IMMERSIVE WELL PATH PLANNING:
THE ADDED VALUE OF INTERACTIVE IMMERSIVE VISUALIZATION

by

KENNY GRUCHALLA

B.S. New Mexico Institute of Mining and Technology, 1995

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Masters of Science

Department of Computer Science

2003

This thesis is entitled:

IMMERSIVE WELL PATH PLANNING:

THE ADDED VALUE OF INTERACTIVE IMMERSIVE VISUALIZATION

written by Kenny Gruchalla

has been approved for the Department of Computer Science

____

Chair of Committee

____

Committee Member

Date ______________

The final copy of this thesis has been examined by the signators, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Kenny Gruchalla (M.S. Computer Science)

Immersive well path planning: The added value of interactive immersive visualization

Thesis directed by Professor Clayton Lewis

The benefits of immersive visualization are primarily anecdotal; there have been few controlled users studies that have attempted to quantify the added value of immersion for problems requiring the manipulation of virtual objects. This research quantifies the added value of immersion for a real-world industrial problem: oil well path planning. An experiment was designed to compare human performance between an immersive virtual environment (IVE) and a desktop workstation with stereoscopic display. This work consisted of building a cross-environment application, capable of visualizing and editing a planned well path within an existing oilfield, and conducting an user study on that application. This work presents the results of sixteen participants who planned the paths of four oil wells. Each participant planned two well paths on a desktop workstation with a stereoscopic display and two well paths in a CAVE^{TM}-like IVE. Fifteen of the participants completed well path editing tasks faster in the IVE than in the desktop environment, which is statistically significant ($p < 0.001$). The increased speed in the IVE was complimented by an increase correct solutions. There was a statistically significant ($p < 0.05$) increase in correct solutions in the IVE. The results suggest that an IVE allows for faster and more accurate problem solving in a complex interactive three-dimensional domain.

Acknowledgments

Thanks to

Clayton Lewis, my thesis advisor, for his support and guidance throughout this work;

Liz Bradley, Geoffrey Dorn, and Kim Touysinhthiphonexay for their help and guidance offered as members of my thesis committee;

the staff at the B.P. Center for Visualization, for giving me the opportunity to conduct this research;

Gwen Pech for working tirelessly to keep the system running in an effort to support my work;

Jonathan Marbach for his countless contributions to this work, we spent many hours working together, laughing together and learning from each other;

Jo Solonika and Michael Gruchalla for their fearless proof-reading;

Michelle for her unwavering love, patience, and support;

and finally,

all the research participants for volunteering their time.

Contents

Tables

Figures

Chapter 1

Introduction

An immersive virtual environment (IVE) is a combination of hardware and software that provides a psychophysical experience of being surrounded by a computer generated scene. An IVE physically immerses users in a virtual world, where they can explore complex spatial systems by looking through them, walking around them, and viewing them from different perspectives.

There is a common assumption that IVEs provide an improved interface to view and interact with three dimensional structures, over more traditional desktop graphics workstations [van Dam 2000]. After all, an IVE differs greatly from traditional desktop graphics workstations in that it provides users a three-dimensional interface to view and interact with three-dimensional objects in a virtual world. In contrast, most three-dimensional desktop applications only use two dimensions, mapping two-dimensional input from a mouse into a three-dimensional virtual world. The three-dimensional interface provided by an IVE would seemingly provide a more natural and intuitive means for viewing and interacting with a three-dimensional virtual world.

Immersive applications have been envisioned for a variety of industrial application areas, including architectural walkthroughs [Brooks 1992], mechanical engineering [Yuan 1997], medicine [Foresberg 2000], and geophysical exploration [Winkler 1999, Frohlich 1999], to name a few. However, immersive technology has been slow to move outside the research laboratory and into industry. One of the main barriers in promoting immersive technology to industry is the fact that the benefits are primarily anecdotal. Very few formal studies have been performed to quantify the added value of immersion [Mizell 1990]. This work takes initial steps to explore and quantify the added value of immersion. The goal is to quantify the performance and usability of an IVE compared to a desktop graphics workstation for a real-world industrial task involving a complex three-dimensional domain.

Oil well design and optimization is a real-world task that requires the understanding of a complex three-dimensional domain. A cross-environment application, capable of visualizing and editing a planned well path within an existing oilfield, was designed and implemented for this study. Nineteen participants were asked to plan the path of four oil wells. Two well paths were planned on a desktop workstation with a stereoscopic display and two well paths were planned in an IVE. Each well path displayed a complexity value which loosely quantified the difficultly of drilling the path. Participants were given a goal complexity value for each path and instructed that the final path should not intersect any existing well paths. The participants' solutions were timed and the correctness of each solution was evaluated.

To evaluate the usability of an IVE, issues of cybersickness also need to be addressed. Some IVE users experience symptoms that parallel symptoms of classical motion sickness. Published estimates suggest that as many as 60% of the users experience some adverse effects and as many as 20% experience moderate to severe dizziness and nausea in IVEs [Potel 1998]. However, these published rates are misleading because not only has immersive hardware improved vastly since many of these studies, but most published cybersickness data was gathered from either military simulator experiments or experiments specifically designed to induce cybersickness. Neither type of experiment is representative of how the technology will likely be used in most industrial and academic settings. Therefore, cybersickness data was collected as part of this study. Participants were asked to evaluate the level of cybersickness symptoms they were experiencing before and after both the desktop and immersive treatments of the experiment.

1.1 Hypothesis

An immersive virtual environment will allow for faster and more accurate problem solving in a complex interactive spatial domain.

1.2 Contributions

This research makes two main contributions to the fields of virtual environments, three-dimensional interaction, and human computer interfaces:

- Adds to the current state of knowledge by quantifying the impact of immersion on a solution of real-life industrial problem.
- Adds to the current state of knowledge by collecting cybersickness data during an

experiment that is more representative of industrial IVE use.

1.3 Related Work

Most human performance virtual environment studies have focused on comparing various navigation and manipulation techniques within the same virtual environment. Several human performance studies have evaluated different virtual environment attributes, such as stereoscopic and head-tracked displays. Stereoscopic displays tend to improve performance in both immersive and desktop environments. Conversely, head-tracked displays only improve performance in immersive environments, and actually degrade performance in desktop environments. Only a few studies have attempted to compare IVEs with traditional desktop environments. This work has focused on comparing navigation and identification in the two environments. Only one study, to our knowledge, attempted to compare an interaction task between the two different environments and its results were inconclusive. These previous research studies will be presented in the following sections, although the discussion of the prior navigation and manipulation research will be postponed until Chapter three.

1.3.1 Mizell, et al.

Mizell et al. conducted an experiment to determine what features of an IVE provide users with a better understanding of complex three-dimensional geometry [Mizell 2000]. Participants were shown a virtual sculpture inside an IVE and were tasked with assembling a physical replica of the virtual sculpture. The sculptures used for the experiment were a set of abstract shapes that had no intrinsic meaning. Participants were given an empty peg board and set of rods and instructed to construct a physical replica of the virtual sculpture they were being shown, as quickly and as accurately as possible. The experimenters compared four IVE display modes by independently varying head-tracked verses non-head-tracked, and stereoscopic verses monoscopic modes. The results indicated that the head-tracked mode consistently produced faster solutions with lower error rate than the non-head-tracked mode. Participants took over forty percent longer and committed three times as many errors when controlling their perspective with a joystick. The study found no statistically significant difference between stereoscopic and monoscopic displays.

