, at the end of each set of 8 trials, we will ask you some questions regarding your level of comfort at that time. Again, you will be asked to give a rating between 1 and 10 to assess your level of arm strain, hand strain, dizziness, and/or nausea. Please try to be consistent with these ratings during the experiment.

At the end of the experiment, we would love to hear any comments you might have about the techniques we are testing, the experiment itself, or your performance of the tasks. Again, thank you for your participation.

The World-in-Miniature Technique:

The selection and manipulation technique that you will be using today is called the World-in-Miniature (WIM) technique. The WIM is a scaled down version of the virtual environment (5% scale in our case). The objects in the WIM are avatars of the objects in the virtual life-sized environment. Moving an object in the WIM will also move the actual object virtual life-sized environment.

To select an object in the WIM, use your virtual hand. Once the tip of the index finger (or the transparent cube on top of the index finger) coincides with the object, press the top left button on the wand to select the object. Your virtual hand will highlight to give you feedback that the object has been hit and it is now “selectable”. Once you select the object, it will be attached to the index finger and you can manipulate it. To place the selected object, hit the upper left button on your wand again. The object is now deselected and will be placed at that particular position.

B2: Pre-questionnaire for Experiment 2

9. Age:

10. Gender: Male Female

11. Left-handed or right-handed? Left handed Right handed

12. Do you have perfectly uncorrected vision or wear glasses/contacts?

    Uncorrected vision Wear glasses/contacts

13. How often do you play 3D games like Doom, Quake, etc.?

    Never Occasionally Often

14. Would you consider yourself a beginner or an experienced user with Virtual Reality devices like the CAVE and HMD?

    Beginner Experienced

B3: Post-questionnaire for Experiment 2

1. What were some of the difficulties you had while selecting the darkened cube in the HMD?

2. What were some of the difficulties you had while manipulating the darkened cube in the HMD?

3. What were some of the difficulties you had while selecting the darkened cube in the CAVE?

4. What were some of the difficulties you had while manipulating the darkened cube in the CAVE?

5. Can you describe some changes that can be made to the WIM technique in the HMD and/or CAVE to make it easier to select and manipulate objects in the environment?

B4: Results of Experiment 2

Size1 = 0.4 m cubes   Size2 = 1.0 m cubes
Target1 = 1.5x   Target2 = 2.5x

Average selection times for all users in the HMD and the CAVE (all times in seconds):

Average Selection Times for all Users

	CAVE	HMD
Size 1	5.513	5.944
	5.916	5.597
	6.632	7.029
	5.589	5.699
	5.314	5.964
	4.259	4.771
	6.010	5.631
	6.541	5.677
Size 2	3.015	4.322
	3.856	3.819
	5.109	2.980
	2.764	1.512
	4.652	2.821
	4.107	3.276
	3.666	4.145
	4.281	4.095

ANOVA for selection times:

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	35.4206649	1	35.4206649	57.7098113	2.8274E-08	4.19598223
Columns	0.48502016	1	0.48502016	0.79022859	0.38160473	4.19598223
Interaction	0.78731642	1	0.78731642	1.28275068	0.26699412	4.19598223
Within	17.1856153	28	0.61377198			
Total	53.8786168	31				

Average manipulation times for all users in the HMD and the CAVE (all times in seconds):

Average Manipulation Times for all Users

	CAVE	HMD
Target 1	3.601	2.945
	4.077	4.238
	5.255	2.915
	3.737	2.380
	4.357	3.550
	4.637	4.116
	4.862	2.761
	4.264	5.282
Target 2	3.221	3.651
	5.391	2.686
	4.292	5.371
	3.780	2.960
	4.247	3.856
	3.068	2.598
	3.569	3.630
	5.704	4.770

ANOVA for manipulation times:

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	0.001062	1	0.001062	0.001353	0.970924	4.195972
Columns	3.348495	1	3.348495	4.265104	0.048272	4.195972
Interaction	0.254534	1	0.254534	0.324209	0.573629	4.195972
Within	21.98255	28	0.785091			
Total	25.58664	31				

Selection errors for all users in the HMD and the CAVE:

Selection errors		
	CAVE	HMD
Size1	1	3
	4	4
	0	2
	2	1
	1	2
	2	2
	4	2
	5	2
Size2	2	1
	0	0
	2	1
	1	0
	2	0
	1	0
	0	0
	1	1

