Migrating Three Dimensional Interaction Techniques

Brian Elvis Badillo

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

Master of Science
In
Computer Science

Dr. Doug A. Bowman, Chair
Dr. D. Scott McCrickard, Committee Member
Dr. Manuel A. Pérez-Quiñones

23rd April 2007
Blacksburg, Virginia
Keywords: Virtual Environments, 3D Interaction, Migration, Selection, Manipulation

Migrating Three Dimensional Interaction Techniques

Brian Elvis Badillo

ABSTRACT

Multiplatform virtual environment (VE) development is fast-becoming a realization for today's developers. 3D user interfaces (3DUIs) can easily be ported to a variety of VE systems. However, few researchers have addressed the need to intelligently migrate 3DUIs across VE systems. We claim that the naïve migration of 3D interaction techniques (3DITs) to other VE systems could result in decreases in usability. We also claim that device specificity can be used to increase usability on these other VE systems. In this thesis, we have chosen three manipulation 3DITs to naively migrate across a set of four VE systems. We use an exploratory usability study to identify any usability issues stemming from our naïve migrations. After finding decreases in usability in select migrations, we redesigned two of the 3DITs for device specificity. We investigated the benefits of our redesigns with usability studies on the original, naïve, and redesigned implementations of both 3DITs. Results from our studies are mixed. In one case we demonstrate that device specificity can be used effectively to increase 3DIT migratability. As a result from our experience in this work, we have learned several lessons in device-specific design as well as 3DIT migration.

Acknowledgements

Thanks to my wife, Joy, for all her patience and understanding. I'm very grateful for her support throughout.

Thanks to all my family for their support in all of my endeavors. Their comfort was invaluable to me in times of stress.

Many thanks to Dr. Bowman for giving me his time and effort when I needed it most. His understanding and helpfulness is greatly appreciated.

Table of Contents

List of Figures

List of Tables

Chapter 1 – Introduction

In this chapter, we introduce virtual environment (VE) systems and 3D user interfaces (3DUIs). We present an argument for the intelligent migration of 3DUIs, present the problem, and outline our research approach to this topic.

1.1 – Motivation

Virtual environments have several applications for work in the real world including scientific visualization, designing and prototyping, simulation, collaboration, entertainment, and more. Today, there are many specific examples of the benefits of working with VEs in these applications. In scientific visualization, we see VEs being used in geotechnical engineering and geology with the stereoscopic GeoWall display [1]. Immersive VR [2] and collaborative VEs [3] are used for visualization in medical applications. Scientists and engineers can use 3D graphics paired with 3D user interfaces to interactively obtain insight from complex data [4]. In the design domain, there exist VE tools for 3D modeling [5] and CAD [6, 7]. We can see VEs at work designing and prototyping architecture [7] and automobiles [8] where the immersive nature of VEs allow designers to more effectively visualize their designs. The virtual aspect of VEs makes them good for training. Medical professionals enjoy the benefits of VEs with remote surgical training [8]. Firefighters, pilots, and military personnel, and operators of various kinds can also be trained in this way because VEs allow them to build experience while remaining in the safety of a VE [8, 9]. Immersive VEs are effective in the psychiatric treatment of phobias such as fear of heights [10], arachnophobia [11] and other mental ailments such as post-traumatic stress disorder [12]. Collaborative work also benefits from the use of immersive VR [13]. And inevitably VEs are used for entertainment [14]. With so many different applications covering so many disciplines we begin to see how VEs are relevant to various sectors in society. These applications are valuable to many organizations in research, academia, industry, medicine, entertainment, the military, and more.

Virtual environments are very different from traditional desktop systems and as such have different requirements in user interface (UI) design. Traditional desktop user interfaces use commonly accepted displays and input devices such as the monitor, keyboard, and mouse. They have user interfaces and interaction techniques that have undergone significant design and evaluation with regard to usability and study is continuing. However, these interfaces and interaction techniques are meant for the two dimensional realm while VEs use three dimensional spatial contexts. Given this difference, VEs require a new set of displays, input devices, and interaction techniques to successfully interface with them [15]. These displays, input devices, and interaction techniques are the components of 3D user interfaces (3DUIs).

3DUIs are the means by which a human interacts with a VE in a 3D spatial context [15]. In this capacity, 3DUIs are an integral part of VEs ranging from the 3D desktop application to the fully immersive application. A 3DUI for a 3D desktop CAD application might include 3D interaction techniques implemented for a standard mouse such as a WIMP-based (Windows, Icons, Menus, Pointing device) interface for system control tasks, widgets for object manipulation, and viewpoint control using separable axes in conjunction with mouse drags. However, consider a similar example where a fully immersive VE with an HMD and a six degree of freedom device is used for 3D modeling. A 3DUI for this application must support the same tasks of system control, manipulation, and navigation as was the case before. However, the devices used afford different interaction techniques and an altogether very different 3DUI.

3DM, a 3D modeling tool developed by Butterworth et al. built for an HMD and a 6DOF input device, shows what a 3DUI for these devices could look like [5]. A user stands on a “carpet” that defines the tracked area they are standing in. To navigate the environment, they simply walk around in this area, inspecting objects they are creating from different angles in a natural manner. If they wish to navigate to an area not covered by the tracked area users have the option of moving their “carpet” around the world by “grabbing” the air and dragging their viewpoint to different locations and angles. When the user is ready to work with models in the environment, they can select modes and commands from a toolbox virtual window in the environment simply by reaching out their hand and grabbing a tool. This same direct manipulation metaphor is used to select and manipulate objects and vertices. We can see in these two examples that using 6DOF devices for input can greatly enhance interaction in 3D contexts. But the methods for interaction must change to accommodate the differences in device. 3DUIs are closely connected to the VE devices on which they are built.

VE systems are comprised of a wide set of input and output devices [6, 15, 16]. Display devices such as head-mounted displays (HMDs) and immersive projection technologies (IPTs), along with spatial input devices such as wands, gloves, and trackers, etc. make up a formidable combination space for VE systems. Types of displays range from simple fish tank VR [17] (CRT monitors in conjunction with stereoscopic glasses to), the fully immersive HMD, to the truly three-dimensional volumetric displays [18]. Among the most common display equipment for virtual environments are the wearable HMDs and the immersive projection technologies (IPTs). These are highly used in all sectors of society that have a need for VEs. This is because both of these sets of equipment provide fully (or near fully) immersive VR and can produce stereoscopic images from a head-tracked view.

There are some variations to the HMD. Some HMDs occlude the real world while others can provide augmented reality by letting the user see the real world and a fabricated image. Among these latter HMDs are video see-through displays, projection see-through displays [19], and virtual retinal displays [20].

There are also some variations to the IPT. In some VE applications, a plain large display wall such as the GeoWall [1] is sufficient. However, more common are the CAVE [21], RAVE [22], or ReaCTor displays. These wall displays surround the users on multiple sides to provide better immersion. Still another semi-immersive wall display is the Ensphered Vision display developed at Tsukuba, Japan [23]. This is just a hemispherical wall that surrounds the user’s entire field of vision.

VE input devices can vary in the amount of DOFs they provide, the type of input they generate, what metaphors they afford, and their size or shape. Some examples of input devices are the simple palette of buttons, joysticks, ShapeTape, the SpaceMouse, the SpaceBall, Pinch Gloves, Data Gloves, 6DOF trackers, etc. Input devices available for use in VEs are large in number [16] and new devices are still emerging [15].

There is a huge explosion of devices available because they are sometimes created to support very specific applications [24]. The application domain may dictate or inspire the creation of a special display or input device to be used. For example, VE designers in the medical application domain have created displays, input devices, and 3DITs specific to tasks related to the domain. Many VE tasks in the medical domain involve the inspection of CT or MRI scans by a group of medical professionals. Usually the scans take the form of volumetric data and inspection tasks involve the rotation, scaling, and clipping of the data. Kitamura et al. has created a specific display to facilitate carrying out these tasks called IllusionHole [3]. This display supports multiple medical professionals sharing a stereoscopic view of the same object while also supporting quick and usable rotations of the display. The display is specifically tailored to support interaction for medical inspection tasks. CavePainting is an excellent example of using a custom input device [25]. CavePainting is an art application that supports the paint metaphor through the use of specialized devices such as a bucket, paintbrush, and palette. Similarly, Hinckley et al. has developed a two-handed device to easily support rotation and clipping tasks on volumetric data of the head using the passive haptics and specific device affordances of a six degree-of-freedom (6DOF) doll's head [26]. This input device has been created specifically to facilitate interaction with volumetric head data.

High-fidelity VE systems are still expensive to put together. This exacerbates the problem of standardizing VE devices because an organization may not have the funds available to obtain certain VE devices over others. They might use cheap devices that don't have the same functionality as their more expensive counterparts. For example, lower DOF inertial trackers might be chosen over more powerful 6DOF trackers. Some organizations also try to keep costs down by using fewer devices. For example, Fish-Tank VR is a commonly used VE system. This setup uses head tracking to make the environment semi-immersive but usually only provides a standard keyboard and mouse for interaction. Physical space can also become a cost when deciding which devices will make up a system. This is especially the case when considering displays with large footprints such as the CAVE or devices that may require special architectural planning such as a motion platform.

For the many reasons outlined above – a lack of standard displays and input devices, differing application domains, and cost constraints – different organizations tend to own widely varying VE systems. Even though VE systems are so different between organizations, these organizations may want or need to use applications that have been used or created in other VE systems. If a VE application is specific to a single organization's VE system then its value is lost by its limited accessibility and use. In contrast, if a powerful VE application can be used by several organizations it gains value and usefulness. For this reason, we want to be able to use VE applications independent of device. Consider the practicality of VE device independence in a scenario involving an architectural VE application in both the commercial and academic sectors. An architecture firm may own a VE system composed of an HMD and a pen-and-tablet. The architecture department at a university has access to a CAVE and a wand that is a common resource of the institution. Both of these organizations wish to work with VE applications that will prototype structures rapidly and in an immersive way. It would be practical if we could migrate this architectural VE application equally well on the wide variety of VE devices available to these two organizations.

Today there are several development tools available that facilitate multiplatform VE development [27-31]. A VE platform here is defined as the set of development tools, libraries, and the OS/hardware platform that a VE is implemented on. Multiplatform VE development tools make developing applications for multiple VE platforms easier with minimal effort by abstracting platform-specific development. For example, rendering can be accomplished via the tool in a platform-independent fashion while input can be handled by generic or abstract events such as a discrete button push or a 6DOF device poll. Some of those that are freely available include DIVERSE and VR Juggler [27, 29]. Others that are in advanced stages of development include CHASM and Envir3D [28, 30]. These tools make porting VE applications and 3DUIs easier for developers.

In spite of this, if we are concerned with maintaining usability, simply porting VE applications and 3DUIs across devices is not sufficient. VE devices have considerable effects on the usability of 3DUIs [32-37] and they should influence 3DUI design [38]. So we cannot simply port VE applications to different VE systems; we must reevaluate the design of their 3DUIs too. We must migrate 3DUIs intelligently so as to provide more effective interfaces for widely used VE applications. Today, intelligent 3DUI migration has been left as an open-ended research question [15].

1.2 – Migrating 3DUIs

As mentioned before, 3DUIs are made up of several building blocks including VE devices and 3D interaction techniques (3DITs) [15]. 3DITs are the components of 3DUIs that allow users to accomplish specific tasks in VEs. However, we believe 3DITs are not inherently migratable for reasons involving the characteristics and properties of displays and input devices combined with the design specificity of interaction techniques. In our research, we specifically address the migratability of 3DITs.

VE displays and input devices have their own intrinsic characteristics that make them advantageous in certain situations and not well-suited for others. For example, Pinch Gloves are an input device that supports several discrete input events in a natural way through the simple pinching of one's fingers. This affords a natural metaphor for object manipulation [39]. However, this device lacks a method for specifying continuous input, like a joystick, for applications involving specifying velocity.

In spite of the fact that immersive display devices have different advantages and disadvantages and input devices afford different interaction metaphors, there are well-established general purpose 3D interaction techniques (3DITs) that should be used on all of these systems. Some examples of these include the HOMER and Go-Go object manipulation techniques or the WIM technique for navigation [40-42]. These techniques are designed to be easy to use, easy to understand, and thus usable methods of 3D interaction for any VE application. Lots of hard work has been put into designing and quantifying the benefits of these techniques.

However, they have each been developed on a specific VE system and it is possible that their design was influenced by the system on which they were developed (we call this the primary VE system). For example, the WIM technique was developed primarily for a VE system composed of a six degrees-of-freedom (6DOF) HMD for display and two 6DOF trackers with a buttonball for input device. The technique involves displaying a miniature version of the virtual world in the tracked non-dominant hand of the user. In an HMD, this works quite well. However, in a CAVE the user's actual hands occlude their view of the virtual world making viewing something in the hand a potential usability issue. This subtle difference may negatively affect the usability of the technique when it is implemented in the CAVE (we would call this a secondary VE system).

Can these techniques be used without detrimental effects when implemented over a variety of secondary systems? Little work has been done to show either way. If there are heavy usability flaws when naively implementing 3DITs onto secondary VE systems, can the 3DITs be redesigned to leverage specific characteristics of the system? 3DITs may not be inherently migratable if they are implemented on secondary VE systems in naivety. However, we can make them migratable by intelligently redesigning the 3DITs on the VE systems that breakdown usability.

In this thesis, we aim to show that 3DIT migration is not a simple problem. We also would like to show that using device specificity, or redesigning 3DITs specifically for secondary VE systems, can be a solution to the problem. Therefore, in this thesis we seek to address the issue of 3DIT migration through empirical study of the effects of VE devices on 3DITs and the benefits of VE device specificity.

1.3 – Definitions

In this section we present terms we use and their definitions in the context of this work.

Virtual Environment (VE) – A synthetic, spatial (usually 3D) world seen from a first-person point of view. The view in a VE is under the real-time control of the user [15].

Virtual Reality (VR) – Synonymous with VE. We use the term VE in this document because the term VR is associated with unrealistic hype generated by the media [15].

3D User Interface (3DUI) – A user interface that involves human-computer interaction in which the user's tasks are performed directly in a 3D spatial context [15].

3D Interaction Technique (3DIT) – A method allowing a user to accomplish a task via the 3DUI. Implementation of a 3DIT is responsible for mapping the information from the input device (or devices) into some action within the system, and for mapping the output of the system to a form that can be displayed by the output device (or devices) [15].

Usability – The characteristics of an artifact such as a display, input device, 3DIT, or 3DUI that affect the user's use of the artifact [15]. There are several components of usability including learnability, efficiency of use, memorability, resistance to error, and subjective satisfaction in usage [43].

Specificity – Consideration in the design of 3DITs to increase usability by decreasing over-generalities. Specificity can be applied to the domains, tasks, displays and input devices, and users of VEs [38].

VE System – A combination of display and input devices that are used to interact with a 3D user interface. This is the hardware makeup of a virtual environment.

Primary VE System – In the context of this work, this is the VE system on which a 3DUI or 3DIT was originally designed.

Secondary VE System – In the context of this work, this is the VE system to which a 3DUI or 3DIT has been migrated.

3DIT Migration – The process of moving 3DITs in a device-independent way from one VE system to another. 3DIT migration entails both a software design and usability design component. The software design component is responsible for the details of simply porting a 3DIT across VE systems. The usability design component is responsible for maintaining usability of 3DITs on various VE systems.

Naïve Migration – The migration of a 3DIT that neglects the usability design component of 3DIT migration where design choices are made. For example, when migrating a 3DIT across widely varying devices a designer might face a decision in mapping input. If this mapping is done arbitrarily or without regard to potential effects on usability, the process is said to be a naïve migration.

Migratability – In the context of this work, this is a measure of how well usability is maintained when a 3DIT undergoes naïve migration.

1.4 – General Approach

We claim that VE device specificity is a solution to the problem of 3DIT migration. Our approach to this work attempts to break down the problem into smaller parts to obtain observable results to support our claim. We bring the problem down to a smaller and manageable scale by choosing a common VE task, set of VE devices, and set of 3DITs to empirically study. We must also choose a multiplatform development tool for developing the environments and 3DITs for this work.

1.4.1 – Research Questions

We have developed some research questions to help ourselves investigate the problem of migrating 3DITs. The following questions represent the focus of this research.

1. *How do different display devices affect the usability of well-established general purpose interaction techniques?*
   Displays are the front line of visual feedback to the user. They are the means by which a user gains understanding of how their input affects the environment. It is likely that certain properties of the display used will affect the way a user understands their interactions.

2. *How do different input devices affect the usability of well-established general purpose interaction techniques?*
   Input devices are the primary means for a user to act within the virtual environment. These devices enable a user to provide input in very different ways that span from simple discrete input such as buttons to 6DOF continuous input such as tracker data to any combination in between. An important characteristic of input devices are the affordances they carry that make them more understandable to use. This aspect of input devices could cripple the usability of 3DITs if the metaphor breaks down in the absence of matching affordances.

3. *If the VE device used affects the usability of the interaction technique, can the interaction techniques be made display-specific or device-specific to maintain usability?*
   If we find that the usability of certain interaction techniques severely breaks down after being used on a secondary VE system, can we fix it? It may or may not be possible to redesign certain aspects of a technique to make it as usable as the original implementation on its primary system.

1.4.2 – Intelligent 3DUI Migration Hypotheses

Through this research we hope to show that 3DIT migration is a problem that can be addressed and that VE device specificity is useful for solving this problem. From the research questions presented above we can make hypotheses. We note that this research does not employ traditional hypothesis-testing methods, because our work considers only a small set of displays, devices, and 3DITs and investigates these in a practical manner. However, these hypotheses serve the purpose of expressing our beliefs on 3DIT migration.

We present the following three hypotheses for 3DIT migration:

- *In addressing research questions 1 and 2 we will find that naively migrating 3DITs across VE systems will result in usability breakdowns.*
  When 3DITs are developed they are intentionally or unintentionally designed to leverage characteristics of the primary VE system they are created on or for. When the 3DIT is implemented on a secondary VE system, the assumptions made in the design process may be broken. As a result usability will degrade.

- *In response to research question 3 we will find that VE device specificity can be used to maintain usability when migrating 3DITs.*
  Redesigning general purpose 3DITs in a display/device-specific way will avoid losses in usability on secondary VE systems by including characteristics of the VE system in the design process.

- *From research questions 1, 2, and 3 we hypothesize that we can use design knowledge to migrate 3DITs in an intelligent way to avoid naïve migration issues.*
  Through the study of design properties of 3DITs and characteristics of VE systems we can identify usability issues that can negatively affect 3DIT migratability. VE designers should be able to intelligently choose which 3DITs to use with particular VE system devices when they are designing 3DUIs. We should also be able to create reusable design principles to help the design process for 3DIT migration.

1.4.3 – Specificity

In the past, 3DUI research has proceeded in the direction of creating and evaluating generic or universal 3DITs. 3DITs such as HOMER [40] and Go-Go [41] are novel interaction techniques for manipulation tasks. They were developed to extend and make more powerful the direct manipulation metaphor. After being developed, their design was validated by extensive user study and their acceptance into many 3DUIs constitutes their best-practice status. 3DITs such as these were created for the "typical user" to perform general manipulation tasks in any VE application [38]. Of note, they are also assumed to work on any 6DOF input device and any display [38, 40, 41].

Creating 3DITs in the way described above produces a highly general-purpose 3DIT that can be more widely used. However, we are increasingly finding that over-generalities are not desirable design characteristics [38]. For example, we have seen in the past that the strength of these generic 3DITs can vary among types of tasks when ray-casting was shown to be difficult in tasks involving smaller objects [44]. We also have seen that the strength of a 3DIT can be affected by varying VE devices. For instance, Manek found that using a CAVE display rather than an HMD for the WIM technique resulted in degraded performance [34].

A current trend in research strives to create and evaluate 3DIT flavors (or small variations in generic 3DITs) and address the benefits of 3DIT specificity [38, 45]. There are four areas in the design of a 3DIT where we can design for specificity. There is application/domain specificity, task specificity, user specificity, and device specificity [38]. It has been shown that by designing 3DITs from the ground up to address specific characteristics of the VE on which they are implemented, we can create more effective 3DITs. For example, Chen and Bowman show that a 3DIT specifically designed for cloning tasks in an architecture application called PORT-C outperformed other general purpose 3DITs that were modified to support cloning tasks as well [46].

This work addresses device specificity as part of the larger scheme to study the benefits of 3DIT specificity. We claim that in the same way that Chen and Bowman have used design specificity to enhance 3DIT usability on the task and domain level we can use design specificity to maintain 3DIT usability across VE devices.

1.4.4 – Empirical Study of the Effects of VE Systems on 3DITs

The purpose of this research is to further understand and overcome the issues involved in migrating 3D interaction techniques over various VE system setups. To reach this end, we must actually take individual interaction techniques and redesign them if possible to make use of the various VE systems. This brute force approach is, in this case, susceptible to a combinatorial crisis. Taking the number of possible popular VE output devices multiplied by the number of possible popular VE input devices and researching problems associated with possible interaction techniques on these devices would be exhaustive, but also exhausting [15]. So a subset of display devices, input devices, and interaction techniques should be defined. We constrain ourselves to working only with manipulation techniques because manipulation tasks are so salient in virtual environments and there are many well-established best-practice manipulation 3DITs to choose from.

The issues that arise when an interaction technique is migrated away from the original VE system it was designed on should be experienced empirically. So we need to evaluate each interaction technique's migratability through observation. An initial formative experiment will be designed that requires participants to perform tasks with each of these 3DITs using both the primary system (on which it was designed) and secondary systems. In this way we can survey changes in usability between the 3DITs on their primary and secondary systems and thus measure their migratability. A redesign of the interaction techniques that have lower migratability among the various VE devices will then undergo a final evaluation to compare their migratability with their unaltered counterparts. This is consistent with user-centered VE evaluation as proposed by Gabbard et al. where a formative and qualitative evaluation precedes an iterative redesign of the system and a comparative evaluation demonstrates the usability of the iterated system [47].

1.4.4 – CHASM

A development system for migrating 3DUIs must provide extensible device support and API abstraction. Additionally for this study we must have the ability to make speedy changes to 3DITs so a tool that supports rapid prototyping and reusability is also needed.

CHASM (Connected Hierarchical Architecture of State Machines) is a development tool that does just this [28]. Development is done on the level of concepts and well-defined states. The tool supports abstraction to such a degree that lower level concepts can be implemented before higher level concepts are even designed. In addition, concepts can be implemented in terms of the designer's desired resolution. In this capacity, a developer can support device migration by designing 3DIT behavior using higher level concepts and leaving interchangeable low level concepts to interface with various VE devices.

One of the main foci of CHASM is that of code reuse, which in turn causes rapid development to be a forte of the system. These aspects of CHASM make it attractive for use in a development system for migrating 3DUIs. During the course of the research we will have to re-implement some parts of the interaction techniques we are examining. We will use CHASM to implement our interaction techniques because it will allow us to change only what is needed and reuse the code that is already written.

1.5 – Potential Impact

This work will potentially enable VE designers to create better 3DUIs through the use of device specificity. If we can prove that 3DITs can be effectively migrated through device-specific redesigns then we can assert that general 3DIT design benefits from device specificity. This would validate another area where specificity is useful and spark research in this area.

From this work, some device-specific 3DIT redesigns will have been created and evaluated. These could become best-practice 3DITs on the VE devices for which they have been created. In addition, some insight into VE device specific redesign will be gained from the design process. These lessons-learned can be used by subsequent developers to guide design for migratability.

On the large scale, this work will make VE applications more powerful and useful to the many organizations that use them. Organizations that use the same VE applications on different VE systems will benefit from having usable and migratable 3DUIs.

1.6 – Overview

In chapter two, we define the current state of 3DIT migration research. We describe work that identifies the effects of VE devices on usability. We also show that the current state of 3DIT migration research is very young and that our work fills this space by actually addressing migration specifically. In this chapter, we also describe several systems that facilitate and guide the design of migratable interfaces.

As mentioned before, there is a huge combination of VE system devices, 3D interaction techniques, and 3D interaction tasks available to a VE designer. In chapter three we describe various characteristics and properties of these. Tackling a problem this large is too big for a master's thesis. So a subset of the 3DUI-VE system space is researched and evaluated in this project. Specifically, the focus of this project will be on migrating manipulation task 3DITs. We use the HOMER, Go-Go, and Voodoo Dolls techniques. The HMD and the CAVE make up the display variable. And the input device variable consists of the wand and Pinch Gloves.

Chapter four presents an informal study we run to identify problem areas in the interactions between the techniques and the devices. The usability of each interaction technique is examined on each display and device combination.

In chapter five, the problem areas that are identified with the informal experiment in chapter four are brought into consideration and redesign for usability is undertaken. The properties of the techniques that contribute to higher or lower degrees of migratability are identified and discussed. Changes are made to the interaction techniques to make them more usable on the devices with which migration issues were found.

A formal evaluation is conducted in chapter six to quantify the benefits of the changes made to the interaction techniques. The goal of the evaluation is two-fold. The evaluation demonstrates the usability gains between the naïve implementations of the interaction techniques and the redesigned implementations. The evaluation also serves the purpose of trying to show that usability can be preserved across VE devices. This is accomplished by pitting the redesigned implementation of the interaction technique on the secondary system that it did not perform well on against the original implementation of the technique on its primary system.

Finally, in chapter seven we discuss the implications of this work and describe potential future research in this area.

Chapter 2 – Related Work

The migration of 3DUIs is a new frontier in the VE research community. These days we can develop one application for multiple VE systems with several VE development tools available that support multiplatform portability. There has been some work to identify the effects of VE displays and devices on usability. However, this work does not explicitly tackle the problem of making interfaces migratable across these systems. Work in creating multiplatform 2D interfaces is further along than its 3D interfaces counterpart. We may be able to gather some knowledge and insight into how they set up guidelines for developing migratable interfaces.

2.1 – The Effects of VE Systems on Usability

Inherent characteristics of VE systems affect the usability of the 3DUIs implemented on them. It is important to identify how and to what capacity these systems affect usability and to show empirically that they in fact do. However, few studies have done this and those that do only show isolated differences between VE systems and their impact on tasks in VEs.

Some studies have addressed the effects of various VE displays on VE task performance and 3DIT usability. Arthur performed an extensive study on the effects of an HMDs field-of-view on performance during walking tasks [48]. Although the work is very controlled, it is also constrained in that it does not explore the effects of other display characteristics such as field-of-regard or stereopsis on other important VE tasks such as selection and manipulation. His dissertation also strictly focuses on measuring performance rather than usability as a whole.

Manek has completed work to show more qualitative effects on a variety of display characteristics such as field-of-view, field-of-regard, and real-world occlusion on selection/manipulation tasks [34]. This work concludes that several attributes can affect the usability of 3DITs on more than the performance level. However, this work only varies display and does not look at the interplay between display and input device.

Steed and Parker conducted an interesting study regarding the interplay between 3D selection and manipulation techniques and VE displays [35]. He performed an experiment that identified performance differences between using virtual hand or virtual pointer metaphors on HMDs or immersive projection technologies (IPTs). Although his experiment very clearly illustrates impacts of display on selection times, his contributions are constrained to only a few guidelines pertaining to choosing VE displays based on task types.

Closer to our own work, Kjeldskov studies qualitative impacts of displays on various VE interactions [33]. His work is broader than other work in that he studies the effects of display on both travel and selection/manipulation tasks. Kjeldskov divides displays into two categories, fully and semi immersive displays, and has identified several issues that affect the usability of various travel and selection/manipulation techniques. However, the categories of displays he has defined may be oversimplified as there are several other characteristics of displays of which their effects on usability haven't been studied. Moreover, the study does not focus on the interplay between displays and various input devices.

Few studies have focused on performance or usability issues associated with VE input devices. On the performance level, Zhai and Milgram have demonstrated empirically the importance of interplay between device characteristics and interaction techniques [37]. With regards to device characteristics, Hinckley et al. has performed a study comparing various methods of performing accurate object rotations [32]. Although the goal of this study was to gauge the usability of some 3D versus 2D techniques, the study did find a large difference in user preference for device form factor in this task. However, even with these studies we know little about the interplay between device characteristics and 3DITs.

All of the studies described above are great for showing the effects of displays or input devices on VE performance or 3DIT usability. These effects could be responsible for varying degrees of 3DIT migratability. Some prior work considers or identifies the aspects of VE systems that pertain to 3DIT migratability such as device specific design. For instance, Lindt considers the notion of limited device exchangeability for 3DITs. She argues that 3DITs are inherently designed for specific VE systems [49, 50]. Manek's work supports this claim with a study on the effects of the CAVE on the World-In-Miniature (WIM) 3DIT. His results suggest that WIM is an HMD specific technique by showing that its usability is negatively impacted by the CAVE. We denote the set of devices for which a 3DIT is designed the primary system and all other VE systems are denoted as secondary systems.

But, there is still a great need for work on migrating 3DITs to the VE systems that show drops in VE performance or usability. Some preliminary work has been done by Manek on migrating the WIM technique from the HMD to a CAVE [34]. He found that users perceived a larger WIM than that in the HMD and speculated that users did not have sufficient depth cues in the CAVE to efficiently accomplish tasks. As a result, he attempted to implement a CAVE-specific version of the WIM interaction technique by making the WIM somewhat smaller and by training users how to use motion parallax to their advantage for gaining better depth perception. The work, however, did not produce significant results. Also of note is LaViola's work involving the Step WIM, a CAVE-specific navigation technique [51]. However, the only part of this technique that separates it as CAVE-specific is the auto rotation based on a scaled 2D Gaussian function. This rotation function confused users at first but was necessary to provide the power of users to fully rotate 360 degrees in a four-walled CAVE. In addition to requiring experience, this feature may also lead to spatial disorientation.

Our work extends previous work pertaining to migration by studying the effects of both display and device on a variety of 3DITs. Previous work does not account for interplay between display and input device. In our work we recognize the problem of 3DIT migration and define the parameters of migratability. Previous work described above only seeks to quantify the effects of certain displays and devices and does not necessarily apply results to the problem of 3DIT migration. As seen in Manek's work, the migration of 3DITs also involves creating display- and/or device-specific implementations of the techniques [34]. We make the claim that device specificity is a good solution to the migration problem and we create display/device-specific 3DIT flavors to back it up.