1.3.2 Barfield, Hendrix, and Bystrom

Barfield, Hendrix, and Bystrom studied the effects of stereoscopic images and head-tracking upon performance on a desktop computer [Barfield 1997]. Similar to the Mizell study, participants viewed a virtual abstract wire sculpture and were asked to select a corresponding sculpture from one of three drawings presented on paper. The treatments included monoscopic head-tracked images, stereoscopic head-tracked images, monoscopic non-head-tracked images, and stereoscopic non-head-tracked images, all on a 19-inch color monitor. In contrast to the Mizell study, the results indicated that neither stereo nor head-tracking on a desktop workstation improved the accuracy of selecting the correct paper representation.

1.3.3 Boritz and Booth

Boritz and Booth investigated the ability to locate points in virtual three-dimensional space on a desktop computer [Boritz 1997]. Four different display modes were studied: a monoscopic head-tracked mode, a stereoscopic head-tracked mode, a monoscopic non-head-tracked mode, and a stereoscopic non-head-tracked mode. Participants were asked to locate a point that was located along the X, Y, or Z axes from a fixed starting position. Like the Barfield, Hendrix, and Bystrom study, head-tracking on a desktop computer had no significant effect. Unlike the Barfield, Hendrix, and Bystrom study, participants did have a much stronger performance on the workstation in the stereoscopic modes than in the monoscopic modes.

1.3.4 Sollenberger and Milgram

Sollenberger and Milgram tested performance in tracing a three-dimensional path on a desktop workstation. Participants were asked to trace a path along a three-dimensional stick-figure tree, from the root to a leaf. Four display modes were studied: a static monoscopic display, a monoscopic display allowing participants rotational control of the scene, a static stereoscopic display, and a stereoscopic display with rotational control. The static monoscopic display resulted in the least successful performance. Results from the static stereoscopic display were slightly better, best of all was the monoscopic display with rotation control. The stereoscopic display with rotational control produced the most successful results.

1.3.5 Arns, Cook, and Cruz-Neira

Arns, Cook, and Cruz-Neira conducted a user study comparing statistical data analysis on a desktop and an IVE [Arns 1999]. The experiment compared both identification and interaction tasks on a desktop and an IVE. During the identification tasks, participants were asked to identify clusters of data and identify the dimensionality of data. During the interaction tasks, participants were asked to “brush” clusters, marking data points with colored glyphs.

The results of the study suggested that IVEs significantly improve productivity for structure and feature detection tasks in the analysis of highly dimensional data. Participants performed almost twice as well when identifying clusters in the IVE, with an eighty percent correct rate verses a forty-seven percent on the desktop. Participants performed equally well identifying the dimensionality in the two environments. The performance in the IVE was as good as or better than the performance on the desktop in the visualization task, but in the interaction tasks the desktop was faster. Participants' brushing times were lower on the desktop than on the IVE. However, drawing any conclusions is difficult, since the brushing times had a large standard deviation.

1.3.6 Slater, et al.

Slater et al. also conducted an experiment comparing performance on a desktop computer and an IVE. [Slater 1996] Participants witnessed a sequence of moves on a virtual Tri-Dimensional Chess board, and were then asked to replicate the sequence on a actual Tri-Dimensional Chess board. Half the participants wore a Virtual Research Flight Helmet to view the virtual chess board, initiating each move with a three-dimensional mouse. The other half of the participants viewed the virtual chess board on a TV monitor, initiating each move with the three-dimensional mouse.

The experiment showed that immersion improved task performance. On average, participants reproduced the correct moves only fifty percent of the time when the moves were viewed on the TV monitor, while participants who viewed the moves with the head-mounted display reproduced the correct moves an average of eighty percent of the time.

1.3.7 Ruddle, Payne, and Jones

Ruddle, Payne, and Jones designed a virtual building walk-through experiment to compare a helmet-mounted display with a desktop monitor display [Ruddle 1999]. Participants would learn the layout of large-scale virtual buildings through repeated navigation. Participants would navigate two large virtual buildings, each consisting of seventy rooms. A repeated measure design was used, where each participant navigated one building four times using the head-mounted display, and navigated the second building four times using the desktop workstation.

On average, participants who were immersed in the virtual environment using the helmet-mounted display navigated the buildings twelve percent faster. The decreased time was attributed to the participants utilizing the ability to "look around" while they were moving when immersed, as the participants spent eight percent more time stationary when using the desktop workstation. Participants also developed a better understanding of the layout of the building, as evidenced by their knowledge of relative distance between locations in the buildings.

1.4 Summary

In this chapter, we have introduced the need to evaluate the added value of immersion, our hypothesis, and the contributions of this work. Chapter two presents a detailed description of the hardware apparatus used for this study. Both the IVE and the desktop environment on which this study was conducted are described. Chapter three presents a detailed description of the software testbed application, or software apparatus, used in this study. This chapter includes a look at the previous work on three-dimensional interaction techniques which influenced the design of the testbed application. This chapter also provides an overview of the well path planning techniques employed by the testbed application. Chapter four presents the experimental design. This includes a description of the experimental method, participant population, and the experimental tasks. Chapter five presents the results of the user study. Chapter six presents the conclusion of this work, including a discussion of the main contributions of this research, suggestions for improvement, and possibilities for future work.

Chapter 2

Environment

This chapter introduces and describes the environments in which this study was performed. The study involved experiments in two environments, an immersive virtual environment and a desktop environment.

2.1 Immersive Virtual Environments

An immersive virtual environment (IVE) is an environment created with a combination of hardware and software that provides the its user with a psychophysical experience of being surrounded by a computer-generated scene. An IVE gives the user a sensation of presence with the objects in the scene providing the user of the system an egocentric view of a scene (i.e., a scene constructed from the user's point of view). This egocentric view is typically created by a head tracked stereoscopic display with a wide field of view. Tracking the position and orientation of the user's head allows the user to move around in the virtual world and see that world from different angles. This egocentric, stereoscopic display creates the illusion that the objects in the scene are three dimensional and in the presence of the user [vanDamn 2000].

The earliest concept of an IVE is often accredited to Plato in 370 B.C. with the writing of Book VII of the Republic. In the “Allegory of the Cave,” Plato describes a physical environment in which the ideas of perception, reality, and illusion could be explored. The first modern concept of an IVE was introduced by Ivan Sutherland in 1965 [Sutherland 1965]. Three years later, Sutherland implemented his concept, a head-mounted display (HMD) with two cathode ray tubes that presented the wearer with a stereoscopic three dimensional view of a simple computer generated scene. The system was head-tracked by coupling the HMD to a six degrees of freedom mechanical sensing device [Sutherland 1968].

Since Sutherland's introduction of the IVE in 1965, it has evolved into many forms. HMDs are still in common use, although the cathode ray tubes have been replaced by eye-glass sized LCD screens. Various limitations of the HMD have given rise to projection-based systems where stereo images are projected onto one or more screens. Projection-based systems allow multiple people to communicate and interact in the virtual environment. The first and the most widely known projection-based system is the CAVE™ (CAVE Automatic Virtual Environment) developed by the Electronic Visualization Laboratory (EVL) at the University of Illinois Chicago. The CAVE™ surrounds the user with four screens (three walls and a floor). The CAVE™ was introduced at SIGGRAPH'92 [Cruz-Neria 1992]. Since that introduction, several projection-based systems have been developed, including the system used for this study.