ANOVA for selection errors:

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	19.53125	1	19.53125	16.02564	0.000417	4.195982
Columns	1.53125	1	1.53125	1.25641	0.271854	4.195982
Interaction	0.78125	1	0.78125	0.641026	0.430081	4.195982
Within	34.125	28	1.21875			
Total	55.96875	31				

Correlation analysis:

	Subject	Age	Gender	Lefty	Vision	D Game	Experie	AvgSelT	AvgManip	MDSelEr	EAvgSel	AvgMani	VESelErrs
Subject	1.0000												
Age	0.2242	1.0000											
Gender	-0.6299	-0.0297	1.0000										
Lefty	0.2474	0.3689	-0.6547	1.0000									
Vision	-0.0563	-0.7560	0.1491	-0.4880	1.0000								
3D Games	-0.2315	-0.1090	0.4082	0.0000	0.3651	1.0000							
VE Experienc	0.1690	0.8621	-0.1491	0.4880	-0.4667	0.0000	1.0000						
HMDAvgSelT	-0.1633	-0.6044	0.1453	-0.4225	0.1658	0.0462	-0.8909	1.0000					
HMDAvgMani	0.3873	-0.1888	-0.5078	-0.1419	-0.1487	-0.7115	-0.4233	0.5207	1.0000				
HMDSelErrs	-0.5224	-0.6536	0.0727	-0.2381	0.2277	0.0000	-0.7482	0.7754	0.3759	1.0000			
CAVEAvgSel	0.1685	-0.1232	0.0298	-0.4256	-0.2516	-0.6742	-0.5138	0.5642	0.7680	0.2807	1.0000		
CAVEAvgMa	0.3259	-0.1519	-0.1828	-0.2924	-0.0827	-0.7259	-0.4223	0.3505	0.7571	0.2474	0.8680	1.0000	
CAVESelErrs	0.5855	-0.4365	-0.5164	-0.1690	0.5774	-0.1581	-0.3464	0.2087	0.5322	0.1690	0.1526	0.4551	1.0000

Appendix C

C1: Instructions for Experiment 3

Introduction & Instructions
Selection/Manipulation Experiment:

Welcome, and thank you for agreeing to participate in our study. We are studying the effect of a display type, the Cave Automatic Virtual Environment (CAVE) in this case, for selecting and manipulating objects in a three-dimensional virtual environment. Through this experiment, you will be helping us to better understand the effect of displays on a selection and manipulation technique. You will be asked to perform some tasks in a 3D virtual environment, using the CAVE and a button device called the wand for input. We may also ask you some questions about how you liked the technique. It is very important that you remember that <u>we are testing these techniques - we are not testing you.</u>

This experiment will consist of some practice time at the beginning to familiarize you with the selection and manipulation technique, and then 32 experimental trials using the technique. You should be able to finish the entire experiment in 30 to 45 minutes.

Your task will be to pick up a specified object and place it in a target location. You will see an array of nine cubes. The eight cubes on the outside of the array will be gray, while the center cube will be blue. The blue cube is the one you are to pick up. After you have grabbed this blue cube, you will try to place it between two target indicators. The target indicators are cylinders, you just need to place the object between them and do not have to match the orientation. In all cases, the virtual hand will change colors when you have placed the blue cube in the correct position. A trial will not end until you have selected the center cube and placed it in the target.