2.2 – Migrating 2D Interfaces

Work in migrating 2D graphical user interfaces (2D GUIs) centers around two fronts: creating toolkits and libraries to facilitate GUI migration or developing guidelines or models for keeping migratable GUIs usable. The latter work is pertinent to the topic of 3DUI migration. Researchers in the 2D realm have found that simply porting 2D GUIs onto different platforms is not sufficient. They have found that design guidelines and transformations for platform specific implementations are needed for successful 2D GUI migration.

Claims are packets of design knowledge that express advantages and disadvantages of a particular design [52]. McCrickard, et al have proposed that claims can be used to express design strategies for platform migration [53]. The interesting advantage of claims is that they can evolve as designs are evaluated and redesigned [52]. It has been shown that the use of claims can in fact encode usability differences in VE displays as McCrickard, et al have used claims to investigate differences in usability between different display devices using identical tasks[53]. Falsifiable claims were used as a method of evaluation as well as a method for deploying rich guidelines for UI developers.

Florins and Vanderdonckt have also proposed extensive guidelines for developers of multiplatform 2D GUIs [54]. They propose a system of graceful degradation or putting in place a set of design guidelines that help developers decide how to transform interfaces in such a way as to preserve continuity among multiple platforms. In their work they have classified certain model transformations as less impacting on continuity than others. The authors suggest employing these transformations first before others to achieve an effect of gracefully degrading continuity among platforms. Also of note are that the authors believe that transformations applied at more concrete levels generate less discontinuity than transformations applied at higher levels.

With the advent of XML (Extensible Markup Language) several XML-based systems have been developed to facilitate multiplatform interaction modeling. UIML, XIML, and UsiXML are all XML-based mechanisms to support model-based design [55-57]. All of these systems have the ability to model interaction in a platform independent fashion. Additionally, these systems use abstract models to separate GUI rendering among widely varying platforms. Through the use of transformations an abstract model can be made concrete and render-able to a specific platform. In this way, intelligent migration of 2D GUIs can be made to widely varying platforms.

These systems for guidelines and model-based multiplatform development tools provide an excellent means for developing 2D GUIs that are usable on many widely varying platforms. In the 3DUI realm, platforms also vary widely in many respects. We can see from work in the migration of 2D GUIs that platform-specific implementations of the GUI have become a good solution. We should use the same solution of platform-specificity in VEs.

2.3 – Migrating 3D Interfaces

Migrating 3D interfaces does not just involve the creation of developing tools to get 3DUIs to run on different platforms. As in the multiplatform 2D GUI development, the migration of 3DUIs should also happen on the design level. There are some tools available that allow for VE system independent 3DUI design specification.

Meant specifically for 3DUI design and based on XML is InTml (Interaction Technique Modeling Language). InTml provides VE developers with a modeling system specifically to model interaction in conjunction with 3D geometries [31]. Although InTml is a system that provides the ability for fast and reusable prototyping on a variety of VE systems, it does not allow for transformations in interaction to take place over VE systems.

Influenced by InTml is the Interface Component Description Language (ICML) [50]. This specification encapsulates a conceptual framework for 3D interaction components (3DICs) which encompass 3D widgets, 3D gadgets, and 3D interaction techniques. The language enables the description of specific and generic 3DICs. The conceptual framework takes into account user, task, and device specificity. As applied to device specificity, the generic 3DIC ICML description is device independent and therefore independent of concrete realization. For realization of the 3DIC a specific 3DIC ICML description should be given that provides specific device mappings.

Comparable to this, a model-based system for VE interface migration that merits special mention is Envir3D [30]. Envir3D works within the model-based paradigm using Abstract Interaction Objects (AIOs) and Concrete Interaction Objects (CIOs). AIOs are a mechanism to abstract interface elements from any concrete representations so that they can be later transformed into platform-specific CIOs. An important contribution of Envir3D is that interaction objects can be transformed over various VE systems resulting in interaction object migration. Additionally, AIOs and CIOs can be evaluated against usability guidelines imposed by Envir3D developers.

Automatic evaluation metrics, however, provide a very limited measure of the system. Usability evaluation of VEs cannot be carried out in the same way 2D interfaces are through these automatic evaluation metrics. 3DUIs will more effectively be evaluated in terms of usability using evaluation methods proposed by Poupyrev, et al.[44] and Bowman, et al. [58]. Another drawback of Envir3D is its focus of use for 3D widget design. For better 3DUI migration, we must not only model interaction objects, but also interaction behaviors.

Systems and Description Languages such as Envir3D [30] and ICML [50] enable VE developers to design 3DITs for migratability. Each of these systems allows developers to design several device-specific realizations of the same general 3DIT. However, work must still be done to establish device specificity as a solution to the migration problem. Additionally, these systems do not provide any design knowledge about the effects of displays and devices on 3DITs and how to redesign these 3DITs for new displays and devices.

Chapter 3 – 3DITs and VE systems

As mentioned previously, we have hypothesized that the usability of 3DITs can degrade as a result of migrating them from their primary VE systems to secondary VE systems. This would be caused by discontinuities in the interplay between 3DITs and the VE displays and devices that they may be implemented on. We have asserted before that displays and input devices have intrinsic characteristics that make them advantageous in certain areas and detrimental in others with respect to 3DITs. In the following sections, we will identify these characteristics of displays and input devices. In these sections, we also identify the displays, devices, and 3DITs that are to be further investigated in this research. At the end of the chapter we describe our naive 3DIT migrations.

3.1 – VE Displays

There is plenty of work that categorizes and creates taxonomies for VE displays. A large percentage of literature divides displays into two types. Some call these types wearable and immovable. Others distinguish displays as head-mounted displays (HMDs) and immersive projection technologies (IPTs) [35]. Still others divide displays into fully-immersive and semi-immersive. [33] In the simplest example, the HMD and the four-sided CAVE, the HMD fits the first term for all of these classifications and the CAVE fits the second term. However, consider that a six-walled CAVE would be classified as an immovable display and yet considered arguably fully immersive. Also consider volumetric or fish tank VR displays which may not fit in some of these classifications. Why should we create a dichotomy of displays if we cannot agree on division criteria? In order to gain an understanding of all displays we must use display characteristics to classify them. In this section, we wish to discuss why most researchers choose to create a dichotomy of displays and identify factors of displays from this dichotomy that may affect interaction.

3.1.1 – Discussion of Display Characteristics

Immersion has been defined as an objective and measurable characteristic of displays that produces fidelity stimuli to the senses [59]. Display characteristics that affect immersion include field of regard, field of view, depth cues, and real world occlusion [15, 60, 61]. Another group of display characteristics is form-factor which includes obtrusiveness and screen geometry [15]. And a final group of display characteristics is display quality which includes spatial resolution, display brightness, color contrast, and frame rate [15]. A description of these VE display characteristics follows.

Field of Regard

Field of regard (FOR) is, "the amount of physical space surrounding the user in which visual images are displayed." [15]. For example, an HMD has a 360 degree FOR because the user will see the virtual world wherever he/she looks. In contrast, a display wall only has a FOR as big as its field of view because the user only sees the virtual world on the confines of the immovable display wall.

While wide FORs are an important aspect for immersion and presence in general, many VE systems in practice use IPTs that do not provide a 360 degree FOR. If the application does not require users to make view rotations, then sacrificing FOR does not create an issue. For example, an application in a four-sided CAVE for visualizing the dashboard of an automobile does not necessarily have to support view rotation to show the back of the car. So, the user is always facing the front wall of the CAVE and never needs to manually rotate his/her view to bring the rear of the car into the FOR. However, in the other case that an application requires the user to see 360 degrees of their world, the VE system must support the task of letting the user perform view rotations manually or a means for the user to access 360 degrees of their virtual world. If these mechanisms are not already a part of a 3DIT, then they must be accommodated in some way. Potentially, accommodating this task in a 3DIT could become problematic and may be a hindrance to migratability.

Field of View

Field of view (FOV) is defined as, "the maximum number of degrees of visual angle that can be seen instantaneously on a display." [15]. For example, a six-sided CAVE theoretically has a FOV as large as the human eye's entire view field. But in practice stereoscopic glasses limit effective FOV significantly. The effects of FOV on performance is made clear by Arthur in his study involving FOV on HMDs [48]. Smaller FOVs result in degraded performance. Furthermore, limited FOVs give less information to the periphery of the user making search tasks in VEs more difficult than in displays with large FOVs. For instance, it was concluded in Manek's work involving usability differences between the CAVE and the HMD that the CAVE was preferred for selecting and manipulating objects in the users' FOV because its FOV is significantly larger than that of the HMD's [34].

Following from this, we may expect that 3DITs that take advantage of a large FOV may result in degraded performance and usability when used on a display with a smaller FOV. Take for example the Step-WIM, a navigation 3DIT designed specifically on the wide FOV of a CAVE [51]. This technique takes advantage of the wide FOV to facilitate search tasks of the virtual world. If the same 3DIT were to be implemented in an HMD, we may expect these search tasks to become more difficult because users cannot utilize their periphery.

Depth Cues

There are many ways in which our bodies survey depth information from our environment. These can be classified as pictorial depth cues, motion parallax, oculomotor cues, and stereopsis [15]. Consistent with the definition of immersion presented above, a user will get a better understanding of the environment as more of these cues are provided with good fidelity.

The stereopsis depth cue provides an example of how migratabilty can be affected by this display characteristic. Some 3DUI components such as floating menus are positioned in the immediate virtual surround of the user. In a monoscopic display, these menus are easily readable. However, the menu may suffer from binocular rivalry (or the phenomenon that occurs which inhibits the fusion of images to the viewer) [15] when used in a stereoscopic display because of the menu's closeness to the user's view. This may render the menu hard to read and hard to use. Thus we see the effects of depth cues on 3DUI usability across displays.

Real World Occlusion

Real world occlusion is the ability of a display to block out the user's actual environment from the virtual environment. There are varying levels of real world occlusion. For example, an HMD occludes the entire real world from the user's view. However, a six-sided CAVE blocks out everything except the user's own body. Additionally, see-though HMDs used for augmented reality block out only the parts of the world occupied by virtual objects.

This characteristic of some displays introduces an incorrect occlusion problem that may become detrimental to usability. In one case, such as in CAVEs, all but the user's body is occluded from the user's view. This, however, allows the user to occlude the virtual world with their body. Although it may be distracting to be able to see disparity between the actual hand and the virtual hand, it is interesting to note that large drops in usability have not been identified because of this display characteristic [34]. Somewhat contrarily, drops in performance have been shown when a virtual hand avatar is absent in interaction in a display such as a CAVE [62]. This evidence seems to support the claim that users use visual feedback far more heavily than proprioceptive senses.

Obtrusiveness

Obtrusiveness has to do with the ergonomics of the display. There are several things that contribute or detract from comfort including bulkiness, weight, and physical accommodation of the wearable device. Another factor that makes a display device obtrusive is the presence of a tether. Still another is whether or not the hands are needed to move or hold the display. All of these characteristics may impede freedom of movement or cause heavy fatigue to a user and should be considered in the design of 3DUIs [15].

Screen Geometry

Screen geometry refers to the configuration and shape of screens in the display [15]. This characteristic affects many things including how much physical space is allowed for the user, occlusion of the real world, field of regard. Additionally, the geometry of screens sometimes imposes visual artifacts or distortion that may hinder user performance due to loss in visual quality [15].

Display Quality

Manek has supported the claim that display characteristics such as brightness and color contrast affect the user in his study involving the CAVE [34]. Others have shown that visual quality and frame rates also affect user performance in VEs [6].

3.1.2 – A Display Dichotomy

As mentioned before, it is the practice of many researchers to create a simple dichotomy of displays and study differences in the division [34-36, 63]. This dichotomy is important, however, to enable researchers to make their claims and findings inductive to more displays than those specifically studied. Typically, researchers are studying the effects of an HMD versus a

CAVE. In this case, the dichotomy is actually fairly effective in studying all the effects of display characteristics as the HMD and CAVE fall in opposing categories in almost all of these. Table 1 shows this contrast. Consequently, researchers can show the effects and differences between many of these display characteristics simply by studying these two displays.

Figure 1: A CAVE

Figure 2: An HMD

The projection screens of the CAVE give a wide FOV (90 to 180 degrees) depending on where the user is looking. In contrast, HMDs typically provide a FOV between 30 and 60 degrees. However, HMDs also provide a 360 degree FOR while the CAVE's FOR is reduced by the lack of a back and top screen.

While an HMD completely occludes the real world from the user's view, the CAVE only partially occludes the world because of the missing back wall. Additionally, the CAVE does not occlude the user's body; rather, the user's body occludes the CAVE screens. This results in incorrect occlusions. A user's real hand would occlude the view of virtual objects both in front of and behind it producing an unnatural effect.

The HMD is a bulky device that protrudes out from the head. The display is front-heavy and is not light. In addition to tracker cables, a video cable is necessary to provide picture. In contrast a user in the CAVE needs only to where head-tracked active stereo glasses. These are lighter than and don't protrude as far as the HMD. The stereo glasses only require a tracker cable.

There are different limitations of the devices in terms of space. The HMD has the potential to be used in large tracked areas so long as there is tracking equipment and cable extensions available. The CAVE screens, however, limit the tracked area to within a 10x10 feet square area.

Both of these devices exhibit some reduced display quality in some area of their display medium. For instance, the CAVE floor can become dirty through use. And display optics in an HMD can distort the edges of the picture.

Table 1: Differences in characterisitics of HMD and CAVE displays

	HMD	**4-sided CAVE**
Field of Regard	360° horizontal and vertical	~270° horizontal ~180° vertical
Field of View	Typically Narrow	Wide
Depth Cues	Supports monoscopic and Stereoscopic depth cues	Supports monoscopic and stereoscopic depth cues
Real World Occlusion	Occludes real world Occludes the user's body	Partially occludes the real world Does not occlude the user's body
Obtrusiveness	Weighty and Bulky	Not very obtrusive
Screen Geometry	Plenty of physical space	Limited physical space
Display Quality	Distortion at edges of image	Variable quality of floor image

3.2 – VE Input Devices

Bowman et al. asserts that empirical studies involving input devices and their appropriateness for certain situations is difficult because of the amount of variables involved [15]. However, we must try to research each device's effects on 3DUIs and their interplay between other VE devices in a systematic, exhaustive way. However, an exhaustive study of VE devices is beyond the scope of this study and should be ongoing work as new devices are developed. We must constrain ourselves to studying only the most common VE devices used in 3DUIs involving manipulation tasks. In this section we discuss characteristics of VE devices, common VE devices for direct manipulation, and the importance of affordances in the interplay between VE devices and 3DITs.

3.2.1 –Input Device Characteristics

Characteristics of input devices are degrees of freedom (DOF), input type, device affordances, and form-factor [15]. These are presented below.

Degrees Of Freedom

DOF is independently controllable movement. A device can have more than one DOF, and tracking devices typically are 6DOF. DOF can be integral or separable and there are tradeoffs associated with this. For example, an integral 6DOF device is useful for fast and learnable location or spatially based tasks. However, if the task involves high accuracy and precision, trajectory control, or coordination, a separable 6DOF device may be more sufficient [15].

Input Type

This characteristic of devices contains many different properties including whether data reported by the device is continuous or discrete, whether the device is active or passive, or how the device senses.

A continuous device reports continuous data such as data reported by a joystick or tracker while discrete data is data such as reported by a button.

Passive devices are devices that require user action to report data. Active devices report data regardless of whether any changes have been made. Related to this property is whether or not a device can be parked. A mouse can be parked because when a user lets it go, it does not move. However, a tracker that is dropped reports values that reflect its fall.

A device can sense through absolute sensing or relative sensing. Additionally a device can be isometric or isotonic [15].

Device Affordances

Affordances play a key role in determining usability when combined with 3DITs because affordances are an important part of interaction metaphors. If the device affords a specific type of interaction that is supported by the interaction metaphor put in place by the 3DITs then a successful match has been made for optimal usability. Sometimes it is better to use specific devices over general devices in order to create a usable VE application doing specific tasks [15].

An example of a device affordance is illustrated by the bend-sensing data glove. This is an input device that measures the flexure of fingers. Consequently, this device affords gestural input. An interaction metaphor supported by this gestural input is a grabbing metaphor for selection.

Form-Factor

This characteristic of devices describes the weight, bulkiness, ergonomics, and comfort of the device. Tethers, cables, weight, and bulkiness can heavily restrict freedom of movement when using the device.

On another facet of form-factor, the types of grasps that the device supports are also important. For instance, there has been some work on the difference between power and precision grasps [64]. A power grasp uses arm muscles to perform coarse actions with a firm grip while a precision grasp uses hand and finger muscles to accomplish finer more delicate movements. Additionally, the device can be one- or two-handed, or even no-handed [15].

3.2.2 – Common Input Devices for VE manipulation

In practice, 3DITs that encompass manipulation tasks tend to also use a common combination of devices. We recognize that specialized input devices tailored to specific 3DITs and task types can be more useful and powerful than the generic input devices that we study here

[15]. However, we wish to constrain our current study to devices that are commonly used in manipulation tasks because an exhaustive study of the wide array of devices available and those still unimplemented would be far beyond the limitations of this work. The input device combination commonly used for manipulation 3DITs is a 6DOF tracker attached to some device capable of discrete input [15].

6DOF tracking provides continuous integrated control of six degrees of freedom. It is simply a natural means of controlling spatial movements. Integrated control of all input dimensions in VE manipulation tasks makes performance easier, more effective, and is similar to real world manipulation [15, 32, 65]. It can be said that the device affords real world spatial coordination. 6DOF tracking by itself does not produce a usable device, however, for two reasons. All of the DOFs supported by the device are used for dimensions in spatial manipulation tasks. So there is no extra channel available on the device to drive other any tasks other than the spatial movement of objects, such as picking up and dropping objects. Additionally, other tasks involved in manipulation tend to only require discrete input. Although a continuous device can be adapted to provide discrete input, a device built for discrete input may be better suited to the task. To produce a usable input device, then, a 6DOF device is usually attached to some other input device that can provide other kinds of input including at least discrete input.

In typical generic devices used for manipulation a 6DOF tracker is attached or built-in to some device used for discrete input. However, this device may vary in any number of other device characteristics. For example, a wand varies from a set of pinch gloves heavily in its form-factor and interaction affordances. But both of these devices do not lack in providing methods of discrete input.

Figure 3: A wand device

Figure 4: A pair of Pinch Gloves

We have chosen to use tracked wand and tracked Pinch Gloves devices for our investigation of 3DIT migration. The criteria for this choice are the availability of the devices in our lab, their support of the 6DOF and discrete input type required for the direct manipulation metaphor, and important differences in form-factor and device affordances that may affect the migratability of 3DITs.

	Tracked Wand	**Tracked Pinch Gloves**
Degrees of Freedom	Integrated 6DOF	Integrated 6DOF
Input Type	Buttons provide discrete input Joystick provides continuous input	Pinches only provide discrete input
Device Affordances	Affords natural hand movement	Affords natural hand movement Affords a pinching metaphor
Form-Factor	Somewhat obtrusive Requires a power grasp	Not bulky, obtrusive, or weighty Hands-free

Figure 5: Differences in characterisitics of the wand and Pinch Gloves devices.

3.2.3 – Discussion of Affordances and Interplay of Devices

Device affordances are an extremely important characteristic to consider that heavily affects the types of interaction metaphors that can be used. An example of this is the pen-and-tablet device. The pen-and-tablet affords a great metaphor for selecting menu items or other widgets in a window. So if the application involves lots of system control that can be easily performed with virtual windows, then the pen-and-tablet is a great choice as a device. However, if the domain of the VE is vehicle simulation, then the pen-and-tablet metaphor breaks down and another input device such as a steering wheel or flight stick should be considered. Another device that is a good example of the importance of device affordances that pertains to manipulation is Pinch Gloves. These gloves support a pinching metaphor that mimics the real-world tasks of picking up and dropping objects [15, 66].

There is also the issue of devices that show certain affinities towards other VE devices. Good interplay between VE devices is important for usability. Some devices display inherent properties that, when used with certain other devices, increase or decrease in usability. An example of this idea is seen in the use of the HMD. The HMD is a great display device for immersive virtual environments because it has a full 360 degree field-of-regard and it occludes the real world in its entirety. However, because the real world is blocked out of view some input devices used in conjunction with this display may become harder to use. Take for example a handheld device with many buttons on it. A user would need to feel around to make sure that the button they are pressing is the button they intend to press or risk pressing the wrong button. In contrast, if the user could see their actual hand holding the device, even just in their periphery, they may be able to make button presses with more confidence and certainty, leading to easier and more efficient task completion. So, a CAVE display may be better suited to use in combination with a wand for tasks that involve many different button presses. Using other constraints, if for any reason the VE display used must be an HMD, then some other discrete input device may be in order. Pinch gloves are a device that supports many discrete inputs but uses more powerful kinesthetic senses to enable users to blindly operate the device [39, 66]. So, pinch gloves may be better suited to use in combination with an HMD for tasks that involve several different discrete input events.

3.3 – Interaction Techniques for Selection and Manipulation

This section is a review of work regarding various selection and manipulation techniques. We have chosen selection and manipulation techniques to research because 3DITs frequently are based on selection and manipulation 3DITs [15]. For example, a grab-the-air navigation technique [67] is based on the direct selection and manipulation metaphor to manipulate a user's viewpoint. Manipulation here is more formally defined as spatial rigid object manipulation [15]. Selection and manipulation 3DITs are not separable in this research because these tasks are so interrelated that they often rely on the same working metaphors. From here on, selection and manipulation is simply referred to as manipulation.

Through this review we can identify more commonly used manipulation techniques and their underlying metaphors. These more common and well-established techniques should be studied for migratability because they are best-practice techniques and will probably be, if they have not already been, migrated to various VE systems. We also present our implementations of the naïve migrations of these 3DITs to the displays and devices we have chosen to study. Lastly through this review, we will identify strengths and weaknesses of these 3DITs and identify issues that may arise when implemented on these displays and devices.

3.3.1 – A Discussion of Manipulation Metaphors

Bowman et al. has divided manipulation into two different paradigms: arm-extension, and ray-casting [40]. These paradigms are very similar in nature to the taxonomic categories developed in by Poupyrev et al. as virtual hand and virtual pointer, with the addition of exocentric manipulation [68]. Although the taxonomies have different names, the techniques divide up among them in the same way. We will discuss the virtual hand, virtual pointer, and exocentric spatial rigid object manipulation metaphors in further detail.

The virtual hand metaphor is an egocentric metaphor that is simply a broad categorization of 3DITs that use the direct manipulation metaphor. A property of the metaphor that can be changed to produce variations of the basic technique is the mapping of the real hand's position and orientation to the virtual hand's position and orientation [68]. Other varying properties are the size and shape of the selection boundaries, and the representation of the virtual hand. A common best-practice 3DIT based on the virtual hand metaphor is the Go-Go technique [41].

The virtual pointer metaphor is also an egocentric metaphor. It is a natural metaphor that deals with the selection of objects at a distance from the user. Although the metaphor can be used to perform object manipulation, it is not ideal for performing object-centered rotations [68]. It is from this flaw that virtual pointing is usually only used as a selection component for more complex manipulation 3DITs [40]. Design properties that can produce variation within this metaphor are the mapping of the ray's direction, the size and shape of the selection volume, the ray's visual representation, and methods for disambiguating selection [68]. Common best-practice selection 3DITs based on the virtual pointer metaphor are ray-casting [69] and image-plane selection [70].

Exocentric metaphors actually use egocentric metaphors to perform selection and manipulation. However, rather than using these component metaphors to interact with the environment from within the environment, users interact with the environment from another view [68]. The power of the exocentric metaphor is in the manipulation of objects at a distance [42, 71-73]. Common best-practice exocentric 3DITs include WIM [42], Voodoo Dolls [71], and Scaled World Grab [73].

We have chosen a single 3DIT from each of the manipulation metaphors on which to focus our study. The Go-Go 3DIT will represent the virtual hand metaphor. The HOMER technique will represent the virtual pointer metaphor. And Voodoo Dolls will represent exocentric metaphors. On a side note, the HOMER manipulation technique is a hybrid technique of ray-casting and the virtual hand metaphor. However, we believe that the performance of virtual pointing techniques for manipulation is impractical and HOMER is a best-practice, widely used technique worthy of further study in migration research that partially represents virtual pointing for selection tasks.

3.3.2 – Manipulation 3DITs and Naïve Migration

We have chosen to study the effects of 3DIT migration to four different VE systems. These are distinct pairs of the CAVE or HMD displays with the wand or Pinch Gloves devices. The resulting display/device pairs are shown in table 2.

Table 2: Display and device combinations investigated in this experiment

	HMD	**CAVE**
Wand	HMD/Wand	CAVE/Wand
Pinch Gloves	HMD/Pinch Gloves	CAVE/Pinch Gloves

When implementing our 3DITs on each of these systems, we attempted to keep the implementations as close as possible to the original implementation. This is the definition of naïve migration. However, there are situations where some design decisions must be made to accommodate differences in VE devices.

For example, a major difference between the HMD and CAVE devices is their respective FORs. User of 3DITs in the HMD can access their entire virtual space by physically turning their body. However, the four-sided CAVE does not allow a 360 degree FOR and we must give the user a mechanism for accessing their entire virtual space as is possible in the HMD. In accommodating differences in the HMD and CAVE, we have added support for a manual view rotation task to our 3DITs in the CAVE display.

With 3DITs involving the Pinch Gloves, the manual view rotation was mapped to other pinches aside from the selection pinch. On single-handed 3DITs such as the HOMER and Go-Go techniques the view rotation pinches were mapped to thumb and middle finger pinches and thumb and ring finger pinches. Left view rotations were mapped to middle finger pinches and right view rotations to ring finger pinches. This mapping was made because it was intuitive assuming the user's palm is facing down, which was encouraged because the hand trackers were located on the backs of the hands. Since both gloves were used in the Voodoo Dolls 3DIT, left and right view rotations were mapped to thumb and ring finger pinches on the left and right hands respectively. The discrete input type of pinches limited the rate of rotation to being constant.

With 3DITs involving the wand, the manual view rotation tasks was mapped to the joystick located at the center of the wand. This channel of input was readily available on the device and provided a sufficient mapping to the task. For instance, if the user pressed the joystick left, the view rotates left, and if the user pressed the joystick right the view rotates right. The rate of rotation was made constant to keep the view rotations consistent with the Pinch Gloves implementations.

Table 3: 3DIT support of manual view rotation implementations on different VE systems

	HMD/Wand	**HMD/Pinch Gloves**	**CAVE/Wand**	**CAVE/Pinch Gloves**
HOMER	Full FOR Available	Full FOR Available	Left = Joystick left Right = Joystick right	Left = Middle finger pinch Right = Ring finger pinch
Go-Go	Full FOR Available	Full FOR Available	Left = Joystick left Right = Joystick right	Left = Middle finger pinch Right = Ring finger pinch
Voodoo Dolls	Full FOR Available	Full FOR Available	Left = Joystick left Right = Joystick right	Left = Left hand ring finger pinch Right = Right hand ring finger pinch

Another implementation detail was the input for selection tasks. On the wand, a single button was chosen to specify selections made. The user was expected to press to select objects and press again to deselect objects. Using Pinch Gloves, selection was performed by pinching together the thumb and forefinger. Here, as long as the pinch was held, the object was selected.

3.3.3 – Potential Migration Issues with Common Manipulation 3DITs

3DITs are usually designed for a specific VE system or set of VE systems [49, 50]. Although the 3DITs we have chosen to investigate for this study are general purpose manipulation techniques, they were designed on (and therefore designed for) a specific VE system.

All of the 3DITs we investigate in this study were originally designed in a VE system with an HMD. In migrating these techniques to the CAVE, we face all of the potential effects of displays on 3DITs described in section 3.1.

The HOMER and Go-Go technique were both originally developed using the wand. The wand easily supports these techniques by providing discrete input for selection and 6DOF input for manipulation. These 3DITs are strong because of their simplicity. Any device that supports at least discrete input and 6DOFs can potentially support the HOMER and Go-Go techniques.

However, we cannot assume usability remains intact. One immense factor that determines the strength of a 3DIT is the degree to which the interaction metaphor matches the affordances provided by the VE system being used. Mismatches in this area can render 3DITs unusable [15]. The HOMER and Go-Go techniques are extensions of the direct manipulation metaphor. And so, as asserted previously in section 3.2.2, it is important to use VE devices that afford the direct manipulation metaphor. 6DOF positional tracking devices do just this. 6DOF isometric devices such as the SpaceBall do not afford this direct manipulation metaphor as strongly and as a result suffer in performance [37].

The Voodoo Dolls technique was originally developed using the Pinch Gloves. The major characteristic of the device leveraged for the technique was the pinching metaphor. This metaphor is used to allow the user to hold virtual dolls and perform bimanual interactions with them in the same way you would in the real world. Attempting to migrate this technique to other devices may cause the metaphor to breakdown and with it the technique's usability.

3.4 – Summary

In this chapter we have presented several characteristics and properties of VE displays and devices. We have given several examples of how these characteristics can impact the migratability of 3DITs implemented on them.

We have also chosen displays, devices, and 3DITs to further investigate in this research. We have given some rationale behind our decisions by providing a dichotomy of displays and stressing the importance of device affordances and form-factor. We describe the primary VE systems of the 3DITs we have chosen and described our implementations of these techniques on secondary systems.

Next we perform a usability evaluation of our 3DITs on the displays and devices we have chosen. This will let us identify drops in usability on certain display and device combinations.

Chapter 4 – First Experiment (Usability Evaluation)

4.1 – Experiment Goals

The purpose of this experiment was two-fold. The first idea that we wanted to show was that there were usability problems associated with naïve migration of interaction techniques from their primary VE systems to secondary systems. Moreover we wanted to uncover what these problems are so that they could be further addressed. The goals of this experiment are exploratory in nature and the design of the experiment should reflect this. We are consistent with user-centered VE evaluation as described by Gabbard et al. in the design of this experiment and we consider the experiment as a whole to be a formative phase of the design and evaluation of display- and device-specific 3DITs [47].