2.1.1 IVE at the B.P. Center for Visualization

The IVE used for this study is located on the University of Colorado campus at the B.P. Center for Visualization. The IVE at the B.P. Center for Visualization is a Mechdyne MD Flex^{TM}, which is a configurable large screen projection-based system. In closed configuration (see figure 2.1) the MD Flex^{TM} is a 12' x12' x10' theater, resembling a CAVE^{TM}-Like system. The MD Flex^{TM} can be re-configured to a 36' x12' x10' open configuration or presentation mode (see figure 2.2). The closed configuration provides a greater sense of immersion, therefore, for the purposes of this study only the closed configuration was used. The MD Flex^{TM} consists of four walls: three rear-projected screens measuring 12' x10' which form the right wall, back wall, and left wall of the IVE, the fourth wall is the 12' x12' floor which is projected from above.

Image Courtesy of Mechdyne Inc.

Figure 2.1: Photograph and drawing of the IVE at the B.P. Center for Visualization shown in closed configuration. The IVE is a MD Flex, which is a 12' x12' x10 projection-based immersive environment.

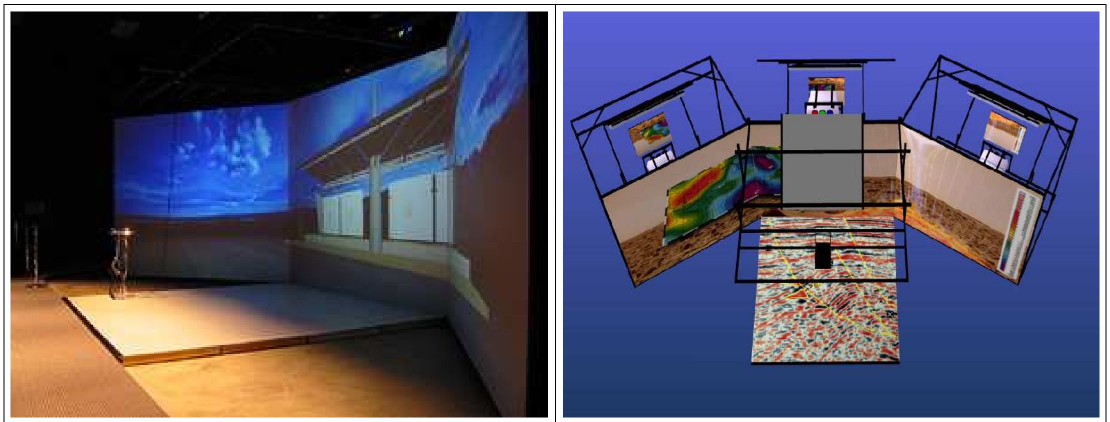

Image Courtesy of Mechdyne Inc.

Figure 2.2: Photograph and drawing of the IVE at the B.P. Center for Visualization shown in open configuration.

The four display screens are driven by one Silicon Graphics Incorporated (SGI) Origin 3800 computer with four SGI Infinite Reality3 graphics pipes. Each pipe feeds a Barco Reality 909 projector. Images from the projectors are bounced off mylar mirrors so that the IVE will fit within a constrained space (the projectors require a ten-foot throw distance). The Barco Reality 909 projectors are capable of up to 1600x1280 stereo resolution; however, due to other hardware constraints, the resolution used for this study was limited to 1024x768.

A three-dimensional effect is created inside the IVE through active stereo projection. This stereo projection is achieved by projecting an image for the viewer's left eye followed by an image for the viewer's right eye. Viewers wear infrared CrystalEyes^{TM} active stereo LCD shutter glasses to view the stereoscopic images. The shutter glasses resemble a large pair of sunglasses (see figure 2.3). Infrared emitters synchronize the glasses with the graphics pipes. When the computer renders the image for the left eye, the right eye shutter is closed. Similarly, when the computer renders the image for the right eye, the left eye shutter is closed. This shuttering action creates the illusion of three-dimensional images.

Figure 2.3: Photograph of the CrystalEyes™ LCD shuttered glasses with an InterSense™ InterTrack motion tracker mounted on top. This pair of glasses was worn by participants during the immersive treatments. A similar pair glasses without the motion tracker was worn by test participants during the desktop treatment.

The IVE has an InterSense™ VET 900 tracking system. An InterSense™ InterTrack motion tracker is mounted to the CrystalEyes™ shuttered glasses, allowing the user's head position and orientation to be tracked by the VET 900. The position and orientation information is used by the software to generate the egocentric perspective. The InterSense™ tracker has a resolution of 1mm for position and 0.1 degrees for orientation.

The sole interaction device used in this study was a wired InterSense wand (see figure 2.4). The wand is a hardware device that can be thought of as three-dimensional, six degrees of freedom mouse. The wand has four buttons and a pressure sensitive joystick. Like the user's glasses, the wand's position and orientation are tracked by the tracker.

Figure 2.4: Photograph of the InterSense^{TM} IS-900 six degree of freedom wand. This was used as the interaction device for the immersive treatments of this study.

2.1.2 Software Libraries

The software application used for this study was built on top of the CAVELib^{TM} and Open Inventor^{TM} libraries. CAVELib^{TM} is a C library that adds an abstraction layer over many of the hardware-specific details of the numerous varieties of IVE display systems. CAVELib^{TM} provides functions to synchronize the screens and generate the correct perspective on each individual screen. The CAVELib^{TM} also provides access to the state of all the tracked devices. The abstraction layer provided by CAVELib^{TM} facilitates the porting of immersive applications from one IVE system to another. [Pape 1996, Czernuszenko 1997].

Open Inventor is an object-oriented three-dimensional application programming interface providing a library of objects and methods used to create interactive three-dimensional graphics applications. Open Inventor is based on the scene graph programming model: a scene graph is a directed acyclic graph that organizes and stores all of the data needed to render a three-dimensional scene. Open Inventor provides a standard for the development of cross-platform (e.g., IRIX, Linux, and Windows), cross-environment (e.g., immersive virtual environments and desktop environments), three-dimensional applications. [Wernecke 1994] That is, an application built with Open Inventor for a desktop environment can be ported to an immersive environment with minimal effort.

2.2 Desktop Environment

In contrast to the IVE, a desktop environment restricts the user to a exocentric view, in which the user is kept on the outside looking in. The desktop equipment used for this study is similar to desktop computers found in many homes and offices. The desktop equipment consisted of a 21-inch SGI monitor, a 3-button mouse, and an SGI keyboard (see figure 2.5). Unlike those in most homes and offices, the desktop interface in this study was connected to a SGI Origin 3800 (the same machine used in this study' s immersive experiments). The monitor' s images, like the screens in the immersive experiments, were driven by a SGI Infinite Reality3 graphics pipe and constrained to a resolution of 1024x768. The images produced on the desktop were rendered in stereo, producing a stereoscopic display when used in conjunction with a pair of CrystalEyes^{TM} active stereo LCD shutter glasses. Unlike the immersive environment, the desktop environment did not include head tracking.

Figure 2.5: Photograph of the desktop workstation used for this study. Although the desktop workstation has four monitors, only one (the second from the right) was used in this study.