Select the blue cube   Place it between the target indicators

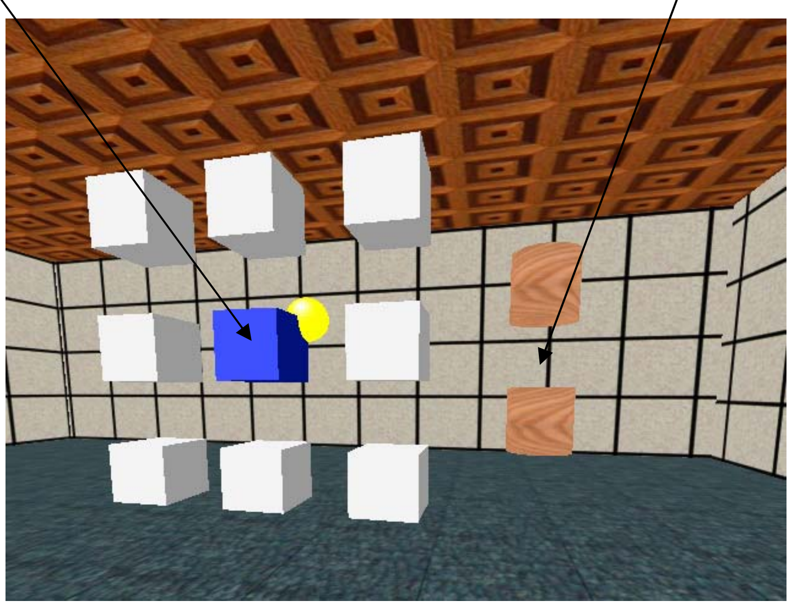

The overall score for each trial will depend on several factors:

- the time you take to grab the blue cube
- the number of incorrect cubes you grab before the blue cube
- the time you take to position and drop the blue cube

Please ask any questions you have before you begin the experiment. If you have questions during a trial, please wait until the trial is over before asking the experimenter. However, if at any time you feel dizzy or nauseous, feel free to take off the stereo glasses immediately.

At the end of each trial, we will ask you some questions regarding the difficulty level of the task. You will be asked to give a rating between 1 and 10 to assess the level of difficult of the task (1 – least difficult, 10 – most difficult). Additionally, at the end of each set of 8 trials, we will ask you some questions regarding your level of comfort at that time. Again, you will be asked to give a rating between 1 and 10 to assess your level of arm strain, hand strain, dizziness, and/or nausea (1 – little or no strain, 10 – lots of strain). Please try to be consistent with these ratings during the experiment.

At the end of the experiment, we would love to hear any comments you might have about the techniques we are testing, the experiment itself, or your performance of the tasks. Again, thank you for your participation.

The World-in-Miniature Technique:

The selection and manipulation technique that you will be using today is called the World-in-Miniature (WIM) technique. The WIM is a scaled down version of the virtual environment (5% scale in our case). The objects in the WIM are avatars of the objects in the virtual life-sized environment. Moving an object in the WIM will also move the actual object virtual life-sized environment.

To select an object in the WIM, use your virtual hand. Once the tip of the index finger (or the transparent cube on top of the index finger) coincides with the object, press the bottom left button on the wand to select the object. Your virtual hand will highlight to give you feedback that the object has been hit and it is now “selectable”. Once you select the object, it will be attached to the index finger and you can manipulate it. To place the selected object, hit the bottom left button on your wand again. The object is now deselected and will be placed at that particular position.

C2: Pre-questionnaire for Experiment 3

15. Age:

16. Gender: Male   Female

17. Left-handed or right-handed?   Left handed   Right handed

18. Do you have perfectly uncorrected vision or wear glasses/contacts?

    Uncorrected vision   Wear glasses/contacts

19. How often do you play 3D games like Doom, Quake, etc.?

    Never   Occasionally   Often

20. Would you consider yourself a beginner or an experienced user with Virtual Reality devices like the CAVE and HMD?

    Beginner   Experienced

C3: Post-questionnaire for Experiment 3

If you are in Group D, please answer ALL the questions.

Group A (Color):

6. What were some of the difficulties you had while selecting the darkened cube in the WIM with the dark blue floor?

7. What were some of the difficulties you had while selecting the darkened cube in the WIM with lighter colored floor?

8. What colored floor did you prefer for selection and manipulation the cube and why?

Group B (Scale):

9. What were some of the difficulties you had while selecting the darkened cube in the larger WIM?

10. What were some of the difficulties you had while selecting the darkened cube in the smaller WIM?

11. Which WIM did you prefer for selecting and manipulating the cube and why? If you feel that a different size of the WIM would have been even better, mention that too.