Through this experiment we expected to find that there are large-scale problems in usability isolated on certain VE systems while using certain interaction techniques. We understand that interaction techniques are designed on specific VE systems or with specific assumptions in mind [49, 50]. We hypothesized that these same techniques, when placed unmodified or naively implemented on other VE systems, will break some of these assumptions and result in less usable interfaces. Along with this general hypothesis we also hoped to find that the issues identified can be remedied in some way to restore their usability.

4.2 – Overview of the Experiment

This experiment evaluates the usability of three interaction techniques on four VE systems. HOMER, Go-Go, and Voodoo Dolls are the three interaction techniques tested. The four VE systems that were considered are distinct combinations of the HMD and CAVE displays with the wand and Pinch Glove input devices. The experiment was designed to identify usability issues associated with secondary VE systems of the 3DIT being evaluated. Participants were asked to perform the same tasks on all combinations of the VE systems. Tasks consisted of selecting and manipulating common furniture of varying shapes and sizes to support the generalization of findings for specific manipulation task scenarios [47].

There are several types of data that can be collected in a usability experiment. We can classify data as qualitative or quantitative in nature. Qualitative data provides rich or descriptive information, but is not easily counted or measured. Quantitative data is information that is measured. Independent of qualitative or quantitative data is subjective or objective data. Subjective data is preference or opinion that can be expressed while objective data is observable. There are two common approaches to performing a usability experiment in VEs; through observation of users performing tasks or through measuring user task performance [47]. The former is used primarily for formative experiments while the latter is useful for comparative experiments. Because we wished to design a formative experiment to identify usability issues and problem areas, we wish to heavily use qualitative data for our experiment. Through observation we can collect objective qualitative data and through collecting the thoughts and opinions of users we can collect subjective qualitative data. We also used some simple performance metrics (quantitative, objective data) and Likert scale surveys (quantitative, subjective data) to reinforce any outstanding findings in the qualitative data [47].

4.3 – Hardware

Participants used a variety of VE equipment including a Virtual Research V8 head-mounted display (HMD) with 640x480 resolution and a 60° horizontal field-of-view and a four walled CAVE (CAVE Automatic Virtual Environment). For the HMD graphics were rendered on a Power Mac G5 running Mac OSX. The CAVE runs on a five PC cluster running Fedora Core 5. The CAVE consisted of four 10' x 10' screens, each having a 1280 x 1024 resolution. An Intersense IS-900 VET tracking system provided tracking for the head, wand, and Pinch Gloves in both systems. The input devices used were a pair of Fakespace Pinch Gloves and a standard Intersense IS-900 wand with four buttons and an analog joystick. Intersense hand trackers could be attached to the backs of the Pinch Gloves. In addition, a tracked wireless mouse was used in the non-dominant hand for two handed techniques that used the wand because of limited access to multiple wands.

Figure 6: A wireless mouse used in conjunction with a 6DOF hand tracker.

4.4 – Software

The environment and techniques were implemented using CHASM [28] on top of Simple Virtual Environments Libraries (SVE) [74]. SVE provided a means for interfacing with peripherals and rendering and loading the environment. CHASM was used to specify the behavior of the interaction techniques as well as the flow of the experiment. By using CHASM we can modify details of technique behavior easily for future design iterations. The SVE library was slightly modified to distribute events among the DADS cluster used in the CAVE. Additionally DIVERSE support was built into the library to make tracker input easier using DTK [75].

4.5 – Participants

The experiment involved a total of six participants. Their mean age was 20. There was one female participant and five male participants. All had near-perfect vision corrected or uncorrected. All but one of the participants was right-handed. The users were all novice users of VE systems. Novice users are defined as users that have never been exposed to the particular technique they were evaluating and the particular devices they were evaluating on.

Tasks were selection and manipulation interactions completed using one of three established 3DITs: HOMER, Go-Go, and VoodooDolls. These tasks were performed using one of two input devices and one of two output devices: Pinch Gloves or wands for the input device, and a CAVE or an HMD for the output device. By multiplying the input device space by the display device space we get that the possible space for input/output device combinations is four.

Each of the interaction techniques was tried over each device pair giving the experiment twelve conditions to study.

Each of the six novice participants were assigned an interaction technique to evaluate among the four VE system device combinations. Two participants were assigned to use HOMER, two were assigned to use Go-Go, and two were assigned to use Voodoo Dolls. Each participant underwent the entire battery of tasks using one of the three different interaction techniques performed over the four display/device combinations (HMD/Wand, HMD/Pinch Gloves, CAVE/Wand, CAVE/Pinch Gloves).

This experiment was necessarily conducted within subjects because this brings uniformity to so many unknowns in the study. In this way, the participant can recount the effects that the display and device may or may not have had on the usability of the technique and directly make comparisons on their experiences using each pairing. Since a majority of the data collected is qualitative data from participants, it is highly desirable to keep a participant's view across the study's space.

However, data from the study should show learning patterns in the latter tasks because the participants become more comfortable with the feel of working in a VE. It has been shown that novice users tend to build interaction preferences, or habits, in coordination with exposure to techniques. For example, users exposed to an environment that forced them to favor manual rotations continued to use manual rotations in an environment that did not [76]. For this reason the order of VE system variable exposure to users were varied. This decision is not only to counterbalance the study, but also to gain small differences in a user's view of a technique/device pairing. For instance, a user exposed to a technique on an HMD might develop interaction preferences for physically rotating their body to change their viewpoint. They might then be more apt to more fully use the side walls of the CAVE because they would first rotate their body to rotate their viewpoint and only manually rotate their view when necessitated by field of regard limits. A user first exposed to the CAVE might instead develop an interaction preference for only manually rotating their viewpoint while using only the front wall of the CAVE.

There were some limitations to the way the exposure to device pairings were varied. There were only two different orderings because there were only two users per interaction technique. Additionally, the orderings will block all of the HMD task sets together and the entire CAVE task sets together. This is because the HMD and CAVE are not collocated in the same lab. A logistical crisis would have occurred if the HMD and CAVE tasks were interspersed.

The VE system variable orderings were as follows (first to last):

- The first user of each technique would be exposed to:
  HMD/Wand, HMD/Pinch Gloves, CAVE/Wand, CAVE/Pinch Gloves.

- The second user of each technique would be exposed to:
  CAVE/Pinch Gloves, CAVE/Wand, HMD/Pinch Gloves, HMD/Wand.

4.6 – Environment and Tasks

The environment created for the experiment was a recreation of a small square room. The room was small enough to see all of the objects contained inside of it well and with ease. This eliminates the possibility that poor visibility of objects may confound the task of positioning objects at the extremities of the room. However, the room was not small enough to make all of the objects contained within naturally reachable. If this were the case, then the virtual magic introduced by the interaction techniques being studied would be unnecessary and the study would be ineffective.

Everyday objects such as a couch, lamps, and tables populated the room. Since the study is designed to be a formative tool we want to introduce variables that may help highlight unforeseeable usability issues. When compared to simple polyhedra, these everyday objects do just this. These objects aren't just convex and don't exhibit symmetry around all of their axes. Coincidently, using these everyday objects emulates an interior design application. In effect, these objects bring in a kind of randomness to the study as well as practicality. The objects' varying shapes and sizes creates a natural variable for evaluating general selection and manipulation task scenarios. The objects in the room also exhibited fairly involved semantic hierarchical structures to facilitate the use of the VoodooDolls technique.

Tasks were designed to force the participant to undergo specific conditions that were of interest to us. Some tasks were designed to be easy manipulations that did not involve large turns around the body. Some tasks were designed to get the user to manipulate objects at various angles around the body. A group of tasks was devoted to near-space selection and manipulation. Some tasks were devoted to far-space selection and manipulation. These tasks covered a variety of issues involved in selection and manipulation and spanned in difficulty from being easy to being hard with respect to target sizes, shapes, positions, and densities. A table of the specific tasks can be found in appendix A.

Participants that were involved in the HOMER and Go-Go conditions were asked to perform a total of twenty-two tasks. Of these, twelve were dedicated to measure performance with six for speed and six for accuracy. Six more tasks were simple controlled manipulations and the four remaining were complex manipulation tasks involving multiple manipulations.

Participants involved in the Voodoo Dolls condition were asked to perform a reduced set of the above tasks. In pilot runs of the experiment, the time commitment for the study was excessive because the 3DIT is complex and pilot users took longer to learn and perform the technique in general. We decided to eliminate some tasks to reduce the time commitment to one comparable to the other two conditions. In this condition, participants were asked to perform a total of ten tasks. Four of these were performance oriented tasks with two for speed and two for accuracy. Five tasks involved simple controlled manipulations and the remaining tasks was dedicated to a complex manipulation.

4.7 – Procedure

The experiment was performed as follows. Before each task set the participant was read instructions on donning and caring for the VE equipment they were currently using. The participant was then given instructions on the use of the interaction technique currently being used and was given time to familiarize themselves with it in the room environment.

The participant was asked to perform a series of practice tasks that illustrated the use of the technique in a variety of situations. Then the participant was given time to practice the technique in the environment at their leisure. After the participant said they were comfortable with the technique and the experimenter was satisfied with their performance, the participant moved on to perform the main series of tasks.

During each task the participant was asked to use a think-aloud protocol [77] (discussed below). After each task the participant was polled for their frustration levels and fatigue levels on a five-point scale. The experimenter gave the participant a task by defining the object to manipulate and a position in the environment to manipulate to. The task goal was made explicit through an orange highlighted copy of the task object in the final position.

At the end of the task set for each condition the participant was given a questionnaire to survey various subjective aspects. They were also asked to record some final comments.

4.7.1 – Collecting Qualitative Data

Qualitative data is useful for finding the source and reason for why usability problems may have occurred [47]. Because the goals of this experiment are to identify usability issues in the naïve migration of 3DITs, we have designed a data rich experiment focused on collecting qualitative data.

Objectively collecting qualitative data was accomplished through experimenter observation of user behavior. The observations were usually recorded at the occurrence of critical incidences [47, 77]. Gabbard et al. defines a critical incidence as, "an event that has a significant effect, either positive or negative, on user task performance or user satisfaction with the interface" [47] During the course of a task, if the user expressed frustration or a displayed a behavior that negatively or positively impacted the performace of the task their behavior was recorded. In this way, we can use this data to analyze where and why a usability problem occurred.

Observation, however, is subject to misinterpretation when analyzing data. For this reason we give the participant many opportunities to express their frustration, confusion, fatigue, and general thoughts or feelings on the system they are using. Qualitative subjective comments made by the user were recorded to give an idea as to what was causing problems in usability. This was achieved through the use of a think-aloud protocol. This is where each user was expected to explain their choice in actions and their thoughts on the interface while they were performing tasks [77]. Because the aim of the study was to identify issues, the experimenter was responsible to further question participants of their actions if they merited more explanation. Furthermore, the think-aloud protocol serves the purpose of identifying to the experimenter when a critical incidence might be occurring.

In addition to the think-aloud protocol, we give the participant ample opportunity to express their subjective views on a condition as a whole by collecting comments after an entire task set has been completed. This both serves the purpose of allowing the user to make general and reflective comments on the condition as a whole and allowing the user to make direct comparisons of previously completed conditions.

4.7.2 – Collecting Quantitative Metrics

Quantitative measures were taken to help indicate that usability problem may be occurring [47]. We can backup our claims that result from the qualitative data with qualitative data that highlights the problem areas.

Actual accuracy and speed metrics were measured by the system in meters and seconds respectively for each task. These objective measures will have little statistical power because only data from at most two participants can be compiled together per technique. Furthermore, in the interest of keeping down total evaluation run times, the amount of tasks designed to measure performance rather than behavior was small. However, these metrics could become a means of strengthening claims made based on feedback and comments provided by participants because strong patterns could still be seen even in small population sets. This data, which includes task run times and accuracy values, was recorded by the system.

It is also important to mention that a small selection of tasks was devoted to measuring user performance rather than behavior. In these tasks, we asked the user to optimize their performance by not providing comments during task completion, but to rather complete the task as fast or as accurate as possible respective of the task type. Some of these tasks were designed to measure speed by asking the user to perform the tasks as fast as possible. And some of these tasks were designed to measure accuracy by giving the user a certain amount of time to perform the tasks as accurate as possible. In these tasks a time constraint of ten seconds was deemed appropriate through pilot testing. Additionally, no visual feedback was provided for satisfactory completion.

Quantitative subjective measures were taken to help identify where general issues between interaction technique and device were perceived by the user. If the user perceives difficulty or ease it is also likely that they are experiencing difficulty or ease. Perceived metrics were measured by Likert scale survey items. Some Likert scale survey questions were also designed to measure perceived problems with technique-device interactions. Subjective satisfaction in the use of a technique is a component of a usable interface and can be detected in this way [43]. We can use this subjective data to back up usability claims that we make when compiling data from observation. The metrics used here were ease of use, frustration, precision, physical comfort, the participants perceived ability to interact with the environment, and the participants perceived ability to view their environment.

Another measure we felt would be interesting to incorporate into this study was a ranking of the VE systems variable. Users were asked to order the VE system device combinations in terms of preference. That is, after the entire study was conducted the participant rated each session in relation to the others. So we get see which devices pairs were the favorites and which were the least liked. This may help confirm that the interaction technique is meant to be used on a specific device combination, strengthen previous input provided by the user, or provide a surprising result that will spark further study.

4.8 – Results

From the evaluation, we see that the naïve migration of these 3DITs to secondary VE systems is mostly successful. This finding is contrary to our hypothesis and we will discuss possible reasons for this later. Metrics from this evaluation do not strongly suggest major or surprising findings such as a completely unusable 3DIT. Barring this, we do identify some outstanding usability problems that, if addressed with 3DIT device-specific redesign, would result in higher degrees of migratability for the 3DITs affected. In addition to these, we do extract some nuances within the space of the experiment that are relevant to issues that may arise when migrating these 3DITs.

The usability issues that are isolated within certain combinations of the technique/display/device independent variables of the experiment are the ones that are most interesting with regards to migratability because they identify issues involved with the interplay of VE devices and 3DITs. Some findings may be relevant to the topic of migration but they are not unique to a specific point within the independent variables of the experiment. Rather, some of these show differences between the HMD and CAVE independent of the interaction device used. And some of these show differences between the wand and Pinch Gloves regardless of the type of display used. In either case, the findings also depend on what 3DIT is being used. However, the HOMER and Go-Go techniques performed similarly and will therefore be referred to as the single-handed interaction techniques.

The evaluation identified a problem area when using single handed interaction techniques with the Pinch Gloves in the CAVE. It is an exemplary instance of loss in usability in the interplay between 3DITs, displays, and devices.

The evaluation also identified key usability issues in the interplay between the Voodoo Dolls and a wand independent of display. The issue explores the need for redesign to maintain usability and will therefore be given further attention. A presentation of these usability issues follows.

4.8.1 – Singe-handed 3DITs/CAVE/Pinch Gloves

We found that the usability of the single-handed techniques, including HOMER and Go-Go, was negatively impacted by a naïve implementation to the CAVE/Pinch Gloves VE system. We present evidence to this claim and interpret results to divine a cause.

Firstly, the subjective ranking of the conditions by the participants unanimously votes the CAVE/Pinch Gloves condition last (seen in table 4), even among the evident indecision seen in the first, second, and third places. The consistent last place ranking indicates that some issue in the CAVE/Pinch Gloves condition of the single-handed interaction techniques is causing negative reactions from the participants.

Table 4: Subjective rankings for the HOMER and Go-Go 3DITs

		Ranking			
		First	**Second**	**Third**	**Fourth**
HOMER	Participant 1	HMD/Wand	CAVE/Wand	HMD/Gloves	CAVE/Gloves
	Participant 2	HMD/Gloves	HMD/Wand	CAVE/Wand	CAVE/Gloves
Go-Go	Participant 3	CAVE/Wand	HMD/Gloves	HMD/Wand	CAVE/Gloves
	Participant 4	HMD/Wand	HMD/Gloves	CAVE/Wand	CAVE/Gloves

We also see that the participants on average scored the CAVE/Pinch Gloves condition lower than other conditions in general (illustrated in figure 7). This does not give us a clue as to where the issue is occurring. However, this does confirm that participants perceived that single-handed 3DITs naively implemented on a VE system composed of a CAVE and Pinch Gloves is the least preferred when compared to its implementation of other VE systems.

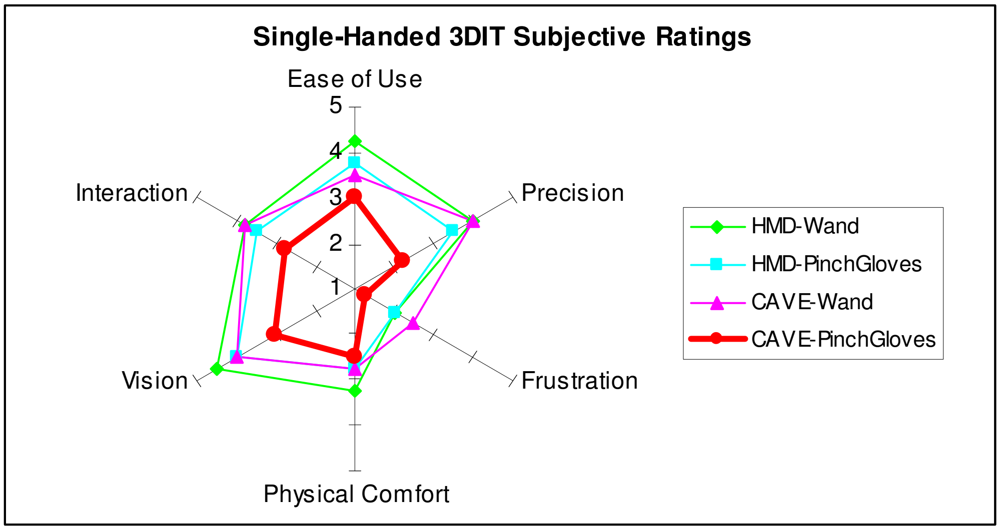

Figure 7: Average subjective ratings of both the HOMER and Go-Go 3DITs. The "Interaction" axis represents the user's subjective feeling on their ability to interact with the environment. The "Vision" axis represents the user's subjective feeling on their ability to see their environment. Refer to Appendix A2 to see the form used to collect these metrics. A smaller frustration rating indicated more frustration in this graph to consistently show that data points closer to the center were badly rated.

We now attempt to interpret our qualitative results to find out what characteristics of the single-handed 3DIT/CAVE/Pinch Gloves combination is causing the usability breakdown. Each user had their own complaint on the interplay between the CAVE and Pinch Gloves with one user commenting, "It was frustrating to have to drop the item to rotate the room." We blame the issues described here on unintuitive and uncomfortable device mappings. The naïve implementation of the single-handed 3DITs entailed mapping manual view rotations to arbitrary pinches using only one Pinch Glove. A pinch to the forefinger with the thumb was already mapped to selection. So, using the next most doable pinches, we mapped view rotations to the middle and ring fingers. A pinch using the middle finger and thumb resulted in a left-hand turn and a pinch using the ring finger and thumb resulted in a right-hand turn. This could be considered an intuitive mapping assuming that the right-hand is being used and the palm is facing down.

Some participants did not, however, find the mapping imposed on them to be appropriate. They were sometimes confused as to which pinch was mapped to which direction. One participant was so frustrated with this dilemma that they even gave an alternative solution to the problem, "The rotation was not intuitive and interfered with grabbing and was imprecise. Maybe add a joystick for the offhand." It is clear that the issue here is an unintuitive device mapping. For example, if the user held their hand in a posture with the palm face up, the initial rationale behind the "intuitive" mapping would lead the user to make inverse rotations.

The users unanimously complained about their inability to perform the main tasks concurrently with manual view rotation tasks. One user exclaimed, "Rotating room using Pinch Gloves and holding onto some object was really hard!" The naïve implementation of the single-handed techniques mapped the tasks of selection and manual view rotations to pinches performed on the same hand. Consequently, when users tried to concurrently perform these tasks, they also tried to concurrently perform these pinches. Although these kinds of pinches are doable, they are often uncomfortable and harder to perform, which makes them less usable and prone to error. These uncomfortable device mappings also drastically affected task times (as in the case of manipulations behind the user) because users were not inclined to perform the main task and manual view rotation tasks concurrently.

We find the major usability problem with this condition is caused by the fact that the back wall of the CAVE is not present. This forces a 3DIT implementation here to support some other means aside from natural rotation to rotate the user's view so that they can see the rear part of the environment. If the input device contains a method of input that supports this rotation that is intuitive and separate from the methods of input needed to interface with other parts of the 3DIT, then the usability remains intact. But if the interface to the 3DIT is diminished by the method of input used to rotate the view, then usability drops significantly. The problem with manual view rotations also goes hand-in-hand with the issues associated with the CAVE seams. Users have been noticed to tend to want to work at the centers of CAVE screens. This compounds the problem of less-than-usable view rotations as the user is then required to make additional manual view rotations for this view-screen alignment task, which is not a trivial task when using the Pinch Gloves.

Figure 8 shows results from the speed performance tasks. These tasks were divided into three types of manipulations including short distance and long distance manipulations as well as manipulations to the back of the user. We see clear trends in manipulation times in all of these cases between the CAVE-Pinch Gloves condition and the rest of the conditions. Interestingly enough, the clearest difference is seen in tasks that involve manipulations behind the user, which coincidently involves manual view rotation.

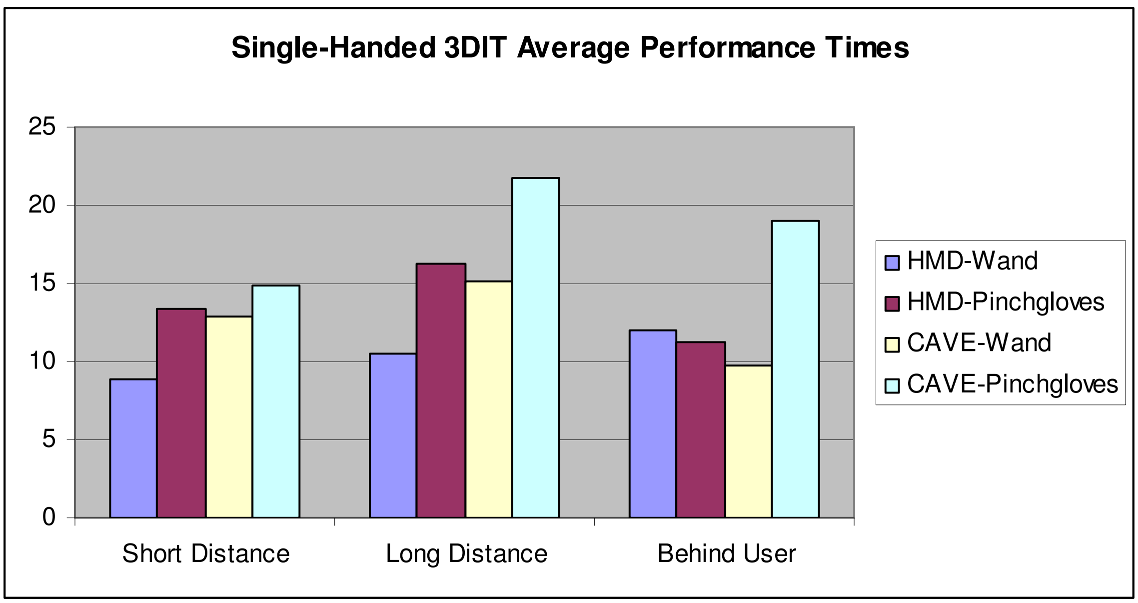

Figure 8: Mean task completion times for all participants using single-handed techniques. The task types were divided into three categories: Short Distance, Long Distance, and Behind User.

4.8.2 – Voodoo Dolls/Wand/Tracked Mouse

Although the Voodoo Dolls interaction technique did not suffer in usability when varying the display, it was plagued with usability issues dependant on the device. Once again, the Voodoo Dolls technique is an exemplary confirmation of the hypothesis that 3DIT characteristics can affect migratability because of assumptions made on their primary VE systems. The primary system for the 3DIT was a combination of an HMD and Pinch Gloves. However, when the 3DIT was used with a wand and tracked mouse, usability dropped. Following we show results that support this claim and interpret qualitative data for reasons for the changes in usability.

We start with the subjective rankings because they provide an ordered general preference of the user. The conditions that involved the wand/tracked mouse were ranked third and fourth, alternatively, by both participants that used the Voodoo Dolls 3DIT.

In looking at the average subjective ratings of Voodoo Dolls over the conditions (figure 9), we can see that both the conditions involving the wand/tracked mouse produced the most frustration among users. We also see that they were both rated lowest in the areas of physical comfort and ease of use. This evidence, along with the users' subjective rankings show us that there is a usability issue occurring when Voodoo Dolls is used in conjunction with a wand/tracked mouse instead of Pinch Gloves.

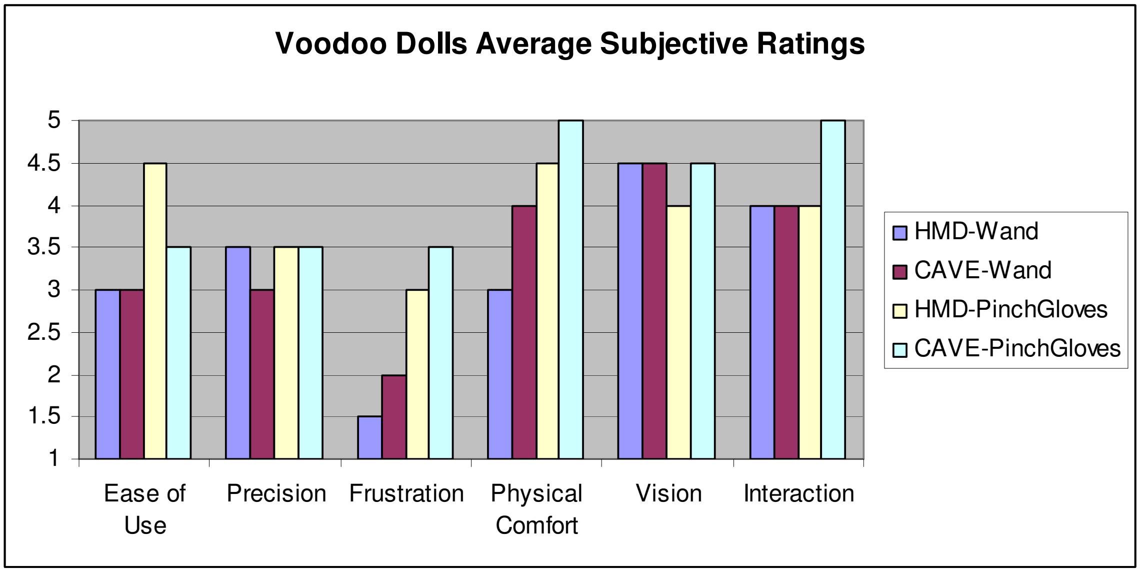

Figure 9: Average subjective ratings of the Voodoo Dolls 3DIT. The "Interaction" axis represents the user's subjective feeling on their ability to interact with the environment. The "Vision" axis represents the user's subjective feeling on their ability to see their environment. Refer to Appendix A2 to see the form used to collect these metrics. A smaller frustration rating indicated more frustration in this graph to consistently show that lower data points are badly rated.

The 3DIT characteristics that were affected most by the change in input device were the two handed interaction and direct selection/manipulation metaphors. Being a two handed manipulation 3DIT that leverages proprioceptive senses, Voodoo Dolls requires users to position their hands in relation to one another in order to coordinate likewise manipulations in the environment. Because the technique also uses direct selection and manipulation metaphors, users find themselves needing to maneuver their hands in close quarters to each other. Having been developed specifically for the Pinch Gloves, we do not see any usability issues associated with this 3DIT characteristic. However, the participants that used this technique with the wand and tracked mouse actually bumped their devices into the front of the HMD and against each other more than several times. One user comments on using the wand with HMD, "Some difficulty with hands running into each other…made easier by switching techniques."

Also because the 3DIT is a two handed technique that employs a direct manipulation metaphor, Voodoo Dolls allows the user to pass the context of manipulation between the dominant and non-dominant hands. This is an important feature of the 3DIT itself that lets the user avoid awkward or crossing hand positions by repositioning the "grip" on the context of manipulation. Using the wand and tracked mouse, users found their selves using awkward hand positioning and blocking out their own trackers by crossing their hands more often, leading to difficulty in completing the tasks and greater levels of frustration. The major differences in devices used that could explain these phenomena were device form-factor and affordances.

The Pinch Gloves are not altogether a bulky device. They are cloth gloves donned on the hand allowing freedom of movement and not adding much to the space inhabited by the hand itself. However, the wand and tracked mouse are devices that must be held and contribute to the overall volume occupied by the hand. The wand is especially obtrusive and users do not have the kinesthetic ability to avoid wand-to-hand contact as well as they do to avoid hand-to-hand contact. Also related to form-factor was the grasp used with the devices. The Pinch Gloves were a worn device that allowed the user to use the tips of their fingers to manipulate objects that "hang" from their hands. In comparison, the wand and the mouse both require the user to have an underhanded grip on objects. In this way, the dolls become "cradled" in the user's hand. It is much easier to position objects within tight spaces if the object is in a "hanging" position rather than if the object is in a "cradled" position. It is the difference between repositioning a piece on a chessboard by pinching the head with two fingers or cupping the pieces in the palms. So, because of the differences in form-factor of the device, users tended to get frustrated with bumping their hands and HMD.