Chapter 3

Immersive Drilling Planner

As mentioned in the previous chapter, the IVE used for this study is located on the University of Colorado campus at the B.P. Center for Visualization. One of the Center's objectives is to conduct research and develop prototypes for drilling visualization and drilling design optimization. The Immersive Drilling Planner (IDP) was started as a long-term project to explore the impact of immersive visualization for drilling, in an effort to reduce drilling costs, risks, and time.

3.1 Basic Drilling Concepts

Modern drilling equipment can be controlled so that a well can be drilled at a predetermined angle and directed toward a predetermined target location. This type of drilling is known as *directional drilling*. [Hyne 2001] The most common use of directional drilling is in offshore fields, where the expense of creating a drilling platform is considerable. An offshore field, particularly those under deeper waters, must be exploited by a small number of fixed platforms. Each platform is capable of tapping a sector of the field through a cluster of wells. Directional drilling is becoming increasingly common onshore in urban and environmentally sensitive areas, since exploiting a field through this method has a much smaller environmental footprint than does exploiting the same field with straight hole drilling. [North 1990]

Oilfields exploited by directional drilling can quickly become a tortuous underground labyrinth of wells, creating a very complex spatial domain (see figure 3.1). When planning a new well in a mature field, the planner must take special care that the new well does not collide with any existing wells. A collision with an existing well can cause a *blow out*, an uncontrolled flow of fluids up a well. Blow outs can lead to fires and explosions resulting in the loss of the the drilling rig and possibly the loss of life. [Hyne 2001, Hyne 1984] One of the design goals of the IDP was to provide well planners a way to plan a safe path for a new well in a mature oilfield.

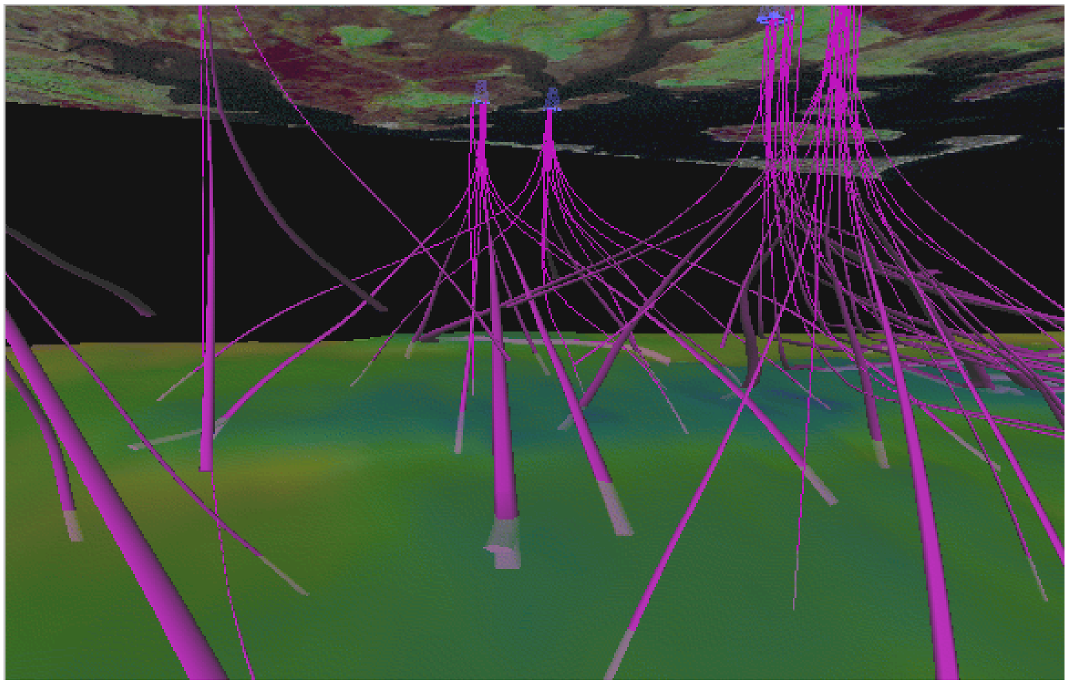

Figure 3.1: Snapshot of a virtual oilfield constructed from a well log dataset donated by British Petroleum. Mature oilfields can be very complex three-dimensional structures.

3.1.1 Well Path Design

This section briefly introduces the basic concepts used by the IDP to describe and design well paths. This section is by no means exhaustive in its coverage of the data structures and algorithms implemented by the IDP; rather, this section attempts to introduce a few basic concepts so the reader can understand how well paths are constructed and edited with the IDP. Southren provides an in depth exploration of the subject. [Southren 2000]

A *well path* is a continuous series of curved and straight sections with boundary locations between them called *salient points* (see figure 3.2). A salient point describes the position and attitude of the well at a particular location on the well path. For the purposes of this study we use three different types of line sections to connect adjacent salient points. The three types of sections are:

2. Straight line section. A straight line section connects two adjacent salient points with a straight line. In this case the attitude vector of the two adjacent points are equal to the vector of the line joining them.
3. Constant radius curve section. A constant radius curve section connects two adjacent salient points with a curve of constant radius (see figure 3.3).
4. Kink section. A kink section accommodates an instantaneous change in attitude between two straight sections. This only occurs at the base of the drill floor.

Using the curve, the path between any two adjacent salient points can be determined by interpolation.

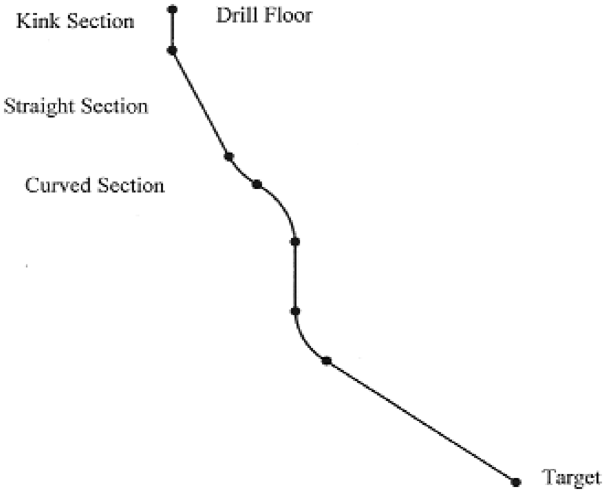

Figure 3.2: A typical well path illustrating the salient points and the sections between them. Three types of line sections are used to connect adjacent salient points: kink, straight, and constant radius curve. Image courtesy of Tech-21 Solutions Ltd. [Southren 2000]

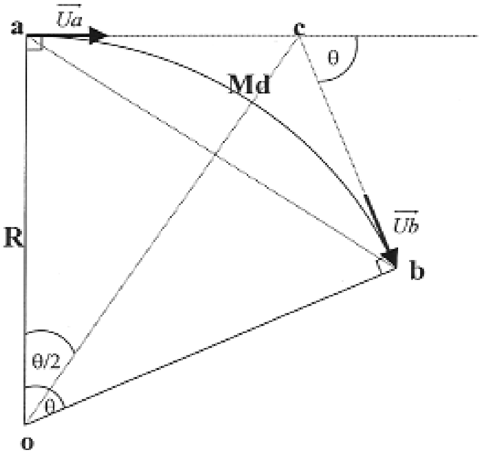

Figure 3.3: Geometrical construction of a constant radius curve. For a curve between points a and b the radius of the curve, R, is perpendicular to the curves starting and ending attitude vectors. Image courtesy of Tech-21 Solutions Ltd. [Southren 2000]

The IDP allows well paths to be edited through the *pull point method*. The pull point method allows a region of the well path to be altered. An edit region is defined by selecting a start and end position on a well path. If a salient point does not exist at either of these locations, a salient point will be created by interpolating between existing salient points. Then a pull point is defined. The pull point has both position and attitude. The section of the well path within the edit region is re-routed to pass through the the pull point' s position, while preserving the attitude of the well path at the start, end, and pull point positions (see figure 3.4).