Group C (Training):

12. What were some of the difficulties you had while selecting the darkened cube without any training?

13. What were some of the difficulties you had while selecting the darkened cube after training?

14. Did you feel that the training helped you select and manipulate the cube easier and faster? Why or why not?

C4: Results of Experiment 3:

Size1 = 0.4 m cubes   Size2 = 1.0 m cubes
Target1 = 1.5x   Target2 = 2.5x

Group A (Brightness Improvement):

Average selection times and manipulation times and selection errors for all 3 users:

Selection time

	Original	Improved
Size1	5.728	4.080
	5.064	4.662
	5.355	4.860
Size2	4.158	3.501
	4.101	4.051
	3.075	4.083

Manipulation time

	Original	Improved
Target1	4.020	3.642
	5.288	5.588
	5.438	4.836
Target2	6.890	4.098
	4.967	4.574
	5.787	7.682

Selection Errors

	Original	Improved
Size1	5	2
	0	3
	2	5
Size2	1	0
	1	1
	1	1

ANOVA for selection times, manipulation times and selection errors:

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	3.831124	1	3.831124	20.32749	0.001979	5.317655
Columns	0.419862	1	0.419862	2.227737	0.173899	5.317655
Interaction	0.67468	1	0.67468	3.579772	0.095128	5.317655
Within	1.50776	8	0.18847			
Total	6.433426	11				

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	2.241108	1	2.241108	1.42537	0.266719	5.317655
Columns	0.323367	1	0.323367	0.205665	0.662235	5.317655
Interaction	0.031046	1	0.031046	0.019746	0.891723	5.317655
Within	12.57839	8	1.572299			
Total	15.17391	11				

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	12	1	12	5.333333	0.049736	5.317655
Columns	0.333333	1	0.333333	0.148148	0.710342	5.317655
Interaction	1.333333	1	1.333333	0.592593	0.46354	5.317655
Within	18	8	2.25			
Total	31.66667	11				

Group B (Scale Improvement):

Average selection times and manipulation times and selection errors for all 3 users:

Selection time

	Original	Improved
Size1	2.835	2.901
	5.169	4.808
	4.920	5.396
Size2	2.046	2.305
	4.996	3.973
	4.987	4.444

Manipulation time

	Original	Improved
Target1	3.014	2.972
	6.234	5.390
	7.743	4.589
Target2	3.790	3.012
	5.091	5.460
	4.549	6.256

Selection Errors

	Original	Improved
Size1	0	0
	1	0
	0	0
Size2	1	1
	0	1
	1	1

ANOVA for selection times, manipulation times and selection errors:

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	0.896123	1	0.896123	0.477975	0.508903	5.317655
Columns	0.105281	1	0.105281	0.056155	0.818635	5.317655
Interaction	0.184388	1	0.184388	0.098349	0.761841	5.317655
Within	14.99867	8	1.874834			
Total	16.18446	11				

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	0.264813	1	0.264813	0.099586	0.760402	5.317655
Columns	0.626147	1	0.626147	0.235471	0.640506	5.317655
Interaction	2.373964	1	2.373964	0.892762	0.372388	5.317655
Within	21.27298	8	2.659122			
Total	24.5379	11				

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	1.333333	1	1.333333	8	0.022204	5.317655
Columns	0	1	0	0	1	5.317655
Interaction	0.333333	1	0.333333	2	0.195016	5.317655
Within	1.333333	8	0.166667			
Total	3	11				

Group C (Training Improvement):

Average selection times and manipulation times and selection errors for all 3 users:

Selection time

	Original	Improved
Size1	3.046	2.605
	7.499	5.484
	4.495	3.902
Size2	2.608	2.224
	3.977	3.971
	5.234	4.311

Manipulation time

	Original	Improved
Target1	4.046	2.878
	4.950	6.045
	5.188	5.171
Target2	3.561	3.780
	3.903	9.402
	4.436	5.418

Selection Errors

	Original	Improved
Size1	0	1
	2	1
	4	1
Size2	3	0
	3	2
	0	1

ANOVA for selection times, manipulation times and selection errors:

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	1.845242	1	1.845242	0.722387	0.420069	5.317655
Columns	1.585678	1	1.585678	0.620771	0.453471	5.317655
Interaction	0.251322	1	0.251322	0.098389	0.761794	5.317655
Within	20.43494	8	2.554367			
Total	24.11718	11				

ANOVA

Source of Variation	SS	df	MS	F	P-value	F crit
Sample	0.411718	1	0.411718	0.142077	0.716026	5.317655
Columns	3.641008	1	3.641008	1.256454	0.294841	5.317655
Interaction	3.841443	1	3.841443	1.32562	0.282832	5.317655
Within	23.18276	8	2.897845			
Total	31.07693	11				