Differences in device affordances also played a part in the migratability of Voodoo Dolls. The Pinch Gloves afford a pinch metaphor that allows for a tacit understanding of dropping and grabbing objects. Users assume or realize very quickly that they can drop and grab objects at any time and with rapid succession. They are also more inclined to pass objects between their hands. A user testifies that it "felt natural to grab an object with the gloved hand." In the primary VE system implementation, users grasped the strategies involved with Voodoo Dolls faster and easier and it was thus successful. However, when using Voodoo Dolls with the wand and tracked mouse, this was not the case. One user confirms the usefulness of the pinching metaphor through contrast to the wand saying, "[It was] a bit harder to pick up objects at times." Using the wand and mouse buttons as interfaces to the selection metaphor was not conducive to the task of passing the context of manipulation. Users were not confident that they could pass these objects between their hands quite as easily as they could with the Pinch Gloves. Thus users did not effectively reposition their context "grips" to avoid awkward or crossing hand or arm positions.

Another issue related to the grip repositioning issue described above involves the preferred interaction flow of the Voodoo Dolls technique. Most users would select objects directly out of the environment with their dominant hand rather than out of the context of manipulation in their non-dominant hand. As a result, users tended to create a doll in the dominant hand before they select a context to put it in. Using this interaction flow the user cannot reposition their grip of the context in their non-dominant hand because their dominant hand is preoccupied with a doll. This, compounded with the form-factor issues described above, led to bumping hands and tracker blockages in some cases where both hands needed to occupy the same space.

Users of the Voodoo Dolls technique find it useful to reposition the grip of their context of manipulation to avoid trying to have two hands occupying the same space. But the affordances of the wand and mouse and the interaction flow of Voodoo Dolls do not facilitate this. Thus we saw many users crossing their arms to complete a task which sometimes led to their trackers being blocked. Then when the tracker was blocked they would experience some tracker error and either find difficulty in completing the task or become less confident in the usability of the system.

We are not able to reinforce the subjective evidence with objective performance metrics. We feel that the Voodoo Dolls 3DIT is an involved and complex interaction to perform. We feel that a good portion of the time spent performing a task involved a large amount of cognitive load on the user. This in addition to the fact that only two participants performed this technique leads us to believe that variability in these results is too high to infer trends from the performance data.

However, we present another form of quantitative objective reinforcement in lieu of performance data. In reviewing observational data, we have identified several recurring usability issues that resulted from critical incidences. We have quantified this data and illustrated it in figure 10. Voodoo Dolls seems to be afflicted by odd hand positions, bumping hands, and tracker blockage, when used with the wand and tracked mouse devices.

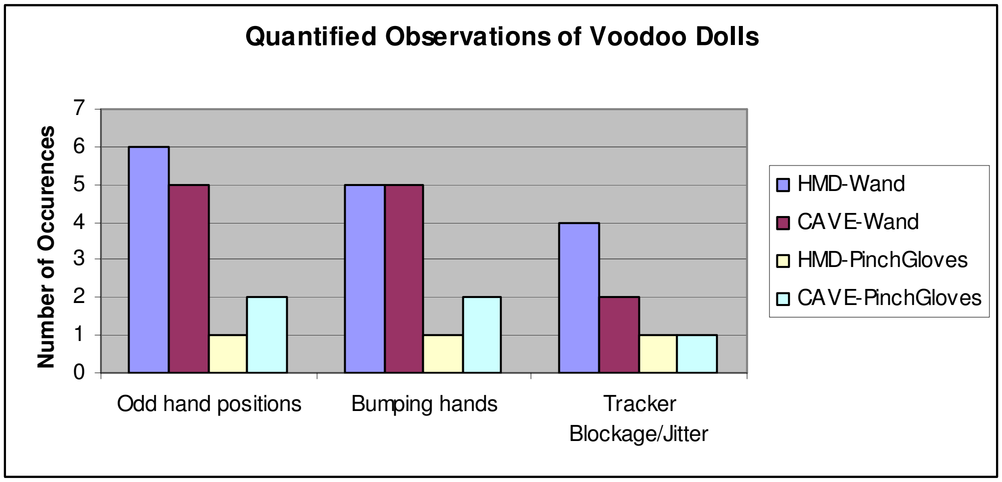

Figure 10: A graph comparing the number of occurrences of specific usability issues.

4.9 – Other Observations

Here we present other general observations that resulted from the study. These observations identify usability issues that arise from characteristics and effects of displays, devices, and 3DITs.

4.9.1 – Observations Regarding Display Characteristics

Differences between the HMD and CAVE were mostly perceived when trying to rotate the view in the cave, use corners of the CAVE, use the top or bottom of the CAVE, or in using the stereoscopic view.

The stereoscopic view in the CAVE helped significantly when a user was trying to figure out certain orientations of an object at close range. One user comments, “I felt like I could visualize the depth and size of the objects much more accurately [with the CAVE] than with the HMD.” Voodoo Dolls seemed to perform well in the CAVE environment because users tended to like the stereo view that the CAVE could offer. Near-space interactions were highly preferred by users in the CAVE because they could easily identify depth and orientation of objects at close range.

However, the benefits of stereopsis diminished at the corners and at the tops of the CAVE screens. One user confirms this saying, “Things that were elevated were harder [and] corners could interfere.” This somewhat contradicts previous work done by Manek with WIM where users were comfortable working with the WIM in the CAVE seams [34]. However, in his experiment a stereo view was not used in conjunction with the cave. In any case, this preference of users generally led to an elevated occurrence of unnecessary manual view rotations because they were constantly centering their interactions on the CAVE screens.

Some strategies were used to avoid device form-factor issues such as weight, bulkiness, or obtrusiveness. The HMD is a heavier, more obtrusive device than the glasses used for the CAVE. One participant's thoughts on the usability of the HMD were, "[The HMD] put a lot of strain on your neck and shoulders." The HMD also has more and thicker cables than those at the CAVE. Sometimes users would turn their bodies opposite of the shortest angular rotation to avoid cable entanglement. The cables became problematic in the HMD because users had to physically turn towards the back of the room where the cables were located and because the cables were not visible to the users. In the CAVE, not only is the user not required to turn 360 degrees, but the user could also see cable issues coming or happening and could avoid or evade them.

4.9.2 – Observations Regarding Devices Characteristics

Differences between the Pinch Gloves and the wand/tracked mouse combination usually pertained to differences in form-factor and device affordances.

It was observed that Pinch Gloves used near the body are easier to use than the wand near the body. Users had trouble rotating and placing objects when their wand was pressed against their body but did not exhibit the same difficulty with Pinch Gloves. This is probably caused by the obtrusiveness of the wand device. Its sensors extend out from the user's hand, while the tracker attached to the Pinch Gloves is relatively out of the way on the back of the user's hand.

Participants usually performed noticeably better with the Pinch Gloves over the wand. Rotations were more intuitive. The difference may be caused by differences in the metaphor that the devices afford. A contrast of this metaphor is seen in the comments of participants. One said that manipulating with the wand was "like moving something around with the end of a stick!" Another participant using Pinch Gloves likened it to "picking up coins."

The natural pinching metaphor is very strong and it encourages people to use multiple rotations to accomplish orientation tasks. With the pinching metaphor, they have more confidence in dropping and reselecting objects.

Missed selections, however, happened more often with the Pinch Gloves. Participant one commented, "[The] pinch action was not very precise." This is not currently labeled as potentially severe because if the Pinch Gloves lead to missed selections simply because the pinching metaphor allows for quick grabbing and dropping events, then a missed selection won't become a huge factor in overall task completion time because any number of grabs can be executed in small amounts of time. This may just lead to higher frustration levels in users if they become impatient.

4.10 – Discussion

Although the above observations demonstrate significant effects of VE systems on 3DITs, the usability issues that result from these effects are not also specific to any one 3DIT. It is for this reason that we have chosen to highlight and further investigate the usability problems found on the single-handed/CAVE/Pinch Gloves and Voodoo Dolls/Wand/Tracked Mouse. At this point we would like to discuss the differences between single-handed and Voodoo Dolls 3DITs and why the usability problems on one of them don't affect the other.

Contrasting the manual view rotation issue that occurs on a CAVE/Pinch Gloves system between the single-handed 3DIT and the Voodoo Dolls 3DIT demonstrates well the interplay between the device/display independent variables. Voodoo Dolls isn't heavily affected by the usability issue that plagues HOMER and Go-Go because both Pinch Gloves are used in the technique. As a result, pinches on the left and right-hands have a good mapping with turning the view left and right. A pinch on the left-hand indicates a left view rotation and a pinch on the right-hand indicates a right view rotation. Although this alleviates the unintuitive mapping problem, Voodoo Dolls still suffers from the uncomfortable mapping problem. Users wishing to move an object from the front of the CAVE to the rear found themselves realizing that they could only make left turns. Users tended to first grab the object to move with their right-hand. Then, realizing that trying to indicate a turn with another pinch on their right-hand led to hardship and discomfort, they made left-hand turns to bring the rear of the room into view. So, Voodoo Dolls stands to gain from a change in the way the view in the CAVE is to be rotated. However, this issue was deemed not critical enough to merit further attention because the user still had the ability, albeit limited, to rotate their view towards the rear of the room and because of characteristics specific to the Voodoo Dolls 3DIT described below.

Another interesting finding that stems from a contrast between single-handed and Voodoo Dolls 3DITs is the confirmation of the existence of 3DIT characteristics that lead to varying levels of 3DIT migratability. The Voodoo Dolls technique works by attaching a context for manipulation to the non-dominant hand. This context can then be moved to any region around the user given the constraints of their reach using a direct manipulation metaphor. This characteristic of Voodoo Dolls allows the user to make object manipulations remote from the actual site of the object. This, combined with the ability of the user to so easily move the remote manipulation context to other regions of space, suppresses the need users felt for constantly manually changing their views to perform view-screen alignments. With Voodoo Dolls, users could move their interactions with the environment to the centers of CAVE screens in a natural way. In turn, the number of manual view rotations was kept to a minimum and fewer critical incidences based on unintuitive or uncomfortable device mappings occurred.

4.11 – Conclusion

From this formative experiment, we have identified two interesting findings pertaining to the problem of 3DIT migration. These findings partially support our hypothesis that the usability of 3DITs is degraded when they are migrated naively to secondary VE systems. We say partially because the 3DITs do not become completely unusable and because negative effects of secondary systems on 3DITs are not guaranteed. However, we have presented two problems associated with usability that have occurred on specific conditions of the study, thus showing that the naïve migration of 3DITs can result in degraded usability.

We have seen that the HOMER 3DIT is heavily negatively impacted by a naïve migration to a VE system involving only the CAVE and Pinch Gloves. We have seen both that the issue creates frustration and lower levels of satisfaction among users and that the issue affects user performance. We have interpreted our qualitative results to show that the issue is caused by a confusing and inadequate device mapping. The pinch mappings that control view rotation in the CAVE cause conflicts with the pinch mapping that controls object selection. These pinch mappings were arbitrarily chosen and we believe more careful consideration in design may have avoided this issue. However, the same tasks were very easily supported by the wand device even though the wand provides 6DOF tracking and discrete input similar (with respect to the developer's point of view) to the Pinch Gloves. From this finding we propose the following guideline: *Devices that provide similar input do not necessarily provide equal usability for the same tasks.*

In addition to the poor device mapping issue, we recognize that we would not have needed to create a new mapping in the first place if a new task had not been introduced. Manual view rotation tasks needed to be supported by 3DITs being used in the CAVE. As a result, we added functionality to the HOMER and Go-Go 3DITs that resulted in poor design. From this finding we propose the following guideline: *Be cautious of situations that force added functionality.*

Results from this experiment also demonstrate the impact of device on the usability of the Voodoo Dolls 3DIT. Results show elevated levels of frustration and lower levels of comfort reported by users of this technique with the wand/tracked mouse device combination. Although we could not support this finding with performance data, we have made conclusions from our observations and other qualitative data. We conclude that the bulkiness and obtrusiveness of the wand/tracked mouse devices caused users to repeatedly bump their devices together and perform difficult postures to accomplish tasks. This finding illustrates the effects of device characteristics on migratability. From these results we propose the following guideline: *The migratability of 3DITs that use bimanual interaction may be negatively impacted by device form-factor.*

We have identified issues with migratability with this evaluation. We now attempt to remedy the problem by redesigning the 3DITs involved with device specificity in the next chapter.

Chapter 5 – Redesign and Rationale

From the general hypothesis that specificity of 3DIT implementation highly affects the usability of the technique, we assert that we can redesign display and/or device specific flavors (or sub-implementations of 3DITs) to maintain usability across VE systems. This chapter discusses the display/device-specific flavors in detail and gives some rationale behind design decisions. Here we discuss the redesign of the HOMER and Voodoo Dolls manipulation techniques.

5.1 – HOMER using CAVE/Pinch Gloves

To briefly recapitulate the usability issue involved with using HOMER on a CAVE/Pinch Gloves system, a single Pinch Glove is not sufficient to make view rotations in a single-handed 3DIT such as HOMER. One problem with this is that the mapping of the rotations with the pinches on one hand may be unintuitive and confusing. Another problem associated with the mapping of pinches is that multiple pinches on the same hand are hard to perform concurrently in a controlled and comfortable way. We need to find more intelligent ways to let the user control their view in a CAVE while using Pinch Gloves and a one-handed interaction technique. Another usability issue specific to the CAVE is that in general users did not like to perform tasks at seams or edges of the CAVE screens.

5.1.1 – Redesign Considerations

To address the issue observed regarding user preference to avoid CAVE seams and edges we decided that providing a better method for manual view rotation would relieve the negative effects of this on the usability of the technique. Negative effects include impact on task completion times and higher task frustration levels. We feel that providing a faster or more intuitive way to manually rotate the user's view would enable the user to more easily center their tasks on a CAVE wall. This would lead to both less frustration associated with the manipulation tasks as well as less frustration performing the view rotation. Additionally, we must provide support for performing the main tasks as well as view rotation tasks concurrently to facilitate faster task completion times.

To provide a concurrent and more intuitive manual view rotation with Pinch Gloves in the CAVE we brainstormed several options. We consider other modes of input available including head and hand tracker data as well as a second Pinch Glove. We also considered other interaction metaphors that support the view rotation tasks that are discussed below.

The camera-in-hand interaction metaphor [78] places the task of reorienting the view in the realm of the direct manipulation metaphor. A tracked device in or on the users hand directly affects the users current view in the world. So, given that the user points the tracking device to their anterior, their view in the world will also be to their anterior. We proposed to use this metaphor in conjunction with the HOMER 3DIT by adding a second tracking device to the non-dominant hand and attaching changes in tracker heading to view heading in the same way. In effect, we would not use the translation capabilities of this metaphor and we would constrain rotations to around the vertical axis. Although this metaphor provides a means to access all regions of space around the user concurrent to manipulation tasks it has several drawbacks. The first is that the user's hand does not effectively rotate 360 degrees around the vertical axis. Our bodies naturally limit this movement. Furthermore, we are adding a second VE device to the VE system as a whole changing the HOMER technique from a one-handed technique to a two handed technique.

Other metaphors we considered to remedy our usability breakdown were non-traditional uses of head tracking studied by Kjeldskov [33]. One of these was a mapping of the head tracker that exaggerated natural view rotations by some factor. In effect, if the exaggeration factor used was two, a user that turns their head 45 degrees to one direction actually rotates their view of the virtual world by 90 degrees in the same direction. Kjeldskov reported that user's felt this interaction to be disorienting. The other untraditional use of head tracking was a mapping that used special zones at the furthest edges of the screen based on head tracker direction to rotate the view. For example when a user directs their gaze to the left edge of the CAVE screen, their view rotates left as long as they hold their gaze. However, Kjeldskov reported that this created a contention between orienting in the real and virtual worlds.

Using a toggle pinches for selection would have enabled users to complete manipulation and view rotation tasks concurrently. This is when the user must perform a pinch to select and hold an object, then make a pinch to deselect and drop and object. Task flow would have proceeded with users first performing a pinch to grab and hold an object. Once an object was selected, the object would be attached to their hand and they would not be performing any pinch gestures. They would then be free to make view rotations indicated by single pinches without needing to hold multiple pinches at once. Then the user would simply make a pinch again to drop the object to complete the task. However, the toggle pinches do not make use of the pinching metaphor that the Pinch Gloves afford and this remedy still does not address the unintuitive mapping issue.

After considering all the above interaction metaphors, using a second Pinch Glove seemed to be our best option. We believe adding a second glove to the VE system does not significantly alter the system in the same way that adding a second tracker does because Pinch Gloves come in pairs. Making pinches on the non-dominant hand can be performed concurrently with pinches made on the dominant hand with ease and comfort. Additionally, we can implement a more intuitive mapping of pinches to view rotation direction by making pinches on the left-hand map to left turns and pinches on the right-hand map to right turns.

5.1.2 – Pilot Testing the HOMER Redesign

To confirm the usefulness and effectiveness of using two Pinch Gloves instead of one for manual view rotations in the CAVE we had two pilot testers use an initial design of the technique and provide initial thoughts or feedback. The initial iteration of the design used a left ring finger pinch for left view rotations and a right ring finger pinch for right view rotations. One of these pilot testers confirmed that while the pinch to direction mapping was very straightforward, they felt like view rotation control needed to be completely mapped to a separate glove from the selection control so that uncomfortable pinches could be avoided during the concurrent performance of both the main task and view rotation tasks.

The second iteration of design then separated view rotation from the selection hand completely and gave the user left and right view rotation control in the left-hand. A pilot user confirmed the intuitiveness of the mapping. Their performance also promised successful results in a future usability study.

5.1.3 – The CAVE/pinchglove-specific HOMER flavor

The redesign still uses a pinch with the index finger on the right-hand as the selection action. It also still implements a middle finger pinch and a ring finger pinch on the right-hand for left and right view rotation respectively. This is an intuitive pinch to direction mapping assuming the user's palms are downward facing. However, the redesign adds a ring finger and middle finger pinch on the left-hand to also turn the view left and right respectively. In this way, the user can make view rotations solely using the left-hand leaving the right-hand completely devoted to selection and manipulation tasks, thus supporting concurrent main and rotation tasks and avoiding uncomfortable pinches. In this redesign, we give the decision of which mappings to use up to the user. If they feel that they are not confused with the mapping of pinches on one hand to make view rotations, then they can avoid the uncomfortable mapping problem completely. If not they can use the intuitive mapping of the ring fingers which coincidently maps left-hand pinches to left turns and right-hand pinches to right turns. However, they will not completely avoid uncomfortable pinch incidents in this case.

5.2 – Voodoo Dolls using the Wand

Regarding the usability issues associated with the Voodoo Dolls technique when implemented with a wand and tracked mouse, we must address two issues. One issue is the form-factor of the device. Bulkiness and Obtrusiveness of the devices frustrated users because they frequently bumped devices. Another issue we should address is making it easier for users to reposition their grip of manipulation contexts. This design point serves two purposes by potentially preventing awkward arm and hand positioning, and therefore reducing the effects of bulkiness and obtrusiveness.

5.2.1 – Redesign Considerations

Although we feel that making it easier for users to change grips on manipulation contexts would alleviate the effects of device bulkiness and obtrusiveness, we believe that we can still provide other features to help. So, to help remedy the form-factor problem of the devices we propose to move the representations of these devices further out in virtual space. Although this change in design would create a bit of disparity between the actual hand positions and virtual hand positions it has been successfully practiced before. Moving the avatars further out along the vectors of their hand directions would provide more space for the user to make close two-handed interactions, thus reducing hand-to-hand bumping.

To address hand-to-HMD bumping issues we perform the same avatar shift along the world's vertical axis. By shifting the virtual hands up from the real hands users don't need to lift their hands so near their faces to get better views of the Voodoo Dolls. Coincidently, this design consideration also alleviates fatigue.

We also feel that we can promote user awareness of devices by providing meaningful and true-to-form device avatars. For instance, if we provide an avatar of the wand that actually looks like the wand, users may become more aware of the dimensions of the wand and be more able to avoid bumping devices.

Figure 11: The wand and hand avatars used in the Voodoo Dolls redesign

We considered several options for providing the user with a means to more easily access various regions of space on the manipulation context. Most of these involved mapping some form of input to rotating or translating the manipulation context in the non-dominant hand so as to circumvent the process of passing the doll between hands.

One option involved tying the orientation of the manipulation context to the user's view orientation. In effect, when the user turns their body to the left, the context of manipulation also rotates to the left, making another region of the context easily available to the user. This would enable users to access different regions of the manipulation context without the overhead of a bimanual task to rotate the context. However, this requires users to also rotate their view away from the view of the main task. Additionally, it does not address the issue that occurs when the user must manipulate objects to the region of the manipulation context that is occupied by the gripping hand.

Another design we considered involved manipulation widgets attached to the manipulation context. These widgets would support a direct manipulation metaphor that allows the user to easily slide or rotate the user's non-dominant hand grip of the context with the user's dominant hand. This can make all regions of the context easily accessible to the other hand. Although this alleviates overhead time involved in passing the context between hands to perform a grip repositioning task, we still require the user here to use their dominant hand for this task rather than the main manipulation tasks.

We felt the best method to avoid passing the context between hands was to automatically start the user's grip at the leftmost edge of the context. After all, this is the predominant strategy that users developed to avoid problems with bumping and crossing arms and hands.

5.2.2 – Pilot Testing the Voodoo Dolls Redesign

During pilot testing using tasks designed for the upcoming experiment, we found that this implementation of Voodoo Dolls ran with a very low frame rate. The polygon counts of objects used in the previous evaluation did not exceed or approach any of those used in the current set of tasks so we hypothesized that the lower frame rate was caused by this rise in object complexity. As a result we reduced the objects in the tasks to being simple cubes, which resulted in better frame rates. We assume that this dramatically low frame rate is caused by the CHASM development tool that we are using as it is still in stages of implementation and not optimized for performance.

5.2.3 – The Voodoo Dolls/Wand/Tracked Mouse Flavor

Our final decision for the Voodoo Dolls redesign involved shifted and true-to-form device avatars. For the left hand we used a hand model to represent the tracked mouse device. For the right hand we used a wand model to represent the wand. These true-to-form models are intended to make the user more aware of the dimensions of the device.

To prevent the user from needing to perform grip repositioning tasks, we decided to make user's grips of the manipulation context start at the leftmost edge of the context. So every time users selected a context for manipulation, their grip of the context was out of the way.

Additionally, to alleviate fatigue and bumping devices, we have shifted device avatars up and out from the hands. Shifting the avatars along the world's vertical axis alleviates arm fatigue while shifting avatars further out from the user's hands should help avoid bumping devices.

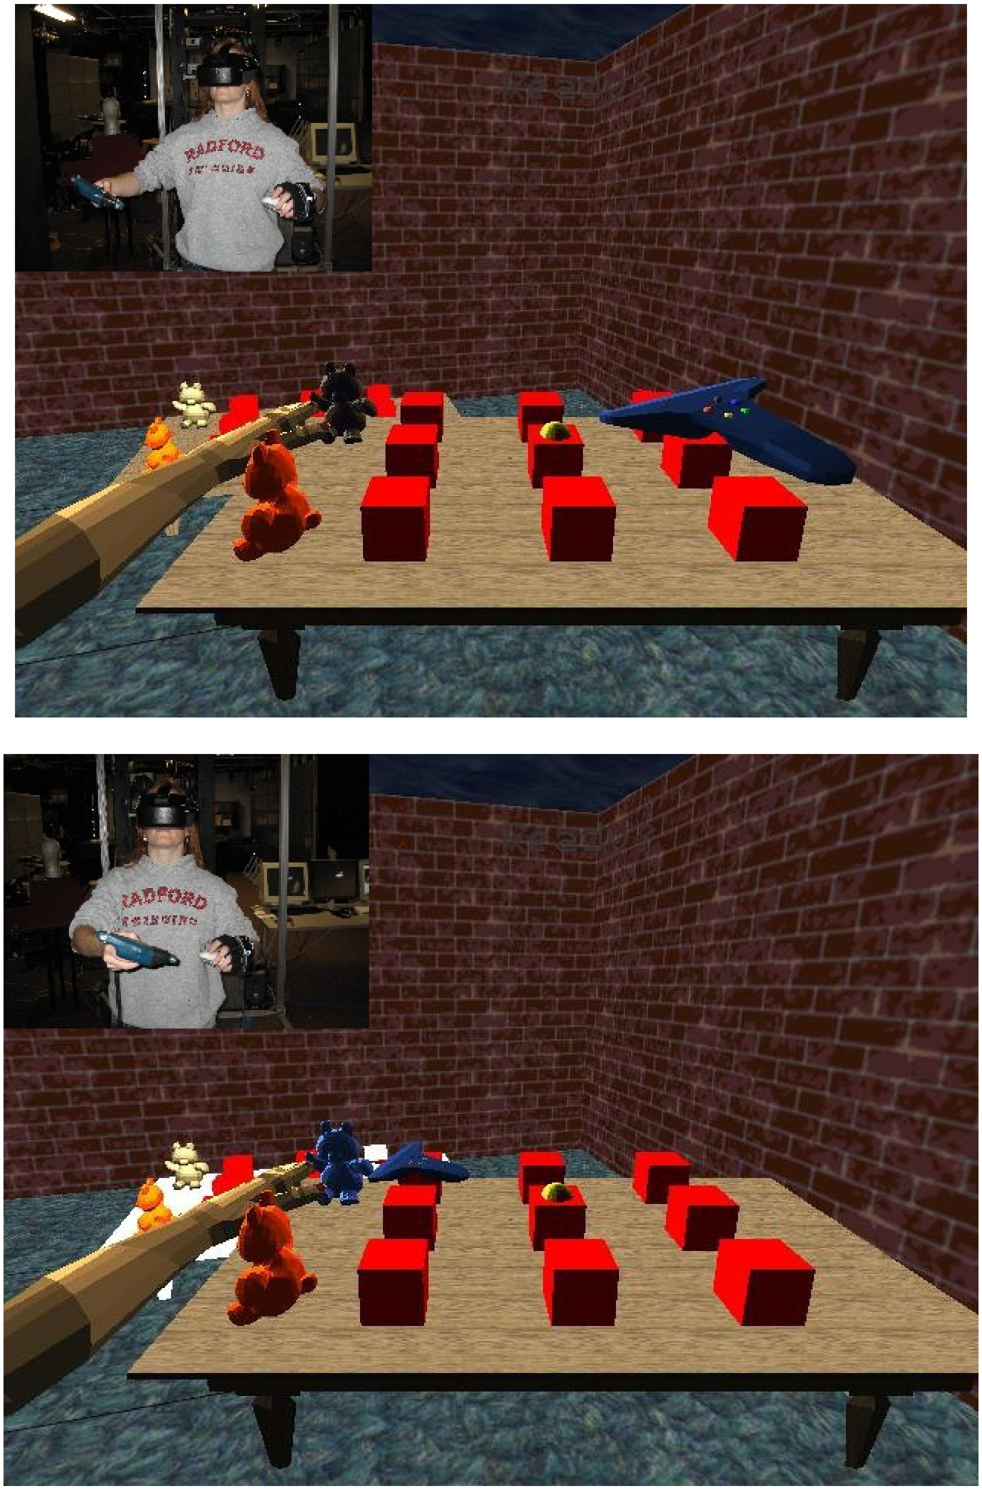

Figure 12: The following sequence of screenshots shows a user performing a manipulation with the Voodoo Dolls redesign. In the first and second pictures, they are creating the manipulation context. The last screenshot demonstrates how the user can manipulate objects on the left side of the context without crossing her hands.

5.3 – Code Reuse with CHASM [28]

CHASM is a developing system designed to support code reusability and behavior abstraction. It was used for developing all of the software involved in this project in this capacity. As such it is appropriate to discuss CHASM's usefulness to the redesign of 3DIT behavior.

Having previously defined the behavior of HOMER, Go-Go, and Voodoo Dolls in varying levels of detail using CHASM, making changes to behavior was a fairly trivial task. The nature of CHASM compartmentalizes small packets of behavior into classes called concepts. If these concepts are well designed, changes can be made to the behavior of any one concept without breaking other concepts that depend on it.

With respect to changes made to the HOMER implementation we modified concepts dealing with Pinch Glove input successfully and with little effort. A low-level concept handles discrete input from the Pinch Gloves themselves. On top of this, a concept encapsulates a pinch event and defines its behavior as a discrete input event called "Button". At a higher level some concepts have been created that map these pinch events to certain isolated behaviors that are part of techniques. For example, a concept that is implemented in the CAVE is an input concept called "ViewRotation". The "ViewRotation" concept is built on top of a concept that encapsulates a joystick or Pinch Gloves. Changes made to the HOMER technique were as simple as adding more Pinch Glove input concepts as components of the "ViewRotation" concept in a CHASM GUI.

Changes made to the Voodoo Dolls implementation were simple as well since CHASM concepts model behaviors with state machines. In order to modify the point in the object that is anchored to the non-dominant hand, the code tied to the non-dominant hand grab state had to be changed. Finding the code was easy because it was in a well-described and easily distinguishable state. Changing the visual components in this 3DIT, however, was handled directly in SVE.

5.4 – Summary

In our formative evaluation we identified some issues in migratability worth the attention of this research. We have made the claim that we can use device specificity to improve the migratability of these techniques. In this section we have looked at various options for device-specific redesign. We have chosen redesigns that we feel effectively address usability issues found in chapter four.

In the following chapter we evaluate the effectiveness of these redesigns by comparing them with the naïve migration and the original version of the 3DIT. With this comparative evaluation we can show that VE device specificity is an effective solution to the problem of 3DIT migration.

Chapter 6 – Evaluation of Display- and Device-Specific 3DITs

6.1 – Experiment Goals

In this experiment, we will be evaluating three distinct flavors (or implementations) of each of the HOMER and Voodoo Dolls 3DITs. The original implementations will be evaluated as a control condition. We are also evaluating the redesigned flavors and naïve migrations of these 3DITs. As a result we can draw three distinct comparisons from the data.

These are expressed in the following three categories. The first category is a comparison between the redesigned 3DIT to the naively migrated 3DIT. We expect to find that the redesign would be more usable than the naïve migration. This would show that specificity, on a display, device, or both, can produce migratable 3DITs. The second category compares the redesigned 3DIT to the original. This comparison will serve to show what degree of migratability we have achieved by using device specificity in our redesign. We hope to see here that the device-specific 3DIT is as usable as the original 3DIT on its primary VE system. The last category compares the original 3DIT implementation on its primary VE system with its naively implemented secondary system counterpart. Results from this comparison can confirm those from the previous experiment.