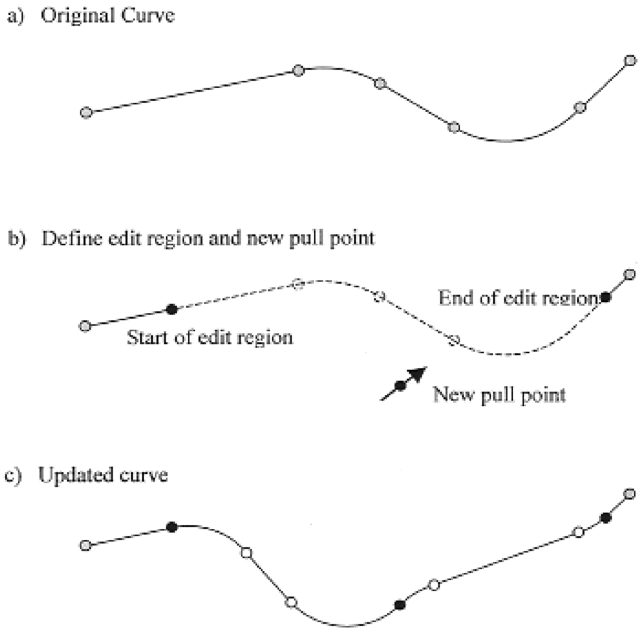

Figure 3.4: Editing a well path using the pull point method. The section of the well path within the edit region is re-routed to pass through the the pull point' s position with attitude at that position. The attitude of the well path at the salient points marking the start and end of region are preserved. Image courtesy of Tech-21 Solutions Ltd. [Southren 2000]

In the process of re-routing the well path through the pull point, intermediate salient points will be added to the well path to ensure that it can be solely constructed with constant radius curve segments, straight line segments, and one kink segment. Intermediate salient points are added to a well path using an iterative curve fitting algorithm. [Southren 2000]

The location of a real well path cannot be known with complete certainty. A position in a well path is determined by surveying instruments that are placed down the drilled hole. The surveying instruments typically measure attitude and *measured depth.* Measured depth is a measure of the length along a well path. As these readings are subject to error, there are uncertainties in a well path' s position that accumulate with depth. The errors are aligned along the well direction, described by three mutually perpendicular unit vectors (see figure 3.5). The error forms an elliptical volume perpendicular to the well path. Accumulating the errors at each point along the well enables an uncertainty surface to be constructed (see figure 3.6). [Southren 2000, North 1990]

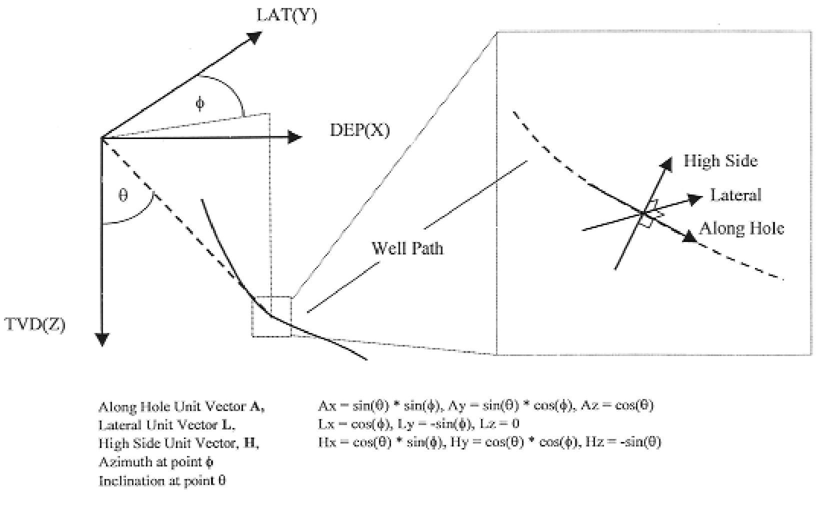

Figure 3.5: Illustration of well path uncertainty vectors. Uncertainties in the well path' s position are described by three mutually perpendicular vectors, *Along Hole, High Side,* and *Lateral.* Image courtesy of Tech-21 Solutions Ltd. [Southren 2000]

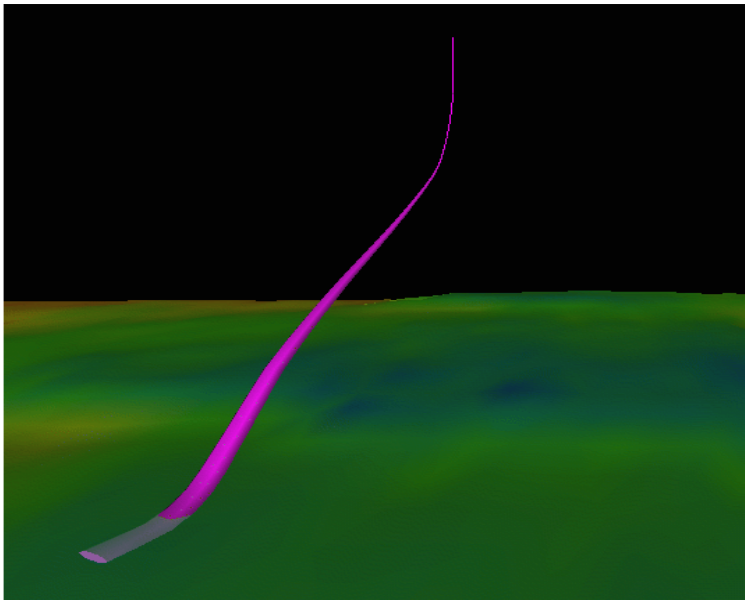

Figure 3.6: Snapshot showing a well path' s uncertainty surface, shown in purple.

Although the direction and angle of the drill can be controlled, the more curvature in a planned well path the more difficult the well will be to drill. In reality, a multitude of geological, geographical, and physical factors drive the complexity of a well, but currently the IDP only provides a simple model: a weighted sum of curvature along the well path. The weight relates to the "sharpness" of the curve; sharper curves have a higher weight than softer curves. This complexity model provides the planner feedback during the planning process.

3.2 Immersive Drilling Planner Design

The IDP development was started at the B.P. Center for Visualization in the fall of 2002 by Kenny Gruchalla and Jonathan Marbach. The IDP capabilities include interactive well planning integrated with geological and geophysical data, visualizations of well uncertainty, and design optimization for the development of mature fields. The vision for the IDP was to create an immersive visualization application that could be used as a testbed to explore the added value of immersion, new interfaces for three-dimensional editing, team dynamics, and collaboration in a real-world application space.