ANOVA

Source of Variation	SS	df	MS	F	P-value	F crit
Sample	0	1	0	0	1	5.317655
Columns	3	1	3	1.5	0.255508	5.317655
Interaction	0	1	0	0	1	5.317655
Within	16	8	2			
Total	19	11				

Group D (All improvements):

Average selection times and manipulation times and selection errors for all 3 users:

Selection times

	Original	O+B	O+S	O+T	O+B+S	O+B+T	O+S+T	All
Size1	5.069	7.994	2.822	6.734	7.294	3.694	7.026	3.033
	3.852	2.695	4.290	6.598	3.438	4.368	8.897	6.656
	6.239	8.771	7.268	7.536	6.851	4.982	4.092	7.413
Size2	2.325	4.221	6.901	3.829	2.763	3.440	4.701	1.131
	3.597	2.560	2.672	3.869	4.221	2.434	2.981	5.701
	5.764	3.128	1.744	4.144	2.545	4.736	1.597	2.177

Manipulation times

	Original	O+B	O+S	O+T	O+B+S	O+B+T	O+S+T	All
Target1	5.390	7.562	5.175	6.097	5.902	3.860	2.992	7.240
	5.315	3.078	9.031	5.686	6.325	6.352	5.728	8.419
	6.871	5.764	7.650	6.437	8.681	7.017	4.535	5.530

Selection errors

	Original	O+B	O+S	O+T	O+B+S	O+B+T	O+S+T	All
Size1	0.000	0.000	0.000	0.500	0.000	0.000	0.000	0.000
	0.000	0.000	0.333	0.000	0.500	0.000	0.667	0.000
	0.500	0.000	0.000	0.500	0.000	0.000	0.000	0.000
Size2	0.500	0.333	1.000	0.000	0.000	0.000	0.000	0.000
	0.000	0.333	0.000	0.500	0.000	0.000	2.000	0.000
	0.000	0.333	2.000	0.000	0.000	0.000	1.000	0.000

O+B = Original + Brightness Improvement
O+S = Original + Scale Improvement
O+T = Original + Training Improvement
O+B+S = Original + Brightness + Scale Improvement
O+B+T = Original + Brightness + Training Improvement
O+S+T = Original + Scale + Training Improvement

ANOVA for selection times, manipulation times and selection errors:

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	61.72589	1.00000	61.72589	17.76445	0.00019	4.14910
Columns	9.05175	7.00000	1.29311	0.37215	0.91175	2.31274
Interaction	12.55706	7.00000	1.79387	0.51627	0.81545	2.31274
Within	111.18994	32.00000	3.47469			
Total	194.52464	47.00000				

Anova: Single Factor

SUMMARY

Groups	*Count*	*Sum*	*Average*	*Variance*
Original	3	17.5755	5.8585	0.7699032
Orig+Color	3	16.40375	5.46791667	5.0926516
Orig+Scale	3	21.855	7.285	3.8168292
Orig+Training	3	18.21875	6.07291667	0.1413235
Orig+Color+Scale	3	20.907	6.969	2.2415507
Orig+Color+Training	3	17.2284167	5.74280556	2.7702984
Original+Scale+Training	3	13.255	4.41833333	1.8819306
All	3	21.1890833	7.06302778	2.1105193

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Between Groups	19.50070	7.00000	2.78581	1.18388	0.36518	2.65720
Within Groups	37.65001	16.00000	2.35313			
Total	57.15071	23.00000				

ANOVA

Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*
Sample	0.520833	1	0.520833	3.202847	0.082979	4.149097
Columns	2.321759	7	0.33168	2.039654	0.080315	2.312741
Interaction	1.821759	7	0.260251	1.600407	0.171111	2.312741
Within	5.203704	32	0.162616			
Total	9.868056	47				

Vita

Dhruv B. Manek was born on 10th March 1980 in Mumbai, India. He graduated from Virginia Tech in May 2002 with a degree in Computer Science. He is currently pursuing a Masters degree in Computer Science and Applications at Virginia Tech. Dhruv will begin a position as a Quality Assurance Tester at Global Computer Enterprises Inc. in July 2004.