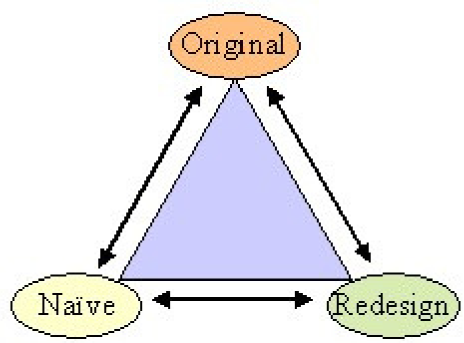

Figure 13: The three comparisons made through the display- and device-specific evaluation.

As a result of the previous migratability study, we identified and attempted to remedy two usability issues related to naïve 3DIT migrations. This experiment addresses these two issues separately to demonstrate empirically both the power of specificity and how it can facilitate 3DIT migration.

One part of the study is devoted to evaluating a display/device-specific implementation of the HOMER technique. The issue here was a usability breakdown when naively implemented on a VE system combining the CAVE and Pinch Gloves. However, through observation of users we hope to show that a CAVE/Pinch Gloves-specific redesign of the HOMER technique is more usable than a naïve implementation of HOMER. Additionally, we want to compare the usability of the CAVE/Pinch Gloves-specific redesign with the original HMD/Wand-specific technique.

The other part of the study evaluates a device-specific implementation of the Voodoo Dolls technique. The usability breakdown in this case happened when Voodoo Dolls was used in conjunction with a wand/tracked mouse device set. Empirically we hope to show that a wand-specific redesign of the Voodoo Dolls technique is more usable than a naïve implementation using the wand. Furthermore, we want to compare the usability of the wand-specific redesign with the original technique.

6.2 – Hardware

Participants used a variety of VE equipment including a Virtual Research V8 head-mounted display (HMD) with 640x480 resolution and a 60° diagonal field-of-view and a four walled CAVE (CAVE Automatic Virtual Environment). For the HMD graphics were rendered on a Power Mac G5 running Mac OSX. The CAVE runs on a five PC cluster running Fedora Core 5. The CAVE consisted of four 10' x 10' screens, each having a 1280 x 1024 resolution. An Intersense IS900 tracking system provided tracking for the head, wand, and Pinch Gloves in both systems. The input devices used were Pinch Gloves and a wand with four buttons and an analog joystick. In addition, a tracked wireless mouse was used in the non-dominant hand for two handed techniques that used the wand due to a limited availability of multiple wands.

6.3 – Software

The environment and techniques were implemented using CHASM on top of Simple Virtual Environments Libraries (SVE). SVE provided a means for interfacing with peripherals and rendering and loading the environment. CHASM was used to specify the behavior of the interaction techniques as well as the flow of the experiment. Using CHASM we have modified details of technique behavior display- and device-specific 3DIT flavors. The SVE library was slightly modified to distribute events among the DADS cluster used in the CAVE. Additionally DIVERSE support was built into the library to make tracker input easier using DTK.

6.4 – Participants

The experiment involved a total of twenty four participants. Twelve were involved in the evaluation of HOMER, the other twelve in the evaluation of Voodoo Dolls. Their mean age was 21. There were six female participants and eighteen male participants. All had near-perfect vision corrected or uncorrected. Four of the participants were left-handed the other twenty were right-handed. All but two of the users were reported novice users of VE systems. Novice users are defined as users that have never been exposed to the particular technique they were evaluating and the particular devices they were evaluating on. After being asked if they had used the 3DIT they were evaluating, the users that reported being familiar with VEs reported consistent with the definition of novice user.

We decided we could divide the study between two subject groups because the goal of the study does not call for comparisons made between the 3DIT independent variable. So, the experiment was divided into two separate studies to alleviate the total runtime of the experiment. However, both of the studies are to be carried out in the same manner aside from differences mentioned herein.

Each of the twenty four novice participants were assigned an interaction technique to evaluate among the three conditions. Twelve participants were assigned to use HOMER, and the other twelve were assigned to use Voodoo Dolls. This experiment was necessarily conducted within subjects because the goals of the study were comparative in nature. In this way, the participant could directly make subjective comparisons of the VE system-specific 3DIT with its naively migrated counterpart, and furthermore with the original 3DIT on its primary system.

Because we are using a within subjects design, data from the study should show learning patterns in the latter tasks because the participants become more comfortable with the feel of working in a VE. Again, it has been shown that novice users tend to build interaction preferences, or habits, in coordination with exposure to techniques [76]. For this reason the order of independent variable exposure to users was varied.

However, there were some limitations to the way the exposure to device pairings were varied. There was the logistical problem associated with using the HMD and the CAVE in the same study. The HMD and CAVE were not located in the same building. The HOMER part of the evaluation had to take place in both the CAVE and the HMD as the migratability problems occur in the CAVE while the primary system of the 3DIT includes the HMD. The Voodoo Dolls part of the evaluation could completely take place using the HMD because the migratability issues explored here are not dependent on the display device.

As a result a latin-squares design was used for the study on Voodoo Dolls, while a best-fit latin-squares design was used for the study on HOMER. With Voodoo Dolls, there were six distinct orderings of exposure to the independent variables. Exactly two participants experience the same ordering. With HOMER, there were only four distinct orderings available due to the logistical constraint. As a result, exactly three participants experienced the same ordering.

6.5 – Environments and Tasks

The environment created for the experiment was the same small square room used in the previous usability evaluation. The room was small enough to see all of the objects contained inside of it well and with ease. This eliminates the possibility that poor visibility of objects may confound the task of positioning objects at the extremities of the room. However, the room was not small enough to make all of the objects contained within naturally reachable. If this were the case, then the interaction techniques being studied would be unnecessary and the study would be ineffective.

In general, tasks used in this evaluation are a refined set of tasks used in the previous formative evaluation. The tasks were designed so that we could collect quantitative data for the purposes of comparing data on a per task basis [58].

6.5.1 – Tasks for HOMER Evaluation

For this evaluation we were mainly concerned with tasks involving manual view rotations in the CAVE. To this end we designed simple selection and manipulation tasks to force the participants to select and manipulate towards and from the missing back wall of the CAVE so that a manual view rotation subtask was required before the user could finish the main task. Since we wanted to control the variables involved in this activity, the tasks were made otherwise simple and straightforward. The tasks consisted of moving geometrically asymmetric objects (colored teddy bears) from a medium distance to a medium distance around the user. Additionally we loosened manipulation task target thresholds to make it easier for participants to finish tasks in the fine movement stages. In this manner, we isolated most of the total interaction time to being associated with manual view rotation tasks.

Figure 14: A screenshot of the front table in the HOMER evaluation world.

Figure 15: A screenshot of the back tables in the HOMER evaluation world.

We identified three cases of tasks involving manual view rotations that we believe generally cover all situations. The first case is moving objects from a region of the world behind the user to within the same region. This kind of task involves an initial manual view rotation and subsequent manipulation task. The second case is moving objects from a region of the world that is visible on the CAVE to a region behind the user that is not. This task involves performing the manipulation task concurrently with the manual view rotation. The third case is moving objects from the region behind the user to the region formerly in front of the user. The latter case involves two iterations of the manual view rotation subtask with the second involving task concurrency. In addition to these tasks we also implemented control tasks that only required the user to perform a manipulation without any manual view rotations at all. Refer to appendix B4 for a table showing each HOMER task and its type.

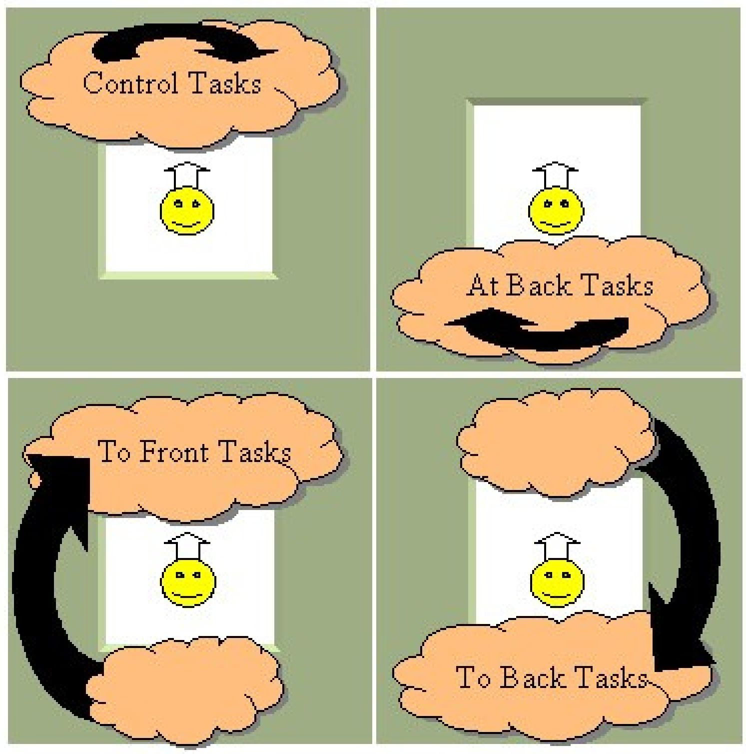

Figure 16: A diagram of the task types used in the HOMER evaluation.

6.5.2 – Tasks for Voodoo Dolls

The major goal of this evaluation was to show that we could help users avoid bumping and crossing their hands by always having the user grab the far edges of objects and providing avatars. Tasks in this evaluation thus consisted of selecting and manipulating small child objects from and to various regions of a parent object. We wished to also use simple tasks in this evaluation to minimize task or interaction technique overhead on task completion times.

For this evaluation we identified four cases of tasks that we believe effectively tests whether the redesign was successful in helping users avoid bumping or crossing arms or hands. The first case involves moving objects to and from areas of the manipulation context that do not require the user to cross or bump their hands in the original 3DIT. With a right-handed user, these tasks would entail moving an object from the right side of the manipulation context to the same side. Appropriately, we name these tasks “right side tasks”. These are also considered the control tasks. The second case of tasks involves moving objects to and from areas of the manipulation context that require the user to cross their hands while using the original 3DIT. With a right-handed user, this occurs when selection and manipulation is performed on the left side of the manipulation context. These tasks are similarly named “left side tasks”. The third case of tasks covers times when a user would bump their hands together when trying to manipulate an object. Incidentally, this usually happens in the center of the manipulation context so we name these tasks “center tasks”. At last we define a “mixed task” case that requires users to select and manipulate at mixed positions in the manipulation context. Refer to appendix B5 for a table showing each Voodoo Dolls task and its type.

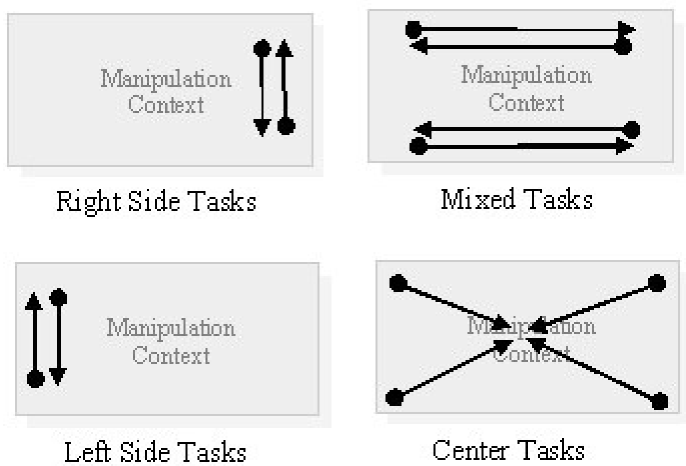

Figure 17: A diagram of the four task types used in the Voodoo Dolls evaluation.

As a side note, we attempted to use the same geometrically asymmetric objects used in the HOMER evaluation for this one. These objects contained larger polygon counts than those included in previous environments using this Voodoo Dolls implementation. Through pilot testing we found that frame rate was dramatically affected by using these objects. To combat the negative effects of frame rate we used simple cubes rather than the more complex geometries previously used. We hypothesize that development with CHASM may be the source of slow frame rate. The tool is still in stages of development.

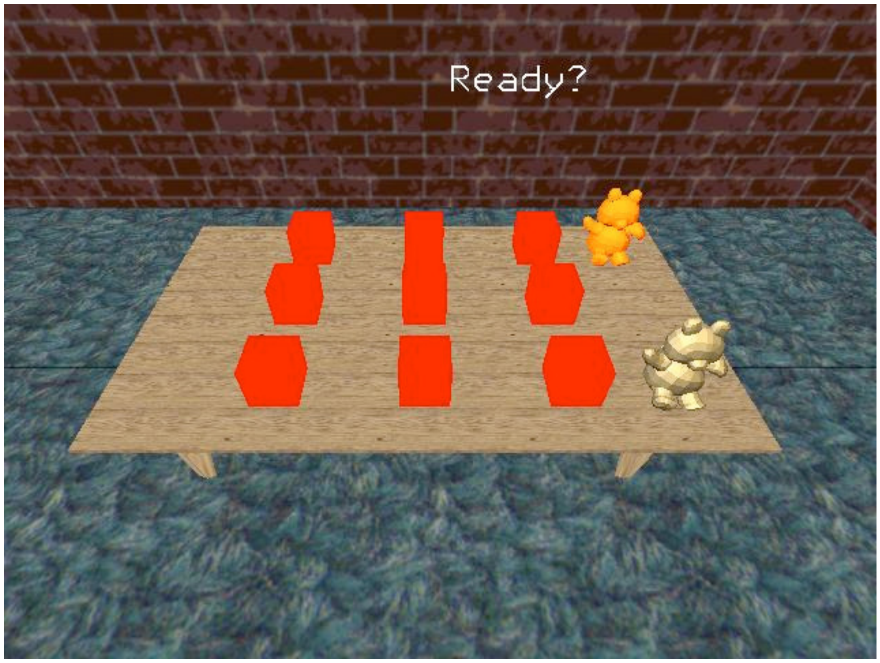

Figure 18: The table and cubes used for the Voodoo Dolls environment.

6.6 – Procedure

The experiment was performed as follows. Before each task set the user was read instructions on donning and caring for the VE equipment they were currently using. The user was then given instructions on the use of the interaction technique currently being used and was given time to familiarize themselves with it in the room environment.

The user was asked to perform a series of practice tasks that illustrated the use of the technique in a variety of situations. Then the user was given time to practice the technique in the environment at their leisure. With experimenter approval of practice task performance, the user was then asked to perform the main series of tasks.

During each task the user was asked to complete the tasks as fast as they could and to the best of their abilities. After each task the user was polled for their frustration levels and fatigue levels on a five point scale. If any of these values were outstanding, the source of change was further investigated through interview with the participant.

At the end of the task set the user was given a worksheet to survey various subjective aspects on a Likert scale. They were also asked to record some final comments, comparing each new condition to the last.

6.6.1 – Collecting Qualitative Data

Qualitative subjective comments made by the user were recorded to give an idea as to what was causing problems in usability. We did not achieve this through the use of a think-aloud protocol because the purpose of this particular investigation was comparative and user performance was given a little more importance. Using a think-aloud protocol may slow down participants. Rather, after participants completed a task, if they performed or reported anything outstanding to the experimenter, the experimenter was then responsible to further question participants of their actions if they merited explanation. Their responses were then recorded. After a condition was finished they were surveyed again for their final thoughts on the 3DIT/display/device combination as a whole and compared to previous conditions they have experience.

6.6.2 – Collecting Quantitative Metrics

Quantitative measures were an important part of the study because we wished to show comparatively that one condition was better than another. As a result, much emphasis was given to collecting and quantifying both objective and subjective measures.

Quantitative subjective measures were taken to help identify where general issues between interaction technique and device were perceived by the user. If the user perceives difficulty or ease it is also likely that they are experiencing difficulty or ease. Perceived metrics were measured by Likert scale survey items. Some Likert scale survey questions were also designed to measure perceived problems with technique-device interactions. Subjective satisfaction in the use of a technique is a component of a usable interface and can be detected in this way [43]. We can use this subjective data to back up usability claims that we make when compiling data from observation. The metrics used here were ease of use, frustration, precision, physical comfort, the participants perceived ability to interact with the environment, and the participants perceived ability to view their environment.

Quantitative objective performance measures were taken in order to give strong support to comparisons being made on the efficiency of the 3DITs. Actual accuracy and speed metrics were measured by the system in meters and milliseconds respectively for each task. These metrics could become a means of strengthening claims made based on feedback and comments provided by participants because strong patterns should be seen in a population set of this size. Efficiency in use, or the performance, of an interface is a factor in determining its usability [43]. This data, which includes task run times and accuracy values, was recorded by the system.

Observation of the behavior of participants is an important part of the usability study being conducted. In a previous formative evaluation, we have identified several outstanding observations that may impede usability. These should then be used in a comparative evaluation to show that we can fix the problems in usability [58]. Because the goals of this study are to prove the usability of one 3DIT flavor over others we wish to quantify to some extent these observations so that comparisons can readily be made. We note that in the previous experiment we developed a shorthand notation for observations of issues that occurred or were reported on frequently. Those frequently observed issues that directly pertained to the goals of this study were gleaned from results of the previous study in order to expedite observations in this study. We can easily quantify the frequency and severity of usability issues that we are specifically looking for. So we will use the following as metrics to quantify usability issues that could indicate higher or lower degrees of usability.

Homer

Table 5 shows a table used to record observations in the HOMER study. We will define each of the items below.

Table 5: Box for observations in the HOMER study.

Usability Issue (In order of importance)	# of Occurrences	Severity (1 to 5)
Avoiding Edges of Screen		
Using Difficult Finger Positioning		
Avoiding Difficult Finger Positioning		
Tracker Loss		
Cable Entanglement		
Avoiding Cables		
Heisenberg Effect		

Avoiding Edges of Screen – This is a strategy of participants in the CAVE to avoid manipulating objects on the CAVE seams.

Using Difficult Finger Positioning – This is when participants try to make multiple pinches on a single hand using the Pinch Gloves (not a trivial task).

Avoiding Difficult Finger Positioning – This is a strategy used by participants to avoid multiple pinches on the same hand in which the object to be manipulated is dropped while the user rotates his/her view.

Tracker Loss – This is the occurrence of jitter or floating hand avatars.

Cable Entanglement – When tethers bother a participant.

Avoiding Cables – When participants are observed to have consciously avoided tethers.

Heisenberg Effect – This is a phenomenon that occurs when a user unintentionally moves the object they are trying to place in the action of deselecting the object.

Voodoo Dolls

Table 6 shows a table used to record observations in the Voodoo Dolls study. We will define each of the items below.

Table 6: Box for observations in the Voodoo Dolls study.

Usability Issue (In order of importance)	# of Occurrences	Severity (1 to 5)
Bumping Hands Together		
Hands Bumping HMD		
Hand Crossover		
Preemptive Repositioning		
Post Hoc Repositioning		
Fine Repositioning		
Difficult Hand Positions		
Tracker Blockage		
Cable Entanglement		
Heisenberg Effect		

Bumping Hands Together – When a participant's input devices collide together.

Hands Bumping HMD – When a participant's input device collides into his/her HMD.

Hand Crossover – When a participant crosses his/her hands. This potentially can result in tracker loss and increases the chances of colliding devices.

Preemptive Repositioning – This is a strategy used by participants to reposition their grip on the manipulation context to prevent colliding or crossing devices.

Post Hoc Repositioning – This is a strategy used by participants to reposition their grip on the manipulation context after a collision or crossing of devices has occurred.

Fine Repositioning – This is a strategy used by participants to reposition their grip on the manipulation context in order to better perform fine manipulation adjustments.

Difficult Hand Positions – Situations where the participant is holding strenuous or difficult to perform hand positions in order to accomplish a task.

Tracker Loss – This is the occurrence of jitter or floating hand avatars.

Cable Entanglement – When tethers bother a participant.

Heisenberg Effect – This is a phenomenon that occurs when a user unintentionally moves the object they are trying to place in the action of deselecting the object.

Another measure we felt would be interesting to incorporate into this study was a ranking of 3DIT flavors. Since the study is comparative in nature it seems natural to allow the user to directly compare and rank the conditions. Users were asked to order the 3DIT flavors in terms of ease of use, frustration, precision, physical comfort, and overall preference. That is, after the entire study was conducted the participant rated each session in relation to the others. So we get see which devices pairs were the favorites and which were the least liked.

6.7 – Results

Results from the study indicate advantages to using device specificity when migrating 3DITs to secondary VE systems. We show that the HOMER redesign is more usable than the naïve implementation.

The study also reconfirms the low degree of migratability of the single handed techniques. We provide strong comparative evidence in this experiment that reinforces our findings of the first experiment that usability of the single handed technique breaks down on a VE system composed of the CAVE and Pinch Gloves.

However, there are also several areas in the results that are contrary to our expected findings. The Voodoo Dolls study results are mixed and do not show strong evidence towards our expectation that device specificity can maintain usability across VE devices. But we do show issues brought up in the previous experiment reoccur in this experiment and that our Voodoo Dolls redesign has addressed these issues.

6.7.1 – HOMER Results

Results indicate strongly that the HOMER redesign is preferred over the naïve migration. Subjective user preferences show that on average users felt the naïve migration of the HOMER 3DIT was not as easy to use or as precise as the HOMER redesign. They also show that users felt the naïve migration to be more frustrating to use. The subjective rankings also drastically show that users preferred to use the redesign over both the naïve migration and the original 3DIT on the primary system (HMD/wand). This coincides with subjective comments made by users about their condition preference with one saying, “I preferred [using both Pinch Gloves] over the single Pinch Glove” and another saying, “I felt that [using both Pinch Gloves] was a lot easier than the technique using the wand.”

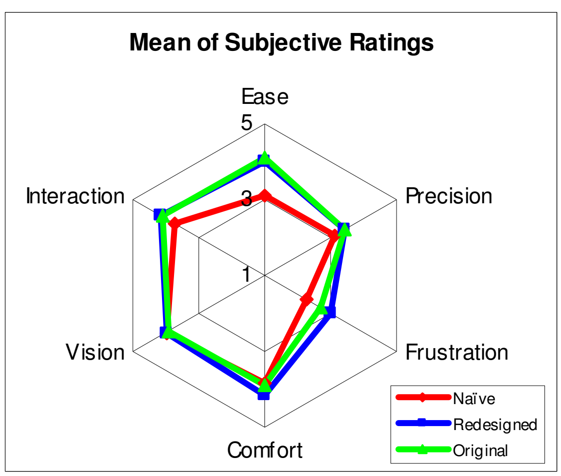

Figure 19: Average subjective ratings of the HOMER 3DIT. The “Interaction” axis represents the user’s subjective feeling on their ability to interact with the environment. The “Vision” axis represents the user’s subjective feeling on their ability to see their environment. Refer to Appendix B2 to see the form used to collect these metrics. A smaller frustration rating indicated more frustration in this graph to consistently show that data points closer to the center were badly rated.

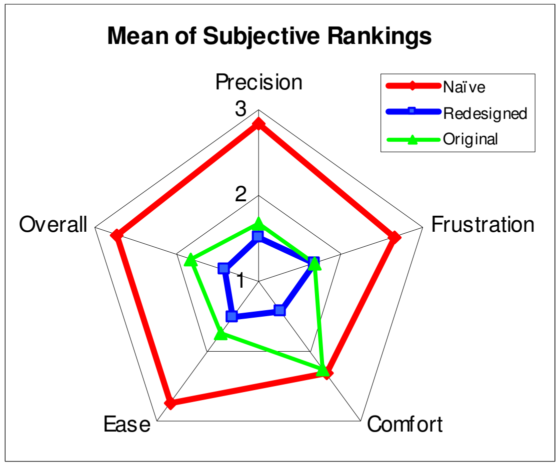

Figure 20: Average subjective ratings of the HOMER 3DIT. Data points closer to the center are better ranked. Refer to Appendix B2 to see the form used to collect these metrics.

Reasons for the success of the redesign can accurately be attributed to the remapping of view rotation control to a second Pinch Glove. When participants were constricted to rotating their view with a single glove while trying to carry an object they became frustrated. One participant vents this frustration, "What really bothered me about this technique was I could not rotate and hold the bear at the same time." It was observed that participants frequently dropped the object to rotate the view, which usually led to slower performance, or attempted to rotate the view and hold an object with a difficult pinch configuration, which usually led to discomfort and degraded performance. Figure 21 shows a comparison of the number of times these observations were made. It is clear that the HOMER redesign avoids discomfort and degraded performance by facilitating users to complete the manipulation tasks concurrently with view rotation tasks. This is clearly expressed by one user, "This time it was easier because there was both a glove to rotate the room and a glove for selecting the bears. I felt like I was able to move the bears quicker."

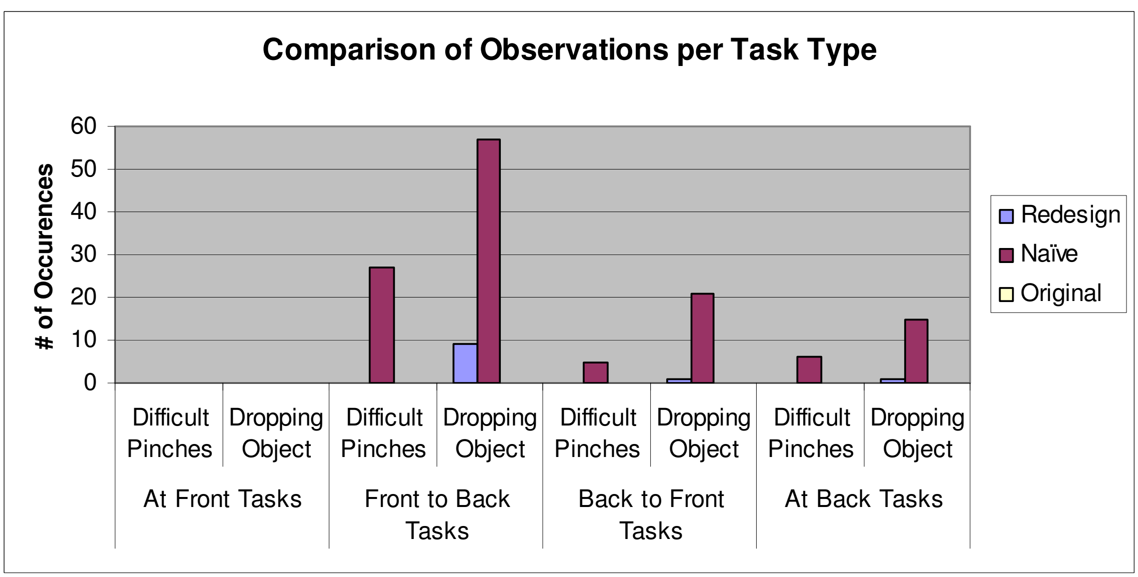

Figure 21: A comparison of the amount of times an observation was made regarding "Difficult Pinches" and "Dropping Object".

We conducted an ANOVA with replication of 3DITs against tasks. Results showed an interaction occurred with 95 percent confidence ($F(2, 11) = 2.428044$, $MS = 127.0046$, $p = 0.000384$). The ANOVA can be found in appendix B6. Following this we performed pair-wise t-tests of each task for each 3DIT to find out which tasks performed with significant difference on the 3DITs. We also performed a t-test comparing the means of the 3DITs regardless of task.

From the t-tests we see trends in task completion times that support the hypothesis that the HOMER redesign is more usable than the naïve migration. In figure 22, we see that for those tasks that produced significantly different means (the tasks with asterisks) the mean task completion times are faster when participants used the HOMER redesign. For example, for task six the mean performance time was over 13 seconds faster with a 95 percent confidence level. We do not see many significant results in this comparison. This is due to relatively high variability in performance times that we believe is due to the participants' varying spatial abilities. However, even without significant difference, a clear trend can be seen in figure 22 supporting our hypothesis.

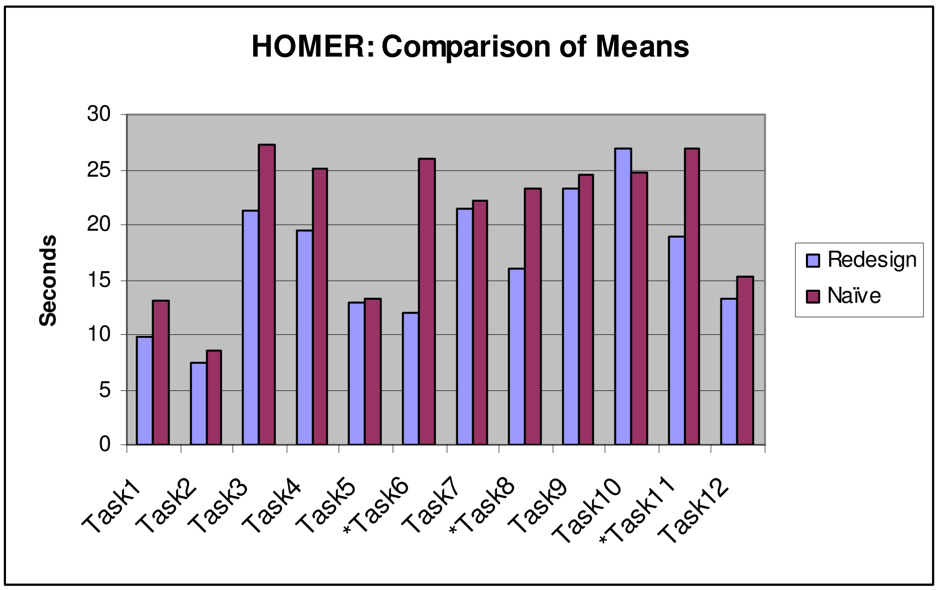

Figure 22: A comparison of task completion time means over all participants between the HOMER redesign and the HOMER naive implementations. The asterisk beside a task label denotes statistical significance ($p = 0.05$). For a table of the data refer to appendix B6.

In the same way, we see trends in task completion times that reconfirm our claims from the previous study that the naïve migration of HOMER is less usable than the original implementation. In figure 23, we see the original implementation of HOMER consistently producing faster performance times than the naïve implementation.