The IDP was designed to operate in a variety of visualization environments, including large screen systems, immersive bench displays, and desktop workstations. To support both immersive environments and desktop workstations, two implementations of the IDP have been created. Both implementations share the same IDP code base and identical scene graphs; the only difference is the front-end user control that allows navigation through the scene and the manipulation of the objects in the scene. The IVE version of the IDP can be run directly on a desktop workstation using the CAVELib^{TM} simulator. However, the simulator was designed as a tool to test immersive applications, not as a production desktop interface. [Pape 1996] Therefore, a separate front-end interface was designed for the desktop version of the IDP.

3.2.1 Desktop Design

Although we exist in a three-dimensional world, there are fundamental difficulties in understanding and interacting with three-dimensional spaces [Herdon 1994]. Interacting with a three-dimensional space through a two-dimensional interface, such as the mouse, only complicates matters. Much work has been done in the area of three-dimensional human-machine interfaces. The Open Inventor library is a result of this work. It provides a three-dimensional viewer and defines a user interface that is becoming an industry standard for interacting with a three-dimensional world on a desktop. The IDP desktop design utilizes the Open Inventor standard. Specifically, the Open Inventor SoXtExaminerViewer is used by the desktop implementation of the IDP as the front-end user interface (see figure 3.7).

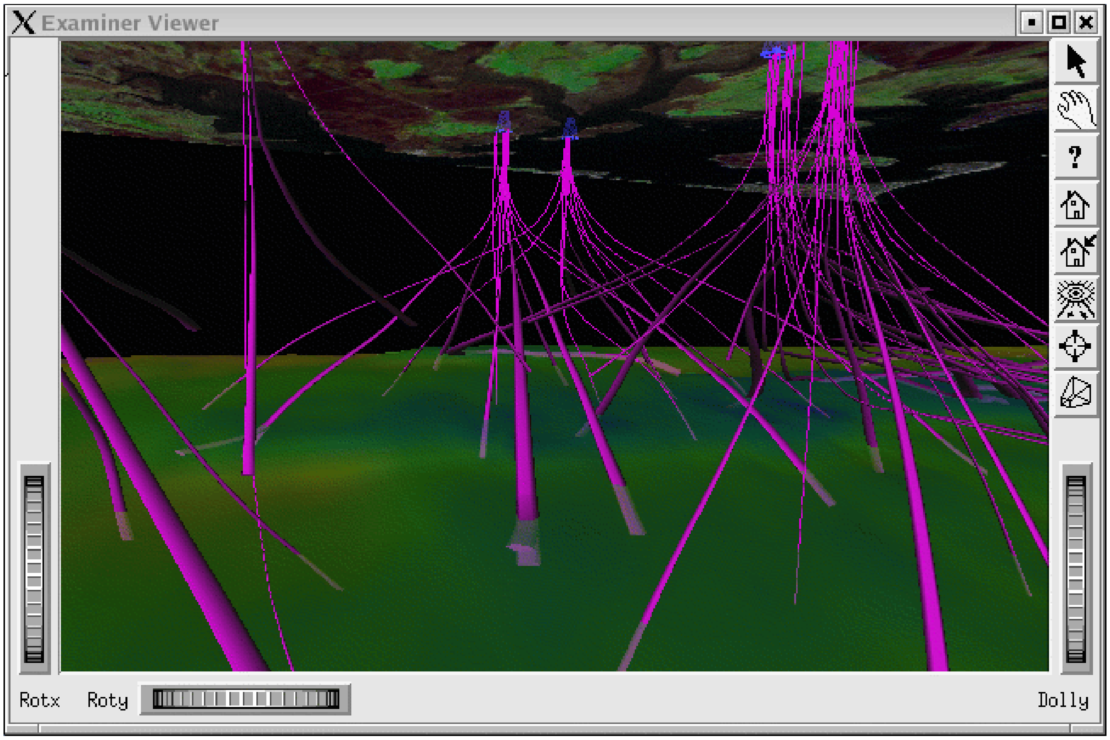

Figure 3.7: Snapshot of IDP desktop interface. The IDP desktop utilized the Open Inventor SoXtExaminerViewer as a front-end interface.

3.2.1.1 Navigation Design

All the functionality to navigate or change the view of the scene, was provided by the SoXtExaminerViewer. The user can manipulate their view of a scene by generating mouse click-and-drag events in the render area (right mouse down rotates the scene, middle mouse down pans the scene, and right and middle mouse down zooms in and out of the scene). The user can also manipulate a scene with three thumbwheel widgets which control zooming and rotation about the X and Y axes.

3.2.1.2 Interaction Design

To interact with objects in the scene, Open Inventor manipulators are used. The manipulators provide a means to position and rotate three-dimensional objects in three-dimensional space with a two-dimensional mouse. A SoHandleBoxManip is used to position interactive objects in the desktop version of the IDP. A SoHandleBoxManip draws a bounding box around the interactive object (see figure 3.8). The SoHandleBoxManip responds to click-and-drag mouse events by translating the interactive object it surrounds. The SoHandleBoxManip also provides scaling functionality, which is not used in the IDP. A SoTrackballManip is used to rotate interactive objects in the desktop version of the IDP. A SoTrackballManip wraps the interactive object with three circular stripes. These stripes are oriented like wheels that can be spun in the X, Y, and Z axes (see figure 3.9). The SoTrackballManip responds to click-and-drag mouse events by rotating the interactive object it surrounds. Clicking in an area between the stripes allows the user to rotate the object freely in three dimensions; clicking on the stripes allows the user to constrain rotation in the X, Y, or Z axes.

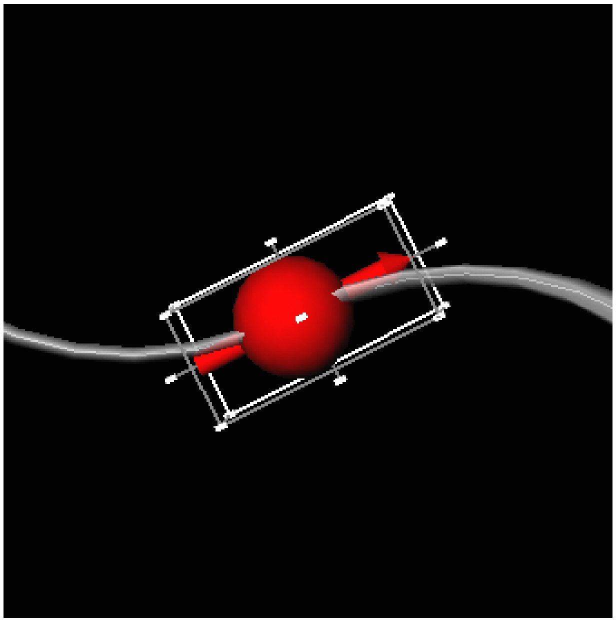

Figure 3.8: Snapshot of SoHandleBoxManip (the white bounding box) which is used translate objects.