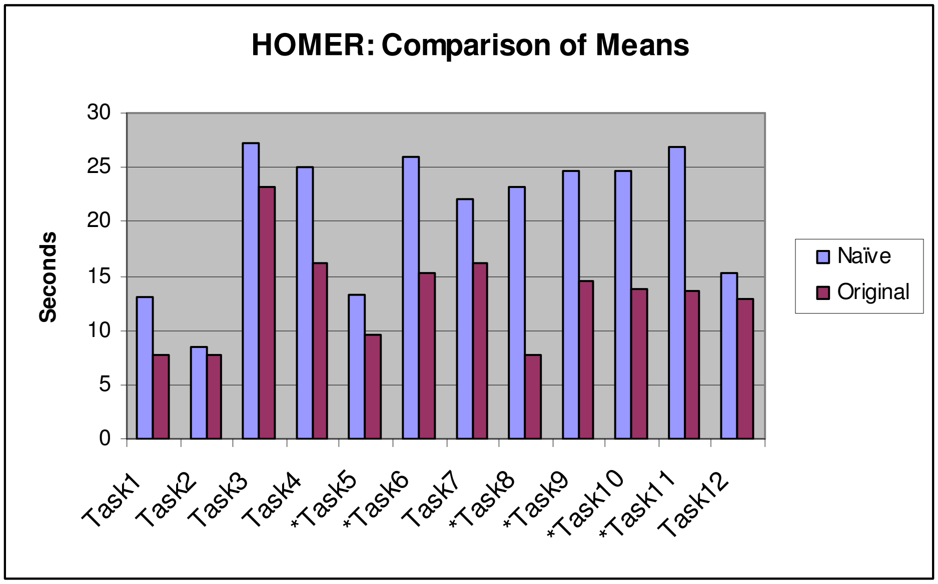

Figure 23: A comparison of task completion time means over all participants between the HOMER naïve and the original HOMER implementations. The asterisk beside a task label denotes statistical significance (p = 0.05). For a table of the data refer to appendix B6.

We also found that users tend to prefer the redesign over the original implementation. However, performance metrics do not coincide with this claim. Figure 24 shows statistical significance in favor of the original implementation for several tasks. Contrary to this, a popular opinion among the users was that, "the Pinch Gloves were easier because you could spin the room using them instead of having to twist your head." It seems that users tended to prefer a lazy style of interaction as opposed to the more realistic style that involves actually turning. Additionally, users complained about cable entanglement with the HMD being a frustration, inspiring one participant to comment, "The wires to the headset got in the way. If it was wireless I would have found the tasks extremely easy." When users are forced to turn their bodies to make view rotations, they also risk running into tethers. Of note, these sentiments can affect both comfort and frustration levels. Coincidently, the comfort and frustration ratings and rankings reflect better on the HOMER redesign than the original implementation.

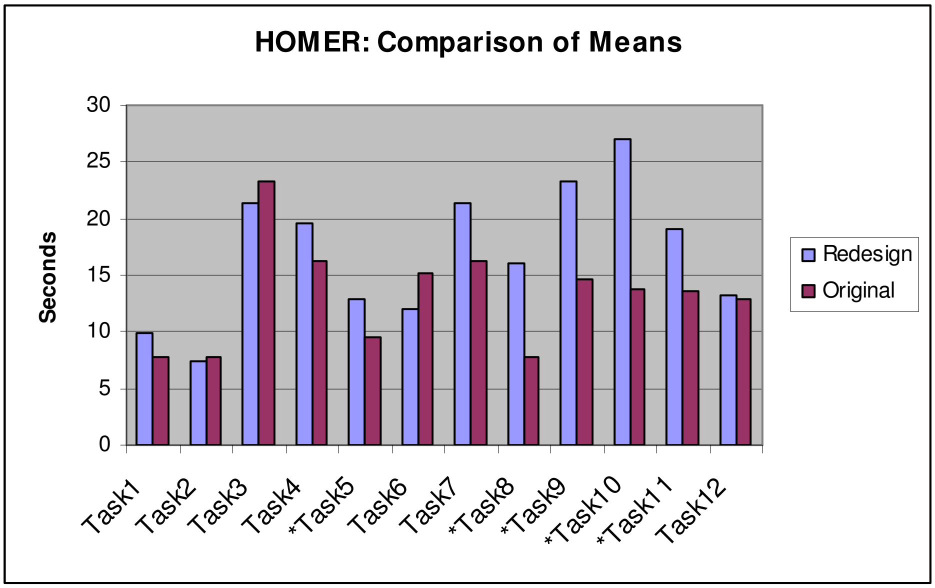

Figure 24: A comparison of task completion time means over all participants between the HOMER redesign and the original HOMER implementations. The asterisk beside a task label denotes statistical significance (p = 0.05). For a table of the data refer to appendix B6.

6.7.2 – Voodoo Dolls

Results are mixed in this part of the study with the performance metrics contradicting other indications of usability. We found that users did not on the whole prefer the redesigned Voodoo Dolls technique to the naïve migration as we previously hypothesized. Subjective ratings and rankings show mixed feelings on this comparison.

The redesigned implementation is actually generally ranked and rated worse than the naïve implementation on all areas barring comfort (Figures 25 and 26). The better comfort rating could be attributed to the feature of the redesign that offsets virtual hand avatars up from the actual positions of the hands. This feature lets the user perform interactions with their hands at a comfortable level in the air. One user actually commented in regards to this feature, "I was able to hold my hands lower, which was more comfortable." However, the original implementation, which does not implement this feature, is ranked and rated superior in comfort.

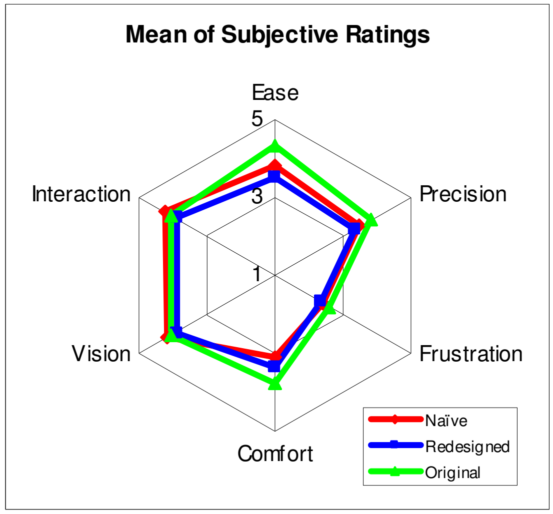

Figure 25: Average subjective ratings of the Voodoo Dolls 3DIT. The "Interaction" axis represents the user's subjective feeling on their ability to interact with the environment. The "Vision" axis represents the user's subjective feeling on their ability to see their environment. Refer to Appendix B3 to see the form used to collect these metrics. A smaller frustration rating indicated more frustration in this graph to consistently show that data points closer to the center were badly rated.

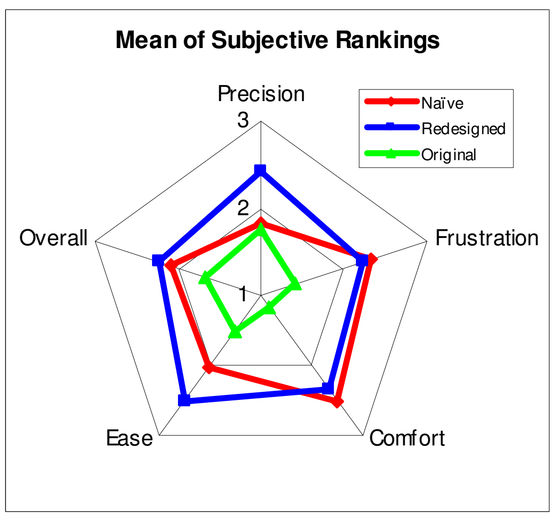

Figure 26: Average subjective rankings of the Voodoo Dolls 3DIT. Data points closer to the center are better ranked. Refer to Appendix B3 to see the form used to collect these metrics.

We believe that the alternative starting grip point on the object is responsible for this indication of comfort. The Voodoo Dolls redesign enabled users to access all areas of the manipulation context by initializing a grip on the leftmost side of the context. This helped users to avoid crossing their hands, bumping their devices, and repositioning the manipulation context. One participant agrees saying, "…I preferred picking up the table from the side vs. the middle."

In the first evaluation of this 3DIT, users frequently crossed or bumped their devices to access areas of the manipulation context. This was because their left-hand grip on the manipulation context was in the center of the area. To avoid the discomfort of crossing or bumping their devices, participants used a strategy of repositioning their grip on the manipulation context. This was a coordinated two-handed task that involved a few selection and manipulation tasks in and of itself. Participants used this strategy in three ways. They would first cross or bump their devices and realize they could benefit from a repositioning. We call this "post hoc repositioning". Or, participants realized they could avoid crossing or bumping their hands and performed the repositioning strategy before the main task. We call this "preemptive repositioning". Many participants realized they could use the repositioning strategy to obtain better fine manipulation control by avoiding difficult hand positions. We call this "fine positioning". We observed and counted the use of these strategies along with other observations such as bumping or crossing events so that we easily compare whether or not our redesign imposed the desired effects.

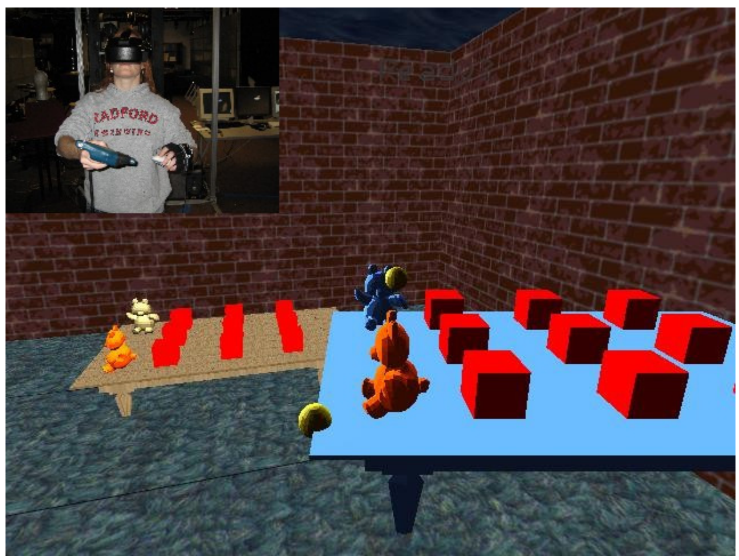

Figure 27: The following sequence of screenshots shows a user repositioning their manipulation context. In the first screenshot, the manipulation context is in the left hand. The user then grabs the context out of the left hand with the right hand in the second screenshot. The third screenshot shows how the user can now reposition their left hand grip farther to the left of the manipulation context. The last screenshot shows the ease with which the user can now manipulate on the left side of the manipulation context.

In the design of our experiment we implemented four kinds of tasks. Here we give them an overview. Right side tasks are the control tasks of the experiment. Participants would find no need to cross their hands or reposition their grip in these tasks. Mixed tasks, left side tasks, and center tasks all involve the selection and/or manipulation of objects to or from areas of the context that require the user to cross their hands or reposition their grip. From figure 28 we can see that the redesign, in its intent, was successful.

In figure 28, we see that our redesigned Voodoo Dolls implementation fulfilled its intent. In the mixed, left side, and center tasks, participants drastically performed less preemptive and post hoc repositioning tasks. We also see that this resulted in reduced device bumping or crossing events.

Incidentally, from figure 28 showing the left side and center tasks, we can also see how the Voodoo Dolls is affected by using bulky and obtrusive devices such as a wand and tracked mouse. Users were not aware of the obtrusiveness of the wand and tended to bump their devices.

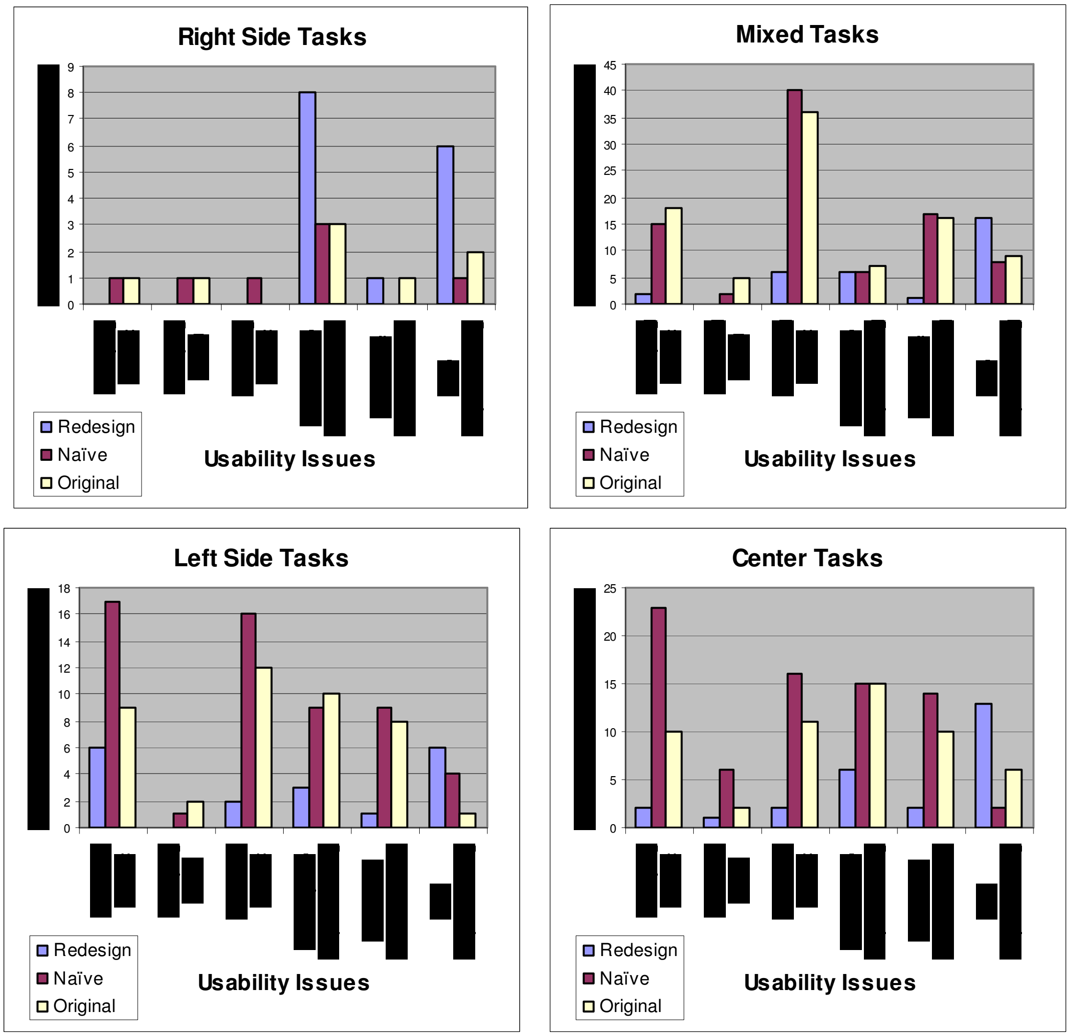

Figure 28: A comparison of the total amount of times an observation was made. These are divided into four task types.

Also from figure 28, we can see that we have introduced an unintentional usability issue that may have kept the Voodoo Dolls redesign from being more successful. In all tasks we notice a large amount of fine repositioning being performed by users. Fine repositioning is a strategy used by participants to gain better control of manipulations by avoiding difficult or straining hand positions. We believe that by placing the user's initial grip of manipulation contexts on the leftmost side that we have made manipulations on the right most side harder to perform. It seemed that participants performed manipulations better when their hands were close to each other. To support this claim we draw from figure 29 showing large amounts of preemptive repositioning and fine repositioning. A participant trying to manipulate objects with our redesign in these right side tasks would have to perform the manipulation with their hands at the width of the entire manipulation context. As a result, many participants found it useful to reposition the manipulation context. We can also support this claim through contrast with left side tasks.

Figure 29 shows that repositioning strategies where used the least on left side tasks where the user was only required to make a selection and manipulation on the leftmost side of the manipulation context where their grip was already located.

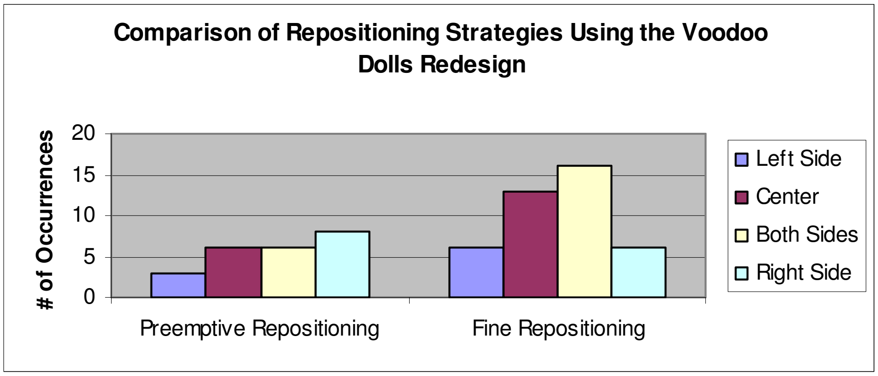

Figure 29: Shows a comparison of the number of observations made for four task types on the Voodoo Dolls redesign.

We conducted an ANOVA with replication of 3DITs against tasks. Results showed an interaction did not occur at a 95 percent confidence level ($F(2, 11) = 1.177982$, $MS = 294.1186$, $p = 0.263441$). The ANOVA can be found in appendix B6. Following this we performed a t-test comparing the means of the 3DITs regardless of task.

In table 7, we see that the redesign technique performed poorly when compared with the naïve migration. In spite of these discouraging results in performance, we would like to point out that there are several other facets of usable interfaces aside from performance. Good performance does not necessarily equal usability. We have plenty of qualitative subjective evidence that reinforces the usefulness of the Voodoo Dolls redesign. Several users commented on the usefulness of having raised virtual hand avatars with one saying (after he had used the redesigned technique), "It agitated me that I hand to raise my hands up high or look down low just to see the cursors." We can pin the changes involving the virtual hand and wand avatars to be partially responsible for low device bumping counts using comments like, "Compared to the basic wand I found it nice to have the hand present and to allow objects to rotate in respect to it as opposed to the left ball cursor."

Table 7: T-Test comparing the means of the 3DITs

Redesign		**Naïve**		**Redesign Vs. Naïve**	
Mean	Variance	Mean	Variance	t-statistic	p-value
29.23107	369.9127	18.76327	130.0893	6.288454	3.66E-09
Redesign		**Original**		**Redesign Vs. Original**	
Mean	Variance	Mean	Variance	t-statistic	p-value
29.23107	369.9127	24.79618	355.4092	2.148558	0.033354
Naïve		**Original**		**Naïve Vs. Original**	
Mean	Variance	Mean	Variance	t-statistic	p-value
18.76327	130.0893	24.79618	355.4092	-3.89488	0.00015

Table 7 also shows the surprising result that the naïve migration also performed better than our original Voodoo Dolls implementation. Looking at comments provided by participants we might blame this on the devices capacity to let participants rotate objects to complete the tasks. One participant says, "Comparing the advanced wand to the Pinch Gloves, I found it easier and more precise because I had something in my hand to rotate. This let me have more of a feel for how much I was moving/rotating…" Other participants agree saying, "My complaints with the gloves are that some of the rotations were hard to perform" and, "The Pinch Gloves seemed more natural for grabbing and releasing objects, but for turning the world and table it seemed a bit more tedious." Figure 25 shows that the naïve migration was rated on average slightly better than the original implementation (naïve = 4.25, original = 4.083) with regards to the interaction metric. This metric was designed to measure how easy participants found moving their virtual hands in the VE. This is only a small difference and all users did not agree on their comments regarding the ease of rotation with their devices with one saying, "It seemed more natural to move the objects." However, we would argue that users may actually have had more rotational control with the wand than with the Pinch Gloves. This rationale stems from the comment made by one participant, "Moving the bears along the [vertical] axis was easier with the wand than the gloves." The grasp on the wand is different from how the Pinch Gloves were used. In order to avoid tracker blockage, users were required to perform interactions with their palms faced down. In this posture users have very limited wrist movement around the vertical axis. However, the wand is held in a posture that enables a higher degree of wrist movement.

Anecdotal evidence reinforces the findings in the previous study that there is a problem in the naïve migration of the Voodoo Dolls 3DIT onto a wand and tracked mouse device set. One user explains that when using the naïve implementation, "for some reason I had more trouble with my hands hitting each other than I did with either of the other two experiments." Along the same lines another user comments that the naïve migration was, "Considerably frustrating when trying to place both hands in the same location. [It was] much more of a problem because the hand-held devices are bigger and more awkward to handle." Additionally, the participant's subjective average rankings do show that the original technique is preferred in all areas.

We clearly see that there are usability problems associated with the naïve migration of Voodoo Dolls to a wand/tracked mouse device set. We also see that the problem occurs specifically when users try to make two-handed manipulations within close proximity using bulky or obtrusive devices. We have attempted to design out the usability issues described above with a device specific implementation of the Voodoo Dolls 3DIT on the wand/tracked mouse device set. Unfortunately, identifying usability issues doesn't show us how to fix them or even if they are fixable. In our attempted redesign we were successful at engineering out the initial usability issues, namely bumping/crossing devices and reducing the amount of grip repositioning tasks. However, we did introduce other usability issues that may have kept the redesign from being successful, namely necessitating grip repositioning tasks on right side manipulation tasks.

User-centered usability engineering is a repeated process that involves iterative design and evaluation [47]. At this point we would like to suggest further redesign. Since we have seen that users tend to prefer close proximity asymmetric bi-manual manipulation, we should make the entire context of manipulation fit within this region. We feel that the context of manipulation in our redesign of Voodoo Dolls was too large to be handled with a grip on the leftmost side and that users could avoid repositioning strategies altogether if the context of manipulation remains in a comfortably reachable volume of space offset from the non-dominant hand.

6.8 – Conclusion

We set out in this experiment to demonstrate the potential advantages of using device specificity for 3DIT migration. From this comparative experiment, we have mixed results, but found one instance in which a device-specific redesign was beneficial. We also confirmed most of the findings from our original evaluation.

In the study involving the HOMER 3DIT, we found our redesign both performed and was rated better than the naïve migration. This result supports the use of device specificity for better migratability. The results also show differences in performance between the original implementation and the naïve migration that support our findings from the formative experiment.

The Voodoo Dolls results, however, did not show our device-specific redesign to be effective. Subjective satisfaction and performance times of participants do not indicate that the redesign is better than naïve implementation. In the face of this, we have shown through observations that our redesign successfully addressed and corrected the usability issues that we were targeting. However, our redesign unintentionally introduced other problems that kept it from performing better than the naïve migration. Specifically, our redesign reduced the occurrence of manipulation context grip repositioning tasks associated with crossing or bumping interaction devices. Characteristics of the redesign, however, caused users to perform the same grip repositioning tasks for easier manipulation. These results suggest that iterative design process, much like that proposed by Gabbard et al. [47], is necessary for usable device-specific design.

Chapter 7 – Conclusions and Future Work

In this chapter we revisit the goals we set out at the beginning of this thesis and reflect on how we tried to address these with practical experiments. We will look at how the results of experiments conducted in this study support our hypotheses for the conditions that we have studied, and focus on lessons we have learned through the course of this research. In the following sections, we mention the contributions this work provides to the field of VE research and 3DUI migration. We also identify areas for future research stemming from this work outside of the scope of this study.

7.1 – Summary

3DIT migration is an issue that can be difficult if not handled intelligently during the design process. In this thesis we have looked at VE device specificity as a solution to the problem of 3DIT migration. This work is part of a larger initiative researching the uses of design specificity in 3DUIs. Our work has been aimed to show benefits of using device specificity in practice through the study of a couple common cases.

The displays and input devices we used in our research were an HMD, a CAVE, a wand, and Pinch Gloves. These VE devices were chosen for their widely varying characteristics. Although they are not nearly representative of all possible VE device combinations, they are very commonly used in virtual environment systems.

The 3DITs chosen for this study were the HOMER, Go-Go, and Voodoo Dolls manipulation techniques. We chose manipulation techniques for the study because the manipulation task is salient in common VE tasks. Specifically, these 3DITs were chosen because they each represented a different manipulation metaphor.

7.1.1 - The migratability of 3DITs

In chapter two we discussed several works that delved into the effects of displays and devices on the usability of a 3DUI. These works positively showed that certain characteristics of displays and devices have large effects on the usability of VE applications. We identified the many characteristics of displays and devices in chapter three and proposed a set of displays and devices to further investigate in this study. These were the HMD, CAVE, wand, and Pinch Gloves. They were chosen both because they are common VE devices and because they are not similar devices with regards to their characteristics. These are discussed in section 3.1.2 for the displays and section 3.2.2 for the devices.

Our investigation, led by our research questions one and two in section 1.4.1, sought to show that displays and devices have negative effects on the usability of 3DITs. In section 1.4.2 we indicated that we believed that 3DITs are developed with a primary VE system and that naïve migration to secondary VE systems (or VE systems with differing display and device characteristics) would be detrimental to usability. In chapter four we evaluated naïve migrations of different 3DITs from their primary VE systems to a few secondary systems. The formative evaluation was qualitative in nature to help us find the causes of any migratability issues found.

The evaluation presented in chapter four demonstrated that the characteristics of the displays and devices used in this study did have some effects on the usability of the 3DITs. However, these effects do not completely break the usability of the 3DIT altogether. None of the 3DITs became unusable and critical incidents did not occur as frequently as expected. However, we did identify some issues with the naïve migrations that resulted in degraded usability of the techniques.

In our first evaluation we found that the HOMER and Go-Go manipulation techniques suffered in user satisfaction and performance when used on a VE system involving the CAVE and Pinch Gloves. The problem was encountered as a result of the introduction of new tasks and poor device mappings. Using the CAVE limits the field of regard and as a result necessitates a means of manually rotating the view. This manual view rotation task in turn had to somehow be incorporated into migrated 3DITs. Although the migration to the CAVE with a wand device was successful, the migration with a set of Pinch Gloves suffered from unintuitive and problematic device mappings. We found that the form-factor of the device made simultaneous pinches on the same hand difficult to perform. This in combination with a device mapping that required users to perform these difficult pinches resulted in the performance and subjective rating drops that we have seen in this experiment.

From our formative evaluation of the HOMER technique we saw that a limitation of the field of regard of the display results in the need for the HOMER 3DIT to support added functionality. In turn, the 3DIT suffered in terms of usability. From this finding we put forth the following guideline for the migration design process:

Be cautious of situations that force added functionality.

This caution should be used in practice by developers during the migration process. If a 3DIT must be used on a secondary VE system, VE developers familiar with the characteristics of devices and their effects should identify whether or not the 3DIT must support functionality that it does not already provide. If this is the case, device-specific redesign should be well thought out and evaluated for usability before deployment of the technique. An example of how this principle can be applied to another migration can be seen in the device-specific redesign of the virtual flying navigation technique on bend-sensing data gloves. Flying is typically performed with a wand device that has a joystick that controls flight velocity. If flying is migrated to a VE system that uses data gloves, the flight velocity could naively be mapped to a finger bend to achieve the same velocity control effect as a joystick. However, when the joystick is released by the user it returns to its origin, allowing the user to perform other actions (such as pressing a button) without unintentionally flying. The joystick is a passive device. When using data gloves the user cannot simply stop providing bend data because it is an active device. As a result, we must design added functionality into the 3DIT to enable the user to stop providing velocity control data so that their hand can be free to perform other functions or gestures. It is not easy to design this added functionality with only one glove. However, with two gloves we may be able to design a gesture for the non-dominant hand to signify when to use or ignore velocity control data provided by the dominant hand.

Also from our formative evaluation of the HOMER technique we observed the negative effects of poor device mappings on the Pinch Gloves. We saw that the same tasks that were easily supported by the wand device were flawed when mapped to the Pinch Gloves even though both devices use discrete input and 6DOF tracking was provided. This was because users found that holding multiple pinches on the same hand (a pinch to hold the object and a pinch to rotate the room) was difficult to perform. From this finding we present the following guideline for migratability:

Devices that provide similar input do not necessarily provide equal usability for the same tasks.

This guideline regards the migratability of devices and should practically be used as a caution. Developers of VEs have developing tools at their disposal to interface with different devices in an abstract way. This abstraction makes devices such as tracked wands and tracked pens look the same from the developer's point of view. The only input the developer needs to deal with in these two cases is 6DOFs and some discrete input values. In our work, we used a wand and a Pinch Glove to perform manual view rotation tasks. This task only required the discrete input of these two devices. From a developer's perspective there is no change apparent. However, although a specific task may use the same kind of input on two different devices, other factors of the device could still affect usability in different ways. In our work, the form-factor of the device led to unintuitive mappings and difficult gestures. As a further example of this point, using a ray-casting 3DIT a user can perform selection tasks with either a tracked wand or a tracked pen. Both of these devices provide six DOFs and both of them provide a button to perform the selection. However, these two devices differ in way they are grasped. It may be easier to use the wand to point rather than pen because the grasp used for the wand is much like that used on a laser pointer.

Also in our first evaluation we found that bulky and obtrusive input devices affected the usability of the Voodoo Dolls 3DIT. The Voodoo Dolls 3DIT was originally designed for use with a set of Pinch Gloves, which allow users to naturally use their hands for bimanual manipulation tasks. We found in our evaluation that this assumption in design led to usability issues when the technique was migrated to a wand and tracked mouse device set. With the introduction of these devices we saw that users became frustrated and less comfortable with the technique. They continually bumped and crossed their input devices during tasks. Through interpretation of our qualitative results we pinned the problem to being caused by the differences in form-factor as well as affordances of the devices. The bulkiness and obtrusiveness of the wand/tracked mouse device set did not facilitate bimanual interaction because users were not as aware of the boundaries of the devices. The device boundaries of the Pinch Gloves were actually the user's own hands which have strong kinesthetic senses. Additionally, the Pinch Gloves afforded a pinching metaphor that encouraged users to pass their manipulation context between their hands more often. As a result, crossing devices wasn't an issue.