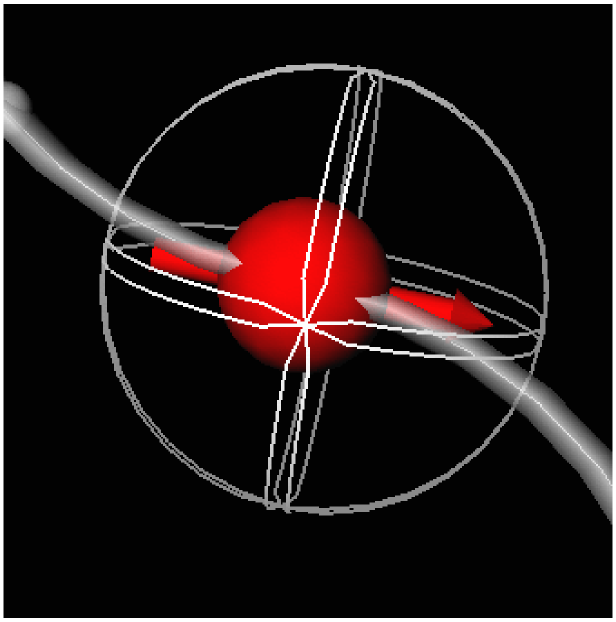

Figure 3.9: Snapshot of SoTrackballManip (the white circular stripes) which is used to rotate objects.

3.2.1.3 Testbed Viewer

For the purposes of the user tests, several modifications were made to the desktop interface. The standard SoXtExaminerViewer provides functionality, through the GUI buttons on left side of the dialog (see Figure 3.7), that is not yet available in the immersive version of the IDP. In an effort to simplify the desktop interface and equalize the functionality between the two environments, the additional functionality offered by the SoXtExaminer viewer was disabled by removing the several buttons (see Figure 3.10). Two buttons were added to the viewer; to provide a mechanism to toggle between the SoHandleBoxManip and SoTrackballManip manipulators. A readout was also added to provide complexity value feedback.

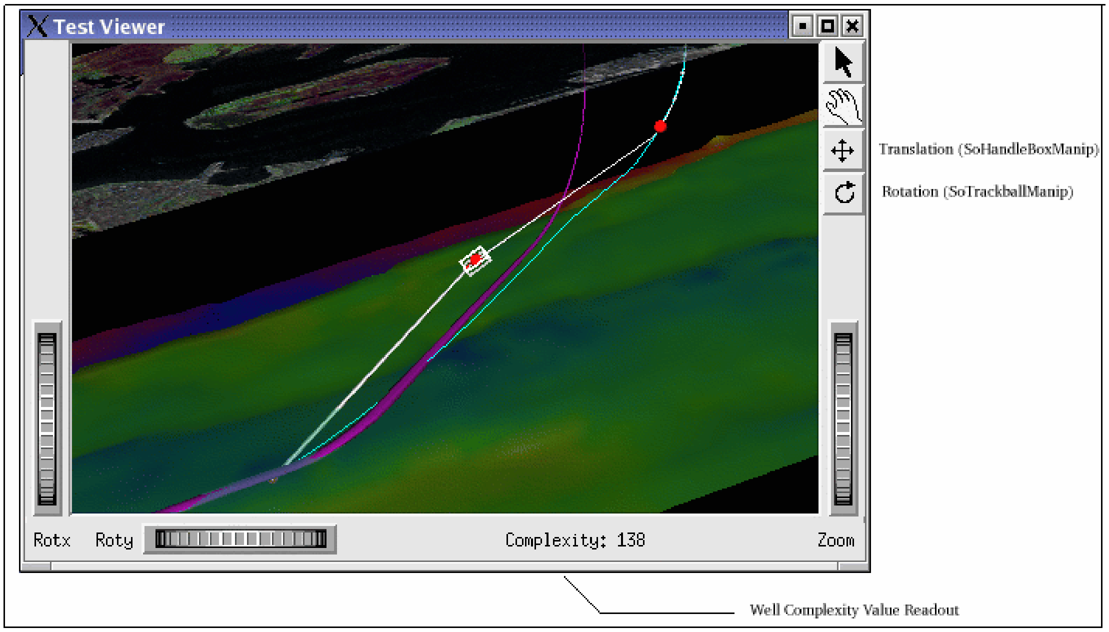

Figure 3.10: Annotated snapshot of the desktop test viewer. The standard Open Inventor viewer, SoXtExaminerViewer, used by the IDP was modified for the user study. The interface was simplified by removing several of the function buttons. Two buttons were added to provide a mechanism to toggle between the two types of Open Inventor manipulators used by the IDP. A complexity readout was also added to the dialog.

3.2.2 Immersive Design

The three-dimensional user interface is a critical component of a immersive virtual environment's usability. Bowman [1999] has shown that immersive interaction techniques based on natural and real-world metaphors often exhibit serious usability problems. Therefore, careful thought must go into the design of user interfaces and interaction techniques for immersive applications. Fortunately, a large body of work in the field of immersive human-computer interaction exists. The design of the IDP is based on many of the specific results and guidelines of that work.

3.2.2.1 Navigation Design

Navigation is the most universal user action in large-scale immersive environments, and consequently several implementations and user studies of immersive navigation techniques have been reported. Mine [1995] provides an overview of the most widely used navigation techniques. Bowman [1999] provides a set of guidelines for the design navigation techniques. This section will briefly discuss the most common navigation techniques and describe the navigation technique implemented in the immersive version of the IDP.

Physical navigation is the simplest and most natural navigation model in an IVE. This model maps a user' s physical movements, such as walking, into corresponding motions in the virtual world. Physical navigation is cognitively simple, requiring no special action on part of the user, and it has been shown to help users maintain spatial awareness of their location in the scene and the objects around them [Usoh 1995]. However, if the size of the virtual world exceeds the physical boundaries of the IVE, physical navigation cannot be used alone. This is the case of the IDP. An oilfield scaled to fit wholly within the physical boundaries of the IVE would be unusably small.

The two most commonly used immersive navigation techniques after physical navigation, *gaze-directed steering* and *pointing,* can both be used to help overcome the limitations of physical navigation [Mine 1995]. In gaze-directed steering, the orientation of user' s head is used to determine the direction of travel. The user' s viewpoint travels along the direction the user is currently looking (approximated by the direction the user' s head is pointing). In the pointing technique, the direction of motion depends upon the current orientation of the user' s hand or hand held device (in the case of IDP, the InterSense wand) [Mine 1995].

User studies have suggested that the pointing technique is superior to the gaze-directed technique for general-purpose applications that require speed and accuracy [Bowman 1997, Coninx 1997]. The pointing technique is also more comfortable and allows the user to look and move in different directions. However, gaze-directed steering has been shown to have distinct advantages in its ease of use and learning, particularly with novice IVE users [Bowman 1997].

The IDP implements a combination of physical navigation and pointing techniques. An IDP user can navigate the portion of the oilfield inside the IVE by simply walking within the IVE. To reach areas of the field outside of the bounds of the IVE, the user points the wand in the direction of desired travel. Pressing forward on the wand' s joystick will “drive” the user in the direction the wand is pointing. Pressing backwards on on the wand' s joystick will “drive” the user in the opposite direction. The joystick is pressure sensitive and the amount of pressure exerted on it maps to the speed of travel. Pressing right or left on the joystick will rotate the scene around the user.

3.2.2.2 Interaction Design

Interaction with a virtual object involves selecting, positioning and rotating the object in the virtual environment. The classical interaction technique provides the user with a virtual hand, whose movements correspond to the movements of the tracked input device (see figure 3.11). Selection and manipulation of objects simply involve touching an object with the virtual hand, then positioning and orienting the virtual hand in the IVE. This technique has been shown to be intuitive [Bowman 2001]. However, the technique has a major limitation in that the

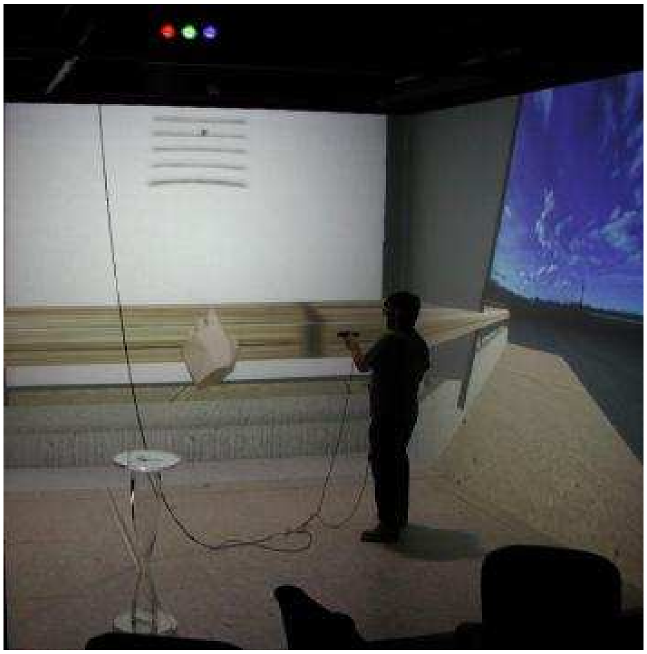

Figure 3.11: Photograph of a immersive application that implements the classical virtual hand interaction technique. Selection and manipulation of objects simply involve touching an object with the virtual hand, then positioning and orientating the virtual hand in the IVE.

user can only manipulate objects that are physically within reach.

Several techniques have been suggested to overcome this limitation, such as the *Go-Go* and the *ray-casting* techniques. The Go-Go technique allows the user an extended reach. When the user extends the virtual hand farther than a predefined threshold distance, a nonlinear mapping stretches a virtual arm thereby extending the user's reach [Poupyrev 1996]. User studies have indicated that this technique is a viable immersive interaction technique. [Poupyrev 1997] Another common interaction technique is ray-casting. With ray-casting, a virtual ray eminates from the input device; when the ray intersects an object, the object can be manipulated. User studies comparing the Go-Go and ray-casting techniques indicate that ray-casting preforms more effectively over a wide range of possible object distances and sizes. [Bowman 1999]

The IDP implements a variation on the ray-casting technique that allows objects to be selected, positioned, and rotated. In this variation, a virtual ray extends from the wand and interactive objects are highlighted when intersected by the virtual ray (see figure 3.12). Once an interactive object is intersected, pressing and holding the lower left wand button will select and drag the object. When the object is selected with the lower left wand button, it is effectively "speared" on the virtual ray. Then, wherever the wand moves, the speared object follows. When the user releases the wand's lower left button, the object is released at its current location. While an object is being dragged, its orientation remains constant, only its position is changed. Once an interactive object is intersected, pressing and holding the lower right wand button will select and rotate the object. When the object is selected with the lower right wand button, it will mimic the orientation of the wand. When the user releases the wand' s lower right button, the object is released at that orientation. While an object is being rotated, its position remains constant; only its orientation is changed.

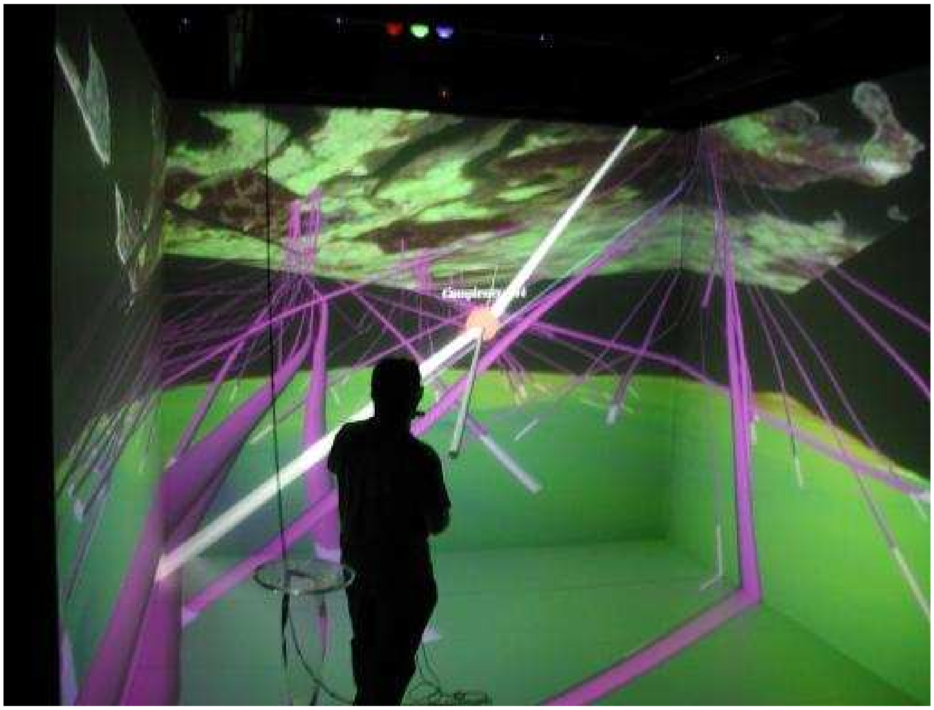

Figure 3.12: Photograph of a IDP user interacting with the virtual world using the ray-casting technique. A virtual ray extends from the wand and virtual objects can be moved and rotated by when intersected by the ray.

3.2.3 Interactive Objects

The IDP scene graph consisted of a number of interactive objects. Only two types of these interactive objects, *pull points* and *well sliders*, were active during this study.

A pull point is represented by a sphere pierced by a three-dimensional arrow (see figure 3.13). A pull point has six degrees of freedom, as it can be dragged to any position in the field and oriented in any direction. The pull point object is used to define the position and attitude for a point in a well path (see the pull point method discussion in Section 3.1.1). The center of the pull point defines the position, while the three-dimensional arrow defines the attitude.

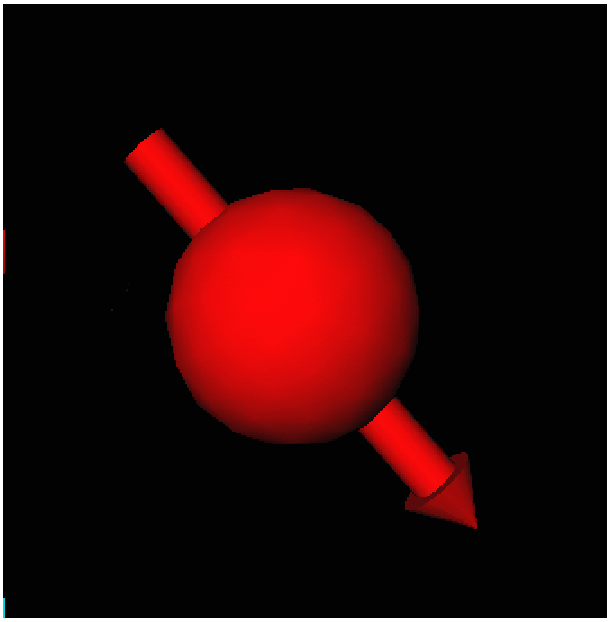

Figure