Our formative evaluation of the Voodoo Dolls technique revealed that users frequently bumped and crossed their input devices when using bulky or obtrusive devices such as the wand and tracked mouse. Users were not as aware of the boundaries of their input devices. From our experience we can assert the following guideline:

The migratability of 3DITs that use bimanual interaction may be negatively impacted by device form-factor.

Another example of this phenomenon can be illustrated through the migration of a specialized 3DIT developed by Hinckley et al. [26] to a pen-and-tablet device. The 3DIT is used for viewing volumetric data of the head using bimanual interaction and employs a tracked doll's head in the non-dominant hand and a tracked "cutting plane" in the dominant hand. If the technique was migrated to a pen-and-tablet device, usability would degrade as a result of large changes in form-factor. The device in the non-dominant hand would change from a small, round, precision grasp object to a larger, power grasp object with far protruding edges. Users would not be as aware of the protruding edges of the tablet and would clumsily bump their devices in attempts to rotate the pen around tablet. Additionally, it would become difficult to make precision movements.

We cannot say that naïve migration guarantees degraded usability on all secondary VE systems. However, we did find that usability degraded on for two of our 3DITs on some of the displays and devices we studied. These issues found in our formative evaluation were specific to certain points in our migration space and were therefore worthy of further study regarding device specificity. Next we used device specificity to fix our usability issues and demonstrated its usefulness with a comparative evaluation.

7.1.2 – Applying device specificity to 3DIT migration

Our work posited that 3DIT migration would benefit from device specificity. In research question three we asked if display-specific or device-specific 3DITs could be designed to maintain usability on secondary VE systems and therefore increase migratability. To investigate this, we redesigned the naïve 3DIT implementations on the VE systems that were causing the usability problems. A discussion of our redesign and rationale can be read in chapter five. In the device-specific redesign we attempted to engineer out these issues in usability by identifying what characteristics of the device were causing the problems.

In chapter six we used a comparative evaluation of the redesigned 3DITs, the naïve migrations, and the original 3DIT implementations to show the effectiveness of device specificity for these situations. The comparisons we made pitted the redesigns against the naïve migrations to show that our redesign fixed the usability issues they were intended to. We also used a comparison of the redesigns against the original implementations to see what degree of migratability we had achieved through using device specificity in design. A third comparison between the original implementations and the naïve migrations would be used to reconfirm our findings in the first evaluation.

Our redesign of the HOMER technique was a success. In our experience, we are led to believe that device specificity can be a good approach to 3DIT migration. Subjective ratings as well as performance metrics indicated that the HOMER redesign was more usable than the HOMER naïve migration.

Our evaluation also showed that users preferred to use the HOMER redesign over the original HOMER implementation. We suspect the cause of this to be the way the HOMER redesign promotes lazy interaction. For example, users can perform all tasks without turning their bodies away from the front screen of the CAVE. However, performance metrics still show significant results in favor of the original HOMER implementation.

The Voodoo Dolls redesign did not net results that favored the use of device specificity. Subjective ratings and performance differences show that the redesign was not more usable than a naïve migration of the technique. However, we do show that we have successfully engineered out the original usability issues that plagued this technique. We also show that other usability issues were introduced into the technique that we believe may have caused our disappointing result.

An interesting development in this evaluation was the evidence that showed that the naïve Voodoo Dolls implementation actually outperformed the original. This was contrary to subjective evidence provided by the users. This subjective evidence was in fact consistent with the previous evaluation of the technique. We blame the discrepancy on actual device limitations that were not perceived by the users. We believe that the first usability evaluation failed to pickup on this discrepancy because it was highly based on the subjective measures of the user.

Our experiences with device specific redesign have provided us with insight into the device-specific design and migration processes. Firstly, we have learned that device specificity is an iterative design process much like that described by Gabbard et al. [47]. In chapter six we show that our initial redesign of Voodoo Dolls was relatively unsuccessful. However, through evaluation we have identified new problem areas that can be addressed to improve usability. We are confident that in section 6.7.2 we have proposed a second iteration of redesign based on the findings of the second experiment that will result in a more usable device-specific implementation of Voodoo Dolls.

Furthermore, we also see an opportunity to use this research in conjunction with research involving the creation of model-based 3DIT development systems that promote reuse on multiple VE platforms. Chapter two discusses work in the 2D GUI realm that uses a transformational, model-based approach to multiplatform GUI development. Some work is also discussed in chapter two that attempts to realize model-based approaches to 3DUI migration. Systems such as Envir3D and specification languages such as ICDL (discussed in section 2.3) can be used to perform 3DIT migration. These systems use abstract models in conjunction with concrete realizations to support the development of 3DITs on multiple platforms. Because these systems support multiple concrete realizations for the same interaction component, these systems are suitable candidates for implementing device-specific 3DITs and performing 3DIT migration.

7.2 – Contributions

Below is a list of contributions made to the study and development of migratable 3D user interfaces through our practical experiences with device-specific design. The following list is useful to the VE community at large.

- **We have demonstrated that VE developers should be aware of 3DIT migratability.** In chapter four we explored the migratability of three best-practice 3DITs over a small set of common VE systems. In our experience, 3DIT migratability was higher than we expected. In section 4.8.1 we saw that out of three secondary systems evaluated for the HOMER and Go-Go 3DITs, only one produced largely degraded usability. And in section 4.8.2 we saw that the Voodoo Dolls technique was only affected by changes in input device and not by changes in display. In our experience, however, naïve migration of 3D manipulation techniques can potentially decrease the usability of the technique. We have identified degraded usability in single-handed interaction techniques on a VE system that uses a CAVE/Pinch Gloves combination (discussed in section 4.8.1). We have also identified degraded usability in the Voodoo Dolls interaction technique when using input devices that are bulky and obtrusive and do not promote natural direct manipulation (discussed in section 4.8.1). Thus we assert that 3DIT migration is a potential problem and that VE developers should be aware of this when designing for multiple VE systems.

- **We have shown VE device specificity to be successful for 3DIT migration.** By creating a successful device-specific redesign (designed in section 5.1.3 and evaluated in section 6.7.1), we have shown that device specificity can be used to promote 3DIT migratability. In the summary of our work (section 7.1.2) we have asserted that device specificity is not only an option, but also a good solution to the problem of 3DIT migratability because it is easily accommodated by VE design tools that use a model-based or transformational approach for multiplatform VE development.

- **We have identified some characteristics of VE devices that may have impacts on the usability of interaction techniques developed with them.** In chapter three we discuss characteristics of displays and devices that have effects on the usability of 3DITs. We have given examples of how these characteristics may cause lower degrees of migratability among 3DITs. In our experience from the first usability experiment in chapter four, device form-factor and device affordances specifically affect the migration process. In addition, changes in display FOR require 3DITs to support a view rotation tasks that potentially affects usability of the 3DIT.

The following list of contributions is specific to the displays, devices, and 3DITs we have evaluated and redesigned.

- **We have created design guidelines for the 3DIT migration process.**
Through experience in the design process of 3DIT migration we have learned several lessons. In chapter four we found situations where VE developers should be wary of 3DIT migratability and make active design decisions. We have compiled these into design guidelines that can be reused by VE developers concerned with the migratability of their 3DUIs. In chapters five and six we learned that display- or device-specific design is not always successful on the first iteration and is an iterative design process.

- **We have made specific improvements on HOMER to increase the migratability of this technique.**
In chapter five we present a device-specific redesign of the original HOMER implementation. This redesign has been shown in section 6.7.1 to be quite usable in a VE system involving a CAVE and Pinch Gloves. User performance and subjective ratings back up the usability of the device-specific redesign. Developers wishing to use the HOMER technique on a CAVE/Pinch Gloves VE system can use this HOMER implementation to maintain its usability.

7.3 – Future Work

In our work, we have started the process of device-specific redesign with a single iteration of design and evaluation. This proved successful with the HOMER 3DIT. However, the Voodoo Dolls 3DIT is still not a usable redesign. More design iterations of the wand-specific Voodoo Dolls redesign need to be made to show the success of device specificity in this case.

Our work has shown that the problem of 3DIT migration is a result of the characteristics of displays and devices. We have presented some work that has looked at the effects of these characteristics on 3DIT usability. However, work can be performed that develops extensive guidelines to help VE developers make intelligent device-specific design decisions.

We have explored the use and usefulness of device specificity. We have shown that it can result in more usable 3DIT designs and that it can be useful as a solution to the problem of 3DIT migration. Our work, however, is part of a larger direction in research and design towards specificity. Already there is work that explores domain specificity [46]. There are some other largely unexplored areas in the topic of specificity. Work can be done to find the uses and usefulness of user and task specificity [38].

In this thesis we have explored the importance of the primary VE system of 3DITs. 3DIT designers leverage the display and device characteristics of the VE system on which 3DITs are developed. Although in our work we have shown that this does not make the 3DIT unusable on all VE systems, we also show 3DITs do have a set of VE systems on which they are most usable. A useful application of this knowledge would be to help VE developers choose 3DITs given the VE system for which they are developing by mapping 3DITs to the displays and devices on which they are usable. Future work in the area of 3DIT migration can focus on researching and compiling this mapping.

We have explored the affinities of three 3DITs towards a small set of VE systems through empirical evaluation. In the context of the entire space of 3DITs, displays, and devices, the mapping we can create is very limited. To more effectively study the subject of 3DIT to VE system mappings, an analytical approach may be useful. Future work can use the display and device characteristics discussed in this paper and extend the impacts these characteristics may have on various interaction components of 3DITs. From this, researchers may be able to create 3DIT to VE system mappings by analyzing the characteristics of the display and the device and identifying potential issues. Here we discuss a simple example using a 6DOF SpaceMouse. This device is not tracked and provides 6DOF control through isometric input. Previous work has identified this kind of input to not be well-suited for the direct manipulation metaphor. With this knowledge we can induct that all 3DITs that use a direct manipulation metaphor, such as Go-Go and Voodoo Dolls, suffer in usability when a SpaceMouse is used for input.

Our work with device specificity can be used in conjunction with research involving the creation of model-based 3DIT development systems that promote reuse on multiple VE platforms. Systems such as Envir3D and specification languages such as ICDL are not simply multiplatform development tools. They can be used to perform intelligent 3DIT migration. These systems use abstract models in conjunction with concrete realizations to support the development of 3DITs on multiple platforms. Because these systems support multiple concrete realizations for the same interaction component, these systems are suitable candidates for implementing device-specific 3DITs and performing 3DIT migration.

Works Cited

1. Steinwand, D., B. Davis, and N. Weeks, *GeoWall: Investigations into Low-cost Stereo Display Systems*. 2002, USGS.
2. Forsberg, A., et al. *Immersive Virtual Reality for Visualizing Flow though an Artery*. in *Proceedings of IEEE Visualization 2000*. 2000.
3. Kitamura, Y., et al. *The IllusionHole for Medical Applications*. in *IEEE Virtual Reality Conference 2007*. 2007. Charlotte, North Carolina, USA.
4. Bryson, S., *Virtual Reality in Scientific Visualization.* Communications of the ACM, 1996. **39**(5): p. 62--71.
5. Butterworth, J., et al. *3DM: A Three-Dimensional Modeler Using a Head-Mounted Display*. in *Proceedings of the 1992 ACM Symposium on Interactive 3D Graphics (I3D '92)*. 1992: ACM Press.
6. Kasik, D., et al., *Evaluating Graphics Displays for Complex 3D Models.* IEEE Computer Graphics and Applications, 2002. **22**(3): p. 56--64.
7. Liang, J. and M. Green, *JDCAD: A Highly Interactive 3D Modeling System.* Computers and Graphics, 1994. **18**(4): p. 499--506.
8. Cruz-Neira, C. and R. Lutz, *Using Immersive Virtual Environments for Certification.* IEEe Software, 1999. **16**(4): p. 26-30.
9. Bliss, J., P. Tidwell, and M. Guest, *The Effectiveness of Virtual Reality for Administering Spatial Navigation Training to Firefighters.* Presence-Teleoperators & Virtual Environments, 1997. **6**(1): p. 73--86.
10. Hodges, L.F., et al., *Virtual Environments for Treating the Fear of Heights.* IEEE Computer, 1995. **28**(7): p. 27--34.
11. Carlin, A., H. Hoffman, and S. Weghorst, *Virtual Reality and Tactile Augmentation in the Treatment of Spider Phobia: A Case Report.* Behavior Research and Therapy, 1997. **35**(2): p. 153--159.
12. Hodges, L., et al., *Virtual Vietnam: A Virtual Environment for the Treatment of Chronic Post-traumatic Stress Disorder.* CyberPsychology and Behavior, 1995. **2**(1): p. 7--14.
13. Blom, K., G. Lindhal, and C. Cruz-Neira. *Multiple Active Viewers in Projection-Based Immersive Environments*. in *Proceedings of the Seventh Annual Immersive Projection Technology Workshop*. 2002. Orlando, Florida, USA.
14. Pausch, R., et al. *Disney's Aladdin: First Steps toward Storytelling in Virtual Reality*. in *Proceedings of SIGGRAPH '96*. 1996: ACM Press.
15. Bowman, D.A., et al., *3D User Interfaces: Theory and Practice*. 2004: Addison-Wesley/Pearson Education.
16. Sherman, W.R. and A. Craig, *Understanding Virtual Reality: Interface, Application, and Design*. 2003, San Francisco: Morgan Kaufmann.
17. Ware, C., K. Arthur, and K. Booth. *Fishtank Virtual Reality*. in *Proceedings of INTERCHI '93*. 1993.
18. Blundell, B.G. and A.J. Schwarz, *Volumetric Three-Dimensional Display Systems*. 2000: John Wiley & Sons.
19. Hua, H., et al. *An Ultra-Light and Compact Design and Implementation of Head-Mounted Projective Displays*. in *Proceedings of IEEE Virtual Reality 2001*. 2001: IEEE Press.

20. McQuaide, S., et al., *Three-Dimensional Virtual Retinal Display System Using Deformable Membrane Mirror*, in *Society for Information Display (SID '02)*. 2002: Boston, MA, USA.
21. Cruz-Neira, C., D.J. Sandin, and T.A. DeFanti. *Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE*. in *Computer Graphics, Annual Conference Proceedings Series*. 1993.
22. Gaver, W., et al., *Realizing a video environment: EuroPARC's RAVE system*, in *CHI '92: Proceedings of the SIGCHI conference on Human factors in computing systems*. 1992, ACM Press. p. 27--35.
23. Iwata, H., *Ensphered Vision*, in *SIGGRAPH '99: ACM SIGGRAPH 99 Conference abstracts and applications*. 1999, ACM Press. p. 175.
24. LaViola, J., et al. *Case Studies in Building Custom Input Devices for Virtual Environment Interfaces*. in *Proceedings of the VR 2004 Workshop on Beyond Wand and Glove Interaction*. 2004.
25. Keefe, D., et al. *CavePainting: A Fully Immersive 3D Artistic Medium and Interactive Experience*. in *Proceedings of the 2001 Symposium on Interactive 3D Graphics (I3D 2001)*. 2001: ACM Press.
26. Hinckley, K., et al. *Passive Real-World Interface Props for Neurosurgical Visualization*. in *ACM Conference on Human Factors in Computing Systems*. 1994: ACM Press.
27. Kelso, J., et al. *DIVERSE: A Framework for Building Extensible and Reconfigurable Device Independent Virtual Environments*. in *IEEE Virtual Reality*. 2002.
28. Wingrave, C. and D.A. Bowman. *CHASM: Bridging Description and Implementation of 3D Interfaces*. in *New Directions in 3D User Interfaces Workshop, IEEE Virtual Reality 2005*. 2005.
29. Bierbaum, A., et al., *VR Juggler: A Virtual Platform for Virtual Reality Application Development*, in *VR '01: Proceedings of the Virtual Reality 2001 Conference (VR'01)*. 2001, IEEE Computer Society. p. 89.
30. Vanderdonckt, J., et al., *Model-based design, generation, and evaluation of virtual user interfaces*, in *Web3D '04: Proceedings of the ninth international conference on 3D Web technology*. 2004, ACM Press. p. 51--60.
31. Figueroa, P., M. Green, and H.J. Hoover, *InTml: a description language for VR applications*, in *Web3D '02: Proceeding of the seventh international conference on 3D Web technology*. 2002, ACM Press. p. 53-58.
32. Hinckley, K., et al., *Usability analysis of 3D rotation techniques*, in *UIST '97: Proceedings of the 10th annual ACM symposium on User interface software and technology*. 1997, ACM Press. p. 1-10.
33. Kjeldskov, J. *Interaction: Full and Partial Immersive Virtual Reality Displays.* . in *IRIS24*. 2001. Bergen, Norway: Bergen, University of Bergen.
34. Manek, D., *Effects of Visual Displays on 3D Interaction in Virtual Environments*, in *Dept. of Computer Science*. 2004, Virginia Tech: Blacksburg, VA.
35. Steed, A. and C. Parker, *Evaluating Effectiveness of Interaction Techniques across Immersive Virtual Environmental Systems*, in *Presence: Teleoperators & Virtual Environments*. 2005. p. 511-527.

36. Bowman, D.A., et al. *Empirical Comparison of Human Behavior and Performance with Different Display Devices for Virtual Environments*. in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. 2002.
37. Zhai, S. and P. Milgram. *Human performance evaluation of manipulation schemes in virtual environments*. 1993.
38. Bowman, D.A., et al., *New Directions in 3D User Interfaces.* International Journal of Virtual Reality, 2006. **5**(2): p. 3--14.
39. Bowman, D.A., et al. *Using Pinch Gloves for both Natural and Abstract Interaction Techniques in Virtual Environments*. in *HCI International*. 2001.
40. Bowman, D.A. and L.F. Hodges. *An evaluation of techniques for grabbing and manipulating remote objects in immersive virtual environments*. in *Proceedings of the 1997 symposium on Interactive 3D graphics*. 1997. Providence, Rhode Island, United States: ACM Press.
41. Poupyrev, I., et al. *The go-go interaction technique: non-linear mapping for direct manipulation in VR*. in *Proceedings of the 9th annual ACM symposium on User interface software and technology*. 1996. Seattle, Washington, United States: ACM Press.
42. Stoakley, R., M.J. Conway, and R. Pausch. *Virtual reality on a WIM: interactive worlds in miniature*. in *Proceedings of the SIGCHI conference on Human factors in computing systems*. 1995. Denver, Colorado, United States: ACM Press/Addison-Wesley Publishing Co.
43. Nielsen, J., *Usability Engineering*. 1994, San Francisco: Morgan Kaufmann.
44. Poupyrev, I., et al. *A Framework and Testbed for Studying Manipulation Techniques for Immersive VR*. in *Proceedings ACM VRST*. 1997.
45. Chen, J. and D.A. Bowman. *Effectiveness of Cloning Techniques for Architectural Virtual Environments*. in *Proceedings of IEEE Virtual Reality*. 2006.
46. Chen, J. and D.A. Bowman. *Effectiveness of Cloning Techniques for Architectural Virtual Environments*. in *IEEE Virtual Reality*. 2006.
47. Gabbard, J.L., D. Hix, and J.E.S. II, *User-Centered Design and Evaluation of Virtual Environments.* Computer Graphics and Applications, IEEE, 1999. **19**(6): p. 51--59.
48. Arthur, K., *Effects of field of view on task performance with head-mounted displays*, in *CHI '96: Conference companion on Human factors in computing systems*. 1996, ACM Press. p. 29--30.
49. Lindt, I. *Exchangeability of 3D Interaction Techniques*. in *IEEE Workshop on New Directions in 3D User Interfaces*. 2005.
50. Figueroa, P., R. Dachselt, and I. Lindt. *A Conceptual Model and Specification Language for Mixed Reality Interface Components*. in *Proceedings of the Workshop "Specification of Mixed Reality User Interfaces: Approaches, Languages, Standardization" at IEEE Virtual Reality 2006*. 2006. Alexandria, Virginia, USA.
51. LaViola, J.J., et al., *Hands-free multi-scale navigation in virtual environments*, in *SI3D '01: Proceedings of the 2001 symposium on Interactive 3D graphics*. 2001, ACM Press. p. 9--15.

52. Sutcliffe, A., *On the effective use and reuse of HCI knowledge.* ACM Trans. Comput.-Hum. Interact., 2000. **7**(2): p. 197--221.
53. McCrickard, D.S., et al. *Use and Reuse in Information and Interaction Design*. in *Proceedings of the 11th International Conference on Human-Computer Interaction (HCII '05)*. 2005. Las Vegas NV.
54. Florins, M. and J. Vanderdonckt, *Graceful degradation of user interfaces as a design method for multiplatform systems*, in *IUI '04: Proceedings of the 9th international conference on Intelligent user interfaces*. 2004, ACM Press. p. 140--147.
55. Phanouriou, C., ed. *Uiml: a device-independent user interface markup language*. 2000.
56. Puerta, A. and J. Eisenstein, *XIML: a common representation for interaction data*, in *IUI '02: Proceedings of the 7th international conference on Intelligent user interfaces*. 2002, ACM Press. p. 214--215.
57. Limbourg, Q. and J. Vanderdonckt. *UsiXML: A User Interface Description Language Supporting Multiple Levels of Independence*. in *Engineering Advanced Web Applications*. 2004: Rinton Press, Paramus.
58. Bowman, D.A., J.L. Gabbard, and D. Hix, *A Survey of Usability Evaluation in Virtual Environments: Classification and Comparison of Methods.* Presence-Teleoperators & Virtual Environments, 2002. **11**(4): p. 404-424.
59. Slater, M., *A Note on Presence Terminology.* Presence-Connect, 2003. **3**.
60. McMahan, R.P., et al., *Separating the effects of level of immersion and 3D interaction techniques*, in *VRST '06: Proceedings of the ACM symposium on Virtual reality software and technology*. 2006, ACM Press. p. 108--111.
61. Bowman, D.A. and D. Raja, *A Method for Quantifying the Benefits of Immersion Using the CAVE.* Presence-Connect, 2004. **4**(2).
62. Sutcliffe, A., et al., *Investigating interaction in CAVE virtual environments.* ACM Trans. Comput.-Hum. Interact., 2006. **13**(2): p. 235--267.
63. Lantz, E., *The future of virtual reality: head mounted displays versus spatially immersive displays (panel)*, in *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*. 1996, ACM Press. p. 485--486.
64. Zhai, S., P. Milgram, and W. Buxton. *The Influence of Muscle Groups on Performance of Multiple Degree-of-Freedom Input*. in *ACM Conference on Human Factors in Computing Systems (CHI '96)*. 1996: ACM Press.
65. Zhai, S., P. Milgram, and A. Rastogi, *Anisotropic Human Performance in Six Degree-of-Freedom Tracking: An Evaluation of Three-Dimensional Display and Control Interfaces.* 1998 ACM Conference on Human Factors in Computing Systems (CHI '98), 1997. **27**(4): p. 518--528.
66. Bowman, D.A., et al., *Novel Uses of Pinch Gloves for Virtual Environment Interaction Techniques.* Virtual Reality, 2002. **6**(3): p. 122-129.
67. Mapes, D. and J. Moshell, *A Two-Handed Interface for Object Manipulation in Virtual Environments.* Presence-Teleoperators & Virtual Environments, 1995. **4**(4): p. 403--416.

68. Poupyrev, I., et al., *Egocentric Object Manipulation in Virtual Environments: Empirical Evaluation of Interaction Techniques.* Computer Graphics Forum, 1998. **17**(3): p. 41-52.
69. Mine, M.R., ed. *Virtual Environment Interaction Techniques*. 1995, Technical Report. UMI Order Number: TR95-018., University of North Carolina at Chapel Hill.
70. Pierce, J.S., et al., *Image plane interaction techniques in 3D immersive environments*, in *SI3D '97: Proceedings of the 1997 symposium on Interactive 3D graphics*. 1997, ACM Press. p. 39--ff.
71. Pierce, J.S., B. Stearns, and R. Pausch. *Voodoo Dolls: Seamless Interaction at Multiple Scales in Virtual Environments.* . in *Proceeding of the 1999 Symposium on Interactive 3D Graphics*. 1999.
72. Pierce, J.S. and R. Pausch. *Comparing voodoo dolls and HOMER: exploring the importance of feedback in virtual environments*. in *Proceedings of the SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves*. 2002. Minneapolis, Minnesota, USA: ACM Press.
73. Mine, M.R., F.P. Brooks, and C.H. Sequin, *Moving objects in space: exploiting proprioception in virtual-environment interaction*, in *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 1997, ACM Press/Addison-Wesley Publishing Co. p. 19--26.
74. Kessler, G.D., D.A. Bowman, and L.F. Hodges, *The Simple Virtual Environment Library: An Extensible Framework for Building VE Applications.* Presence-Teleoperators & Virtual Environments, 2000. **9**(2): p. 187--208.
75. Arsenault, L.E. and J. Kelso. *The DIVERSE Toolkit: A Toolkit for Distributed Simulations and Peripheral Device Services*. in *IEEE Virtual Reality*. 2002.
76. Bowman, D.A., et al. *Empirical Comparison of Human Behavior and Performance with Different Display Devices for Virtual Environments*. in *Human Factors and Ergonomics Society Annual Meeting*. 2002.
77. Hix, D. and H.R. Hartson, *Developing User Interfaces: Ensuring usability through product and process*. 1993, New York: John Wiley and Sons.
78. Boeck, J.D. and K. Coninx. *Haptic Camera Manipulation: Extending the "Camera In Hand Metaphor"*. in *Proceedings of Eurohaptics*. 2002.

Appendix A

A.1 – Pre-Questionnaire for Experiment 1

User Questionnaire

Please help us to categorize our user population by completing the following items.

What is your age? ________

Gender? (circle one)

Male Female

Are you: (circle one)

Right-handed Left-handed

Do you wear glasses or contacts? (circle one)

Yes No

Is your eyesight near-perfect with correction or without correction? (circle one)

Yes No

Rate your familiarity with computers: (circle one)

•-----------------------•-----------------------•-----------------------•-----------------------•

not at all familiar | not very familiar | somewhat familiar | fairly familiar | very familiar

Rate your familiarity with virtual environments: (circle one)

•-----------------------•-----------------------•-----------------------•-----------------------•

not at all familiar | not very familiar | somewhat familiar | fairly familiar | very familiar

A.2 – Sample Post-Questionnaire for Experiment 2

HOMER - User Evaluation Form

HMD and Wand

Please rate the ease of use of the HOMER technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard | hard | neither hard nor easy | easy | very easy

Please rate the amount of precision you feel you could achieve with the HOMER technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no precision | little precision | adequate precision | good precision | perfect precision

Please rate your overall frustration level with this technique using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no frustration | little frustration | moderate frustration | some frustration | complete frustration

Please indicate how physically comfortable you were with the HOMER technique using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very uncomfortable | uncomfortable | neutral | comfortable | very comfortable

Please indicate how easy you found it was to see the virtual world and the effects of your movement in the world with the HOMER technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard | hard | neither hard nor easy | easy | very easy

Please indicate how easy you found it was to move your virtual hand with the HOMER technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard | hard | neither hard nor easy | easy | very easy

HOMER - User Evaluation Form
HMD and Pinch Gloves

Please rate the ease of use of the HOMER technique when using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Please rate the amount of precision you feel you could achieve with the HOMER technique when using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no precision little precision adequate precision good precision perfect precision

Please rate your overall frustration level with this technique using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no frustration little frustration moderate frustration some frustration complete frustration

Please indicate how physically comfortable you were with the HOMER technique using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very uncomfortable uncomfortable neutral comfortable very comfortable

Please indicate how easy you found it was to see the virtual world and the effects of your movement in the world with the HOMER technique when using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Please indicate how easy you found it was to move your virtual hand with the HOMER technique when using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

HOMER - User Evaluation Form
CAVE and Wand

Please rate the ease of use of the HOMER technique when using the

CAVE and wand: (circle one)
•----------------------•----------------------•----------------------•----------------------•
very hard   hard   neither hard nor easy   easy   very easy

Please rate the amount of precision you feel you could achieve with the HOMER technique when using the

CAVE and wand: (circle one)
•----------------------•----------------------•----------------------•----------------------•
no precision   little precision   adequate precision   good precision   perfect precision

Please rate your overall frustration level with this technique using the

CAVE and wand: (circle one)
•----------------------•----------------------•----------------------•----------------------•
no frustration   little frustration   moderate frustration   some frustration   complete frustration

Please indicate how physically comfortable you were with the HOMER technique using the

CAVE and wand: (circle one)
•----------------------•----------------------•----------------------•----------------------•
very uncomfortable   uncomfortable   neutral   comfortable   very comfortable

Please indicate how easy you found it was to see the virtual world and the effects of your movement in the world with the HOMER technique when using the

CAVE and wand: (circle one)
•----------------------•----------------------•----------------------•----------------------•
very hard   hard   neither hard nor easy   easy   very easy

Please indicate how easy you found it was to move your virtual hand with the HOMER technique when using the

CAVE and wand: (circle one)
•----------------------•----------------------•----------------------•----------------------•
very hard   hard   neither hard nor easy   easy   very easy

HOMER - User Evaluation Form
Cave and Pinch Gloves

Please rate the ease of use of the HOMER technique when using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Please rate the amount of precision you feel you could achieve with the HOMER technique when using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no precision little precision adequate precision good precision perfect precision

Please rate your overall frustration level with this technique using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no frustration little frustration moderate frustration some frustration complete frustration

Please indicate how physically comfortable you were with the HOMER technique using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very uncomfortable uncomfortable neutral comfortable very comfortable

Please indicate how easy you found it was to see the virtual world and the effects of your movement in the world with the HOMER technique when using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Please indicate how easy you found it was to move your virtual hand with the HOMER technique when using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Which of these device pairs do you prefer for performing object manipulation tasks with the HOMER technique? Please rank the devices from 1-4, with 1 indicating the most preferred device pair, and 4 indicating the least preferred device pair.

- HMD-wand ______
- HMD-Pinch Gloves ______
- CAVE-wand ______
- CAVE-Pinch Gloves ______

Other Comments

Please write any comments you have on using the HOMER technique in the HMD and with the wand here:

____

____

____

Please write any comments you have on using the HOMER technique in the HMD and with Pinch Gloves here:

____

____

____

Please write any comments you have on using the HOMER technique in the CAVE and with the wand here:

____

____

____

Please write any comments you have on using the HOMER technique in the CAVE and with the Pinch Gloves here:

____

____

____

A.3 – First Usability Evaluation Subjective Results

HOMER	HMD-Wand	HMD-PinchGloves	CAVE-Wand	CAVE-PinchGloves
Ease of Use	4.5	3.5	4	3.5
Precision	4	3	4	2.5
Frustration	2	1.5	3.5	1.5
Physical Comfort	4	3.5	3	3.5
Vision	4.5	4	4	4
Interaction	4	3.5	5	3.5

Go-Go	HMD-Wand	HMD-PinchGloves	CAVE-Wand	CAVE-PinchGloves
Ease of Use	4	4	3	2.5
Precision	4	4	4	2
Frustration	2	2.5	1.5	1
Physical Comfort	2.5	2	2.5	1.5
Vision	4.5	4	4	2
Interaction	3.5	3.5	2.5	2

Voodoo Dolls	HMD-Wand	CAVE-Wand	HMD-PinchGloves	CAVE-PinchGloves
Ease of Use	3	3	4.5	3.5
Precision	3.5	3	3.5	3.5
Frustration	1.5	2	3	3.5
Physical Comfort	3	4	4.5	5
Vision	4.5	4.5	4	4.5
Interaction	4	4	4	5

	Ranking			
	First	Second	Third	Fourth
Participant 1	HMD/Wand	CAVE/Wand	HMD/Pinch Gloves	CAVE/Pinch Gloves
Participant 2	HMD/Pinch Gloves	HMD/Wand	CAVE/Wand	CAVE/Pinch Gloves
Participant 3	CAVE/Wand	HMD/Pinch Gloves	HMD/Wand	CAVE/Pinch Gloves
Participant 4	HMD/Wand	HMD/Pinch Gloves	CAVE/Wand	CAVE/Pinch Gloves
Participant 5	CAVE/Pinch Gloves	HMD/Pinch Gloves	CAVE/Wand	HMD/Wand
Participant 6	HMD/Pinch Gloves	CAVE/Pinch Gloves	HMD/Wand	CAVE/Wand

Appendix B

B.1 – Pre-Questionnaire for Experiment 2

User Questionnaire

Please help us to categorize our user population by completing the following items.

What is your age? ______

Gender? (circle one)

Male Female

Are you: (circle one)

Right-handed Left-handed

Do you wear glasses or contacts? (circle one)

Yes No

Is your eyesight near-perfect with or without correction? (circle one)

Yes No

Rate your familiarity with computers: (circle one)

•-----------------------•-----------------------•-----------------------•-----------------------•

not at all familiar not very familiar somewhat familiar fairly familiar very familiar

Rate your familiarity with virtual environments: (circle one)

•-----------------------•-----------------------•-----------------------•-----------------------•

not at all familiar not very familiar somewhat familiar fairly familiar very familiar

B.2 – HOMER Post-Questionnaire for Experiment 2

HOMER - User Evaluation Form

HMD and Wand

Please rate the ease of use of the HOMER technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Please rate the amount of precision you feel you could achieve with the HOMER technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no precision little precision adequate precision good precision perfect precision

Please rate your overall frustration level with this technique using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no frustration little frustration moderate frustration some frustration complete frustration

Please indicate how physically comfortable you were with the HOMER technique using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very uncomfortable uncomfortable neutral comfortable very comfortable

Please indicate how easy you found it was to see the virtual world and the effects of your movement in the world with the HOMER technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Please indicate how easy you found it was to move your virtual hand with the HOMER technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

HOMER - User Evaluation Form

CAVE and Single Pinch Glove

Please rate the ease of use of the HOMER technique when using the

CAVE and Single Pinch Glove: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard   hard   neither hard nor easy   easy   very easy

Please rate the amount of precision you feel you could achieve with the HOMER technique when using the

CAVE and Single Pinch Glove: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no precision   little precision   adequate precision   good precision   perfect precision

Please rate your overall frustration level with this technique using the

CAVE and Single Pinch Glove: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no frustration   little frustration   moderate frustration   some frustration   complete frustration

Please indicate how physically comfortable you were with the HOMER technique using the

CAVE and Single Pinch Glove: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very uncomfortable   uncomfortable   neutral   comfortable   very comfortable

Please indicate how easy you found it was to see the virtual world and the effects of your movement in the world with the HOMER technique when using the

CAVE and Single Pinch Glove: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard   hard   neither hard nor easy   easy   very easy

Please indicate how easy you found it was to move your virtual hand with the HOMER technique when using the

CAVE and Single Pinch Glove: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard   hard   neither hard nor easy   easy   very easy

HOMER - User Evaluation Form
Cave and Both Pinch Gloves

Please rate the ease of use of the HOMER technique when using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard | hard | neither hard nor easy | easy | very easy

Please rate the amount of precision you feel you could achieve with the HOMER technique when using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no precision | little precision | adequate precision | good precision | perfect precision

Please rate your overall frustration level with this technique using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no frustration | little frustration | moderate frustration | some frustration | complete frustration

Please indicate how physically comfortable you were with the HOMER technique using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very uncomfortable | uncomfortable | neutral | comfortable | very comfortable

Please indicate how easy you found it was to see the virtual world and the effects of your movement in the world with the HOMER technique when using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard | hard | neither hard nor easy | easy | very easy

Please indicate how easy you found it was to move your virtual hand with the HOMER technique when using the

CAVE and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard | hard | neither hard nor easy | easy | very easy

Rank the device pairs that you prefer given the corresponding criteria. Please rank the devices from 1-3, with 1 indicating the most preferred device pair, and 3 indicating the least preferred device pair.

Rank the following for Precision:

- HMD-wand ______
- CAVE-Single Glove ______
- CAVE-Both Gloves ______

Rank the following for Frustration:

- HMD-wand ______
- CAVE-Single Glove ______
- CAVE-Both Gloves ______

Rank the following for Physical Comfort:

- HMD-wand ______
- CAVE-Single Glove ______
- CAVE-Both Gloves ______

Rank the following for ease of use:

- HMD-wand ______
- CAVE-Single Glove ______
- CAVE-Both Gloves ______

Rank the following for overall personal preference:

- HMD-wand ______
- CAVE-Single Glove ______
- CAVE-Both Gloves ______

Other Comments

Please write any comments you have on using the HOMER technique in the HMD and with the wand here:

Please write any comments you have on using the HOMER technique in the CAVE and with the single glove here:

Please write any comments you have on using the HOMER technique in the CAVE and with both Pinch Gloves here:

B.3 – Voodoo Dolls Post-Questionnaire for Experiment 2

Voodoo Dolls - User Evaluation Form

Basic Wand (Naïve Flavor)

Please rate the ease of use of the Voodoo Dolls technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard	hard	neither hard nor easy	easy	very easy

Please rate the amount of precision you feel you could achieve with the Voodoo Dolls technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no precision	little precision	adequate precision	good precision	perfect precision

Please rate your overall frustration level with this technique using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no frustration	little frustration	moderate frustration	some frustration	complete frustration

Please indicate how physically comfortable you were with the Voodoo Dolls technique using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very uncomfortable	uncomfortable	neutral	comfortable	very comfortable

Please indicate how easy you found it was to see the virtual world and the effects of your movement in the world with the Voodoo Dolls technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard	hard	neither hard nor easy	easy	very easy

Please indicate how easy you found it was to move your virtual hand with the Voodoo Dolls technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard	hard	neither hard nor easy	easy	very easy

Voodoo Dolls - User Evaluation Form
Enhanced Wand (Modified Flavor)

Please rate the ease of use of the Voodoo Dolls technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Please rate the amount of precision you feel you could achieve with the Voodoo Dolls technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no precision little precision adequate precision good precision perfect precision

Please rate your overall frustration level with this technique using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no frustration little frustration moderate frustration some frustration complete frustration

Please indicate how physically comfortable you were with the Voodoo Dolls technique using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very uncomfortable uncomfortable neutral comfortable very comfortable

Please indicate how easy you found it was to see the virtual world and the effects of your movement in the world with the Voodoo Dolls technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Please indicate how easy you found it was to move your virtual hand with the Voodoo Dolls technique when using the

HMD and wand: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Voodoo Dolls - User Evaluation Form
HMD and Pinch Gloves (Native Flavor)

Please rate the ease of use of the Voodoo Dolls technique when using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Please rate the amount of precision you feel you could achieve with the Voodoo Dolls technique when using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no precision little precision adequate precision good precision perfect precision

Please rate your overall frustration level with this technique using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

no frustration little frustration moderate frustration some frustration complete frustration

Please indicate how physically comfortable you were with the Voodoo Dolls technique using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very uncomfortable uncomfortable neutral comfortable very comfortable

Please indicate how easy you found it was to see the virtual world and the effects of your movement in the world with the Voodoo Dolls technique when using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Please indicate how easy you found it was to move your virtual hand with the Voodoo Dolls technique when using the

HMD and Pinch Gloves: (circle one)

•----------------------•----------------------•----------------------•----------------------•

very hard hard neither hard nor easy easy very easy

Rank the Voodoo Dolls variation that you prefer given the corresponding criteria. Please rank the devices from 1-3, with 1 indicating the most preferred variation, and 3 indicating the least preferred variation.

Rank the following for Precision:

- Basic Wand ______
- Enhanced Wand ______
- HMD-Pinch Gloves ______

Rank the following for Frustration:

- Basic Wand ______
- Enhanced Wand ______
- HMD-Pinch Gloves ______

Rank the following for Physical Comfort:

- Basic Wand ______
- Enhanced Wand ______
- HMD-Pinch Gloves ______

Rank the following for ease of use:

- Basic Wand ______
- Enhanced Wand ______
- HMD-Pinch Gloves ______

Rank the following for overall personal preference:

- Basic Wand ______
- Enhanced Wand ______
- HMD-Pinch Gloves ______

Other Comments

Please write any comments you have on using the Basic Wand Voodoo Dolls technique in the HMD and with the wand here:

Please write any comments you have on using the Enhanced Wand Voodoo Dolls technique in the HMD and with the wand here:

Please write any comments you have on using the Voodoo Dolls technique in the HMD and with Pinch Gloves here:

B.4 – HOMER Tasks and Observations

Table 8: The tasks used for the HOMER evaluation and their associated task types.

Task #	Tasks	Task Type
1	Move Brown Teddy to front table left	Control
2	Move Green Teddy to front table right	Control
3	Move Pink Teddy to back left table	To Back
4	Move Red Teddy to back right table	To Back
5	Move Cyan Teddy to back left table	To Back
6	Move Yellow Teddy to back right table	To Back
7	Move Gray Teddy to front table left	To Front
8	Move Purple Teddy to front table right	To Front
9	Move Gray Teddy to back right table	At Back
10	Move Purple Teddy to back left table	At Back
11	Move Blue Teddy to back right table	To Back
12	Move Blue Teddy to back left table	To Back

Table 9: A table used to record specific observations during the HOMER evaluation.

Usability Issue (In order of importance)	# of Occurrences	Severity (1 to 5)
Avoiding Edges of Screen		
Using Difficult Finger Positioning		
Avoiding Difficult Finger Positioning		
Tracker Loss		
Cable Entanglement		
Avoiding Cables		
Heisenberg Effect		

B.5 – Voodoo Dolls Task Evaluation Form

Table 10: The tasks used for the Voodoo Dolls evaluation and their associated task types.

Task #	Tasks	Task Type
1	Move front right object to back right side of table	Right Side
2	Move back right object to front right side of table	Right Side
3	Move front left object to back right side of table	Mixed
4	Move back left object to front right side of table	Mixed
5	Move front right object to back left side of table	Mixed
6	Move back right object to front left side of table	Mixed
7	Move front left object to back left side of table	Left
8	Move back left object to front left side of table	Left
9	Move front right object to center of table	Center
10	Move back right object to center of table	Center
11	Move front left object to center of table	Center
12	Move back left object to center of table	Center

Table 11: A table used to record specific observations during the Voodoo Dolls evaluation.

Usability Issue (In order of importance)	# of Occurrences	Severity (1 to 5)
Bumping Hands Together		
Hands Bumping HMD		
Hand Crossover		
Preemptive Repositioning		
Post Hoc Repositioning		
Fine Repositioning		
Difficult Hand Positions		
Tracker Blockage		
Cable Entanglement		
Heisenberg Effect		

B.6 – HOMER Performance Results

Anova: Two-Factor With Replication							
SUMMARY	Task1	Task2	Task3	Task4	Task5	Task6	
Redesign							
Count	12	12	12	12	12	12	
Sum	117.765	89.76	255.8743	234.4255	155.4933	144.495	
Average	9.81375	7.48	21.32286	19.53545	12.95778	12.04125	
Variance	33.45429	6.320873	16.25509	22.42406	8.897487	22.34964	
	Task7	Task8	Task9	Task10	Task11	Task12	Total
Count	12	12	12	12	12	12	144
Sum	256.7067	192.264	279.72	323.328	227.92	158.7867	2436.538
Average	21.39222	16.022	23.31	26.944	18.99333	13.23222	16.92041
Variance	73.54636	12.70778	26.65824	93.04657	31.42331	7.675778	59.62691
SUMMARY	Task1	Task2	Task3	Task4	Task5	Task6	
Naïve							
Count	12	12	12	12	12	12	
Sum	157.3418	101.82	326.4	300.825	159.525	311.4	
Average	13.11182	8.485	27.2	25.06875	13.29375	25.95	
Variance	82.69881	4.117714	177.2057	150.8123	5.47429	67.922	
	Task7	Task8	Task9	Task10	Task11	Task12	Total
Count	12	12	12	12	12	12	144
Sum	265.845	278.83	294.9733	295.728	322.5943	182.61	2997.892
Average	22.15375	23.23583	24.58111	24.644	26.88286	15.2175	20.8187
Variance	49.99711	85.73455	68.43852	128.1849	88.86479	24.04114	110.2252
SUMMARY	Task1	Task2	Task3	Task4	Task5	Task6	
Original							
Count	12	12	12	12	12	12	
Sum	93.69	92.97	279.28	194.6	114.2509	182.62	
Average	7.8075	7.7475	23.27333	16.21667	9.520909	15.21833	
Variance	17.9062	55.48393	100.0073	55.58361	12.25061	104.5991	
	Task7	Task8	Task9	Task10	Task11	Task12	Total
Count	12	12	12	12	12	12	144
Sum	195.26	92.83	174.72	165.33	162.6	154.4	1902.551
Average	16.27167	7.735833	14.56	13.7775	13.55	12.86667	13.21216
Variance	100.3846	2.373208	24.02378	24.8376	34.12173	63.24304	64.94639
SUMMARY	Task1	Task2	Task3	Task4	Task5	Task6	

Total							
Count	36	36	36	36	36	36	
Sum	368.7968	284.55	861.5543	729.8505	429.2692	638.515	
Average	10.24436	7.904167	23.93206	20.27362	11.92415	17.73653	
Variance	47.05156	20.90428	98.37725	85.62813	11.35667	97.66975	
	Task7	Task8	Task9	Task10	Task11	Task12	
Count	36	36	36	36	36	36	
Sum	717.8117	563.924	749.4133	784.386	713.1143	495.7967	
Average	19.93921	15.66456	20.81704	21.7885	19.80873	13.77213	
Variance	77.39442	72.93631	57.84941	111.2477	79.34475	30.94186	
ANOVA							
Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*	
Sample	4166.746	2	2083.373	39.82942	1.73E-16	3.018513	
Columns	10068.36	11	915.3053	17.49859	2.6E-28	1.812854	
Interaction	2794.102	22	127.0046	2.428044	0.000384	1.569143	
Within	20713.73	396	52.30739				
Total	37742.93	431					

HOMER	**Redesign**		**Original**		**Redesign Vs. Original**	
Task Number	Mean	Variance	Mean	Variance	t-statistic	p-value
Task1	9.81375	33.45429	7.8075	17.9062	0.820485	0.429362
Task2	7.48	6.320873	7.7475	55.48393	-0.10982	0.914531
Task3	21.32286	16.25509	23.27333	100.0073	-0.70696	0.494296
Task4	19.53545	22.42406	16.21667	55.58361	1.067999	0.308399
***Task5**	**12.95778**	**8.897487**	**9.520909**	**12.25061**	**2.501362**	**0.029435**
Task6	12.04125	22.34964	15.21833	104.5991	-0.9744	0.3508
Task7	21.39222	73.54636	16.27167	100.3846	1.320282	0.213554
***Task8**	**16.022**	**12.70778**	**7.735833**	**2.373208**	**6.722346**	**3.28E-05**
***Task9**	**23.31**	**26.65824**	**14.56**	**24.02378**	**3.558818**	**0.004482**
***Task10**	**26.944**	**93.04657**	**13.7775**	**24.8376**	**5.552898**	**0.000172**
***Task11**	**18.99333**	**31.42331**	**13.55**	**34.12173**	**2.633328**	**0.023269**
Task12	13.23222	7.675778	12.86667	63.24304	0.148919	0.884313

HOMER	**Redesign**		**Naïve**		**Redesign Vs. Naïve**	
Task Number	Mean	Variance	Mean	Variance	t-statistic	p-value
Task1	9.81375	33.45429	13.11182	82.69881	-0.95234	0.361383
Task2	7.48	6.320873	8.485	4.117714	-1.0149	0.331958
Task3	21.32286	16.25509	27.2	177.2057	-1.45811	0.172762
Task4	19.53545	22.42406	25.06875	150.8123	-1.54465	0.150698
Task5	12.95778	8.897487	13.29375	5.47429	-0.55493	0.590049
***Task6**	**12.04125**	**22.34964**	**25.95**	**67.922**	**-7.80672**	**8.24E-06**
Task7	21.39222	73.54636	22.15375	49.99711	-0.29084	0.776583
***Task8**	**16.022**	**12.70778**	**23.23583**	**85.73455**	**-3.12884**	**0.009597**
Task9	23.31	26.65824	24.58111	68.43852	-0.45262	0.659622
Task10	26.944	93.04657	24.644	128.1849	0.438565	0.669466
***Task11**	**18.99333**	**31.42331**	**26.88286**	**88.86479**	**-2.56194**	**0.026427**
Task12	13.23222	7.675778	15.2175	24.04114	-1.22911	0.244675

HOMER	**Naïve**		**Original**		**Naïve Vs. Original**	
Task Number	Mean	Variance	Mean	Variance	t-statistic	p-value
Task1	13.11182	82.69881	7.8075	17.9062	1.69816	0.117551
Task2	8.485	4.117714	7.7475	55.48393	0.332284	0.74592
Task3	27.2	177.2057	23.27333	100.0073	1.089411	0.299263
Task4	25.06875	150.8123	16.21667	55.58361	2.153723	0.054294
***Task5**	**13.29375**	**5.47429**	**9.520909**	**12.25061**	**3.066401**	**0.010729**
***Task6**	**25.95**	**67.922**	**15.21833**	**104.5991**	**2.885177**	**0.014833**
Task7	22.15375	49.99711	16.27167	100.3846	1.585204	0.141227
***Task8**	**23.23583**	**85.73455**	**7.735833**	**2.373208**	**5.822021**	**0.000116**
***Task9**	**24.58111**	**68.43852**	**14.56**	**24.02378**	**3.479298**	**0.005155**
***Task10**	**24.644**	**128.1849**	**13.7775**	**24.8376**	**2.735804**	**0.019376**
***Task11**	**26.88286**	**88.86479**	**13.55**	**34.12173**	**4.229066**	**0.001415**
Task12	15.2175	24.04114	12.86667	63.24304	0.811383	0.434352

	Redesign		**Naïve**		**Redesign Vs. Naïve**	
	Mean	Variance	Mean	Variance	t-statistic	p-value
*Redesign Vs. Naïve	16.92041	59.62691	20.8187	110.2252	-4.3602	2.47E-05
	Redesign		**Original**		**Redesign Vs. Original**	
	Mean	Variance	Mean	Variance	t-statistic	p-value
*Redesign Vs. Original	16.92041	59.62691	13.21216	64.94639	4.56392	1.07E-05
	Naïve		**Original**		**Naïve Vs. Original**	
	Mean	Variance	Mean	Variance	t-statistic	p-value
*Naive Vs. Original	20.8187	110.2252	13.21216	64.94639	7.904266	6.6E-13

B.7 – Voodoo Dolls Performance Results

Anova: Two-Factor With Replication							
SUMMARY	Task1	Task2	Task3	Task4	Task5	Task6	
Modified							
Count	12	12	12	12	12	12	
Sum	320.1	481.18	433.05	374.99	432.35	308.47	
Average	26.675	40.09833	36.0875	31.24917	36.02917	25.70583	
Variance	313.2096	284.9674	450.0393	342.7436	888.8907	275.5623	
	Task7	Task8	Task9	Task10	Task11	Task12	Total
Count	12	12	12	12	12	12	144
Sum	260.88	473.09	218.69	287.68	314.8036	303.99	4209.274
Average	21.74	39.42417	18.22417	23.97333	26.23364	25.3325	29.23107
Variance	236.873	702.6138	161.7903	152.9058	223.0016	157.3399	369.9127
SUMMARY	Task1	Task2	Task3	Task4	Task5	Task6	
Naïve							
Count	12	12	12	12	12	12	
Sum	100.19	183.58	271.78	204.55	217.72	157.5709	
Average	8.349167	15.29833	22.64833	17.04583	18.14333	13.13091	
Variance	10.48419	45.52605	140.3216	84.61515	63.68501	16.92501	
	Task7	Task8	Task9	Task10	Task11	Task12	Total
Count	12	12	12	12	12	12	144
Sum	203.47	400.66	217.28	217.49	230.22	297.4	2701.911
Average	16.95583	33.38833	18.10667	18.12417	19.185	24.78333	18.76327
Variance	40.0751	325.1542	68.71952	96.40979	131.0672	204.5107	130.0893
SUMMARY	Task1	Task2	Task3	Task4	Task5	Task6	
Original							
Count	12	12	12	12	12	12	
Sum	138.86	296.18	293.98	302.63	329.24	332.73	
Average	11.57167	24.68167	24.49833	25.21917	27.43667	27.7275	
Variance	48.45802	198.4996	230.7114	386.0895	405.8523	284.9013	
	Task7	Task8	Task9	Task10	Task11	Task12	Total
Count	12	12	12	12	12	12	144
Sum	220.58	601.06	192.53	254.67	299.47	308.72	3570.65
Average	18.38167	50.08833	16.04417	21.2225	24.95583	25.72667	24.79618
Variance	89.3068	1182.853	68.09015	184.775	277.7083	213.8034	355.4092
SUMMARY	Task1	Task2	Task3	Task4	Task5	Task6	

Total							
Count	36	36	36	36	36	36	
Sum	559.15	960.94	998.81	882.17	979.31	798.7709	
Average	15.53194	26.69278	27.74472	24.50472	27.20306	22.18808	
Variance	182.5999	273.7704	294.4333	290.5008	481.8031	224.3537	
	Task7	Task8	Task9	Task10	Task11	Task12	
Count	36	36	36	36	36	36	
Sum	684.93	1474.81	628.5	759.84	844.4936	910.11	
Average	19.02583	40.96694	17.45833	21.10667	23.45816	25.28083	
Variance	119.2458	743.8003	94.87658	142.3004	208.2292	181.0738	
ANOVA							
Source of Variation	*SS*	*df*	*MS*	*F*	*P-value*	*F crit*	
Sample	7950.672	2	3975.336	15.92172	2.23E-07	3.018513	
Columns	16979.92	11	1543.629	6.182432	2.1E-09	1.812854	
Interaction	6470.609	22	294.1186	1.177982	0.263441	1.569143	
Within	98873.27	396	249.68				
Total	130274.5	431					

	Redesign		**Naïve**		**Redesign Vs. Naïve**	
	Mean	Variance	Mean	Variance	t-statistic	p-value
*Redesign Vs. Naïve	29.23107	369.9127	18.76327	130.0893	6.288454	3.66E-09
	Redesign		**Original**		**Redesign Vs. Original**	
	Mean	Variance	Mean	Variance	t-statistic	p-value
*Redesign Vs. Original	29.23107	369.9127	24.79618	355.4092	2.148558	0.033354
	Naïve		**Original**		**Naïve Vs. Original**	
	Mean	Variance	Mean	Variance	t-statistic	p-value
*Naive Vs. Original	18.76327	130.0893	24.79618	355.4092	-3.89488	0.00015

B.8 – HOMER Quantitative Subjective Results

Table 12: Average subjective ratings of the HOMER 3DIT. Refer to Appendix B2 to see the form used to collect these metrics. A smaller frustration rating indicated more frustration.

Ratings	Naïve	Redesign	Original
Ease	3.083333	4	4.083333
Precision	3.083333	3.416667	3.416667
Frustration	2.25	3	2.666667
Comfort	3.833333	4.166667	3.916667
Vision	4	4	3.916667
Interaction	3.75	4.166667	4.083333

Table 13: Average subjective rankings of the HOMER 3DIT. Refer to Appendix B2 to see the form used to collect these metrics. A smaller frustration rating indicated more frustration.

Rankings	Naïve	Redesign	Original
Precision	2.833333	1.5	1.666667
Frustration	2.666667	1.666667	1.666667
Comfort	2.333333	1.416667	2.25
Ease	2.75	1.5	1.75
Overall	2.75	1.416667	1.833333
All	2.666667	1.5	1.833333

Table 14: A table showing the total fatigue reported for each task type per HOMER 3DIT flavor. Fatigue was a rating from 1 to 5.

Fatigue	Redesign	Naïve	Original
Control	22	26	24
To Back	70	78	89
To Front	23	28	30
At Back	23	27	28

Table 15: A table showing the total frustration reported for each task type per HOMER 3DIT flavor. Frustration was a rating from 1 to 5.

Frustration	Redesign	Naïve	Original
Control	29	33	29
To Back	93	130	100
To Front	34	41	37
At Back	38	41	31

Table 16: Quantitative observational data from the HOMER evaluation. Each row represents the totals for each usability issue shown along the top for the task types per 3DIT flavor shown on the left.

		Avoiding Screen Edges	Difficult Pinches	Avoiding Difficult Pinches	Pinch Confusion	Tracker Blockage	Cable Entanglement	Heisenberg Effect	Around Front	Avoiding Cables
Redesign	Control	0	0	0	0	1	0	0	0	0
	To Back	63	0	9	9	3	0	0	0	0
	To Front	25	0	1	4	0	0	0	1	0
	At Back	22	0	1	3	3	0	0	2	0
Naïve	Control	0	0	0	0	0	0	0	0	0
	To Back	55	27	57	5	1	0	0	0	0
	To Front	29	5	21	5	0	0	0	0	0
	At Back	27	6	15	3	1	0	0	4	0
Original	Control	0	0	0	0	0	0	1	0	0
	To Back	0	0	0	0	2	32	0	0	3
	To Front	0	0	0	0	3	7	0	0	1
	At Back	0	0	0	0	0	16	0	11	3

B.9 – Voodoo Dolls Quantitative Subjective Results

Table 17: Average subjective ratings of the Voodoo Dolls 3DIT. Refer to Appendix B3 to see the form used to collect these metrics. A smaller frustration rating indicated more frustration.

Ratings	Naïve	Redesigned	Original
Ease	3.833333	3.5	4.333333
Precision	3.5	3.333333	3.833333
Frustration	2.416667	2.333333	2.583333
Comfort	3.083333	3.333333	3.75
Vision	4.166667	3.916667	4.083333
Interaction	4.25	3.916667	4.083333

Table 18: Average subjective rankings of the Voodoo Dolls 3DIT. Refer to Appendix B3 to see the form used to collect these metrics. A smaller frustration rating indicated more frustration.

Rankings	Naïve	Redesigned	Original
Precision	1.833333	2.416667	1.75
Frustration	2.333333	2.25	1.416667
Comfort	2.5	2.333333	1.166667
Ease	2	2.5	1.5
Overall	2.083333	2.25	1.666667
All	2.15	2.35	1.5

Table 19: A table showing the total fatigue reported for each task type per Voodoo Dolls 3DIT flavor. Fatigue was a rating from 1 to 5.

Fatigue	Redesign	Naïve	Original
Control	32	26	30
To Back	62	53	66
To Front	31	31	35
At Back	75	74	70

Table 20: A table showing the total frustration reported for each task type per Voodoo Dolls 3DIT flavor. Frustration was a rating from 1 to 5.

Frustration	Redesign	Naïve	Original
Control	39	30	35
To Back	72	71	68
To Front	33	46	38
At Back	76	76	64

Table 21: Quantitative observational data from the HOMER evaluation. Each row represents the totals for each usability issue shown along the top for the task types per 3DIT flavor shown on the left.

		Bumping Hands	Bumping HMD	Crossing Hands	Preemptive Repositioning	Post Hoc Repositioning	Fine Repositioning
Redesign	Left Side	0	0	0	8	1	6
	Mixed	2	0	6	6	1	16
	Right Side	6	0	2	3	1	6
	Center	2	1	2	6	2	13
Naïve	Left Side	1	1	1	3	0	1
	Mixed	15	2	40	6	17	8
	Right Side	17	1	16	9	9	4
	Center	23	6	16	15	14	2
Original	Left Side	1	1	0	3	1	2
	Mixed	18	5	36	7	16	9
	Right Side	9	2	12	10	8	1
	Center	10	2	11	15	10	6
		Odd Hand Positions	Tracker Blockage	Cable Entanglement	Heisenberg Effect		
Redesign	Left Side	8	1	0	10		
	Mixed	3	1	0	13		
	Right Side	8	2	0	3		
	Center	9	7	0	20		
Naïve	Left Side	6	0	0	5		
	Mixed	6	3	1	6		
	Right Side	11	4	1	7		
	Center	9	5	0	9		
Original	Left Side	3	2	0	7		
	Mixed	9	8	2	13		
	Right Side	10	8	0	2		
	Center	9	7	0	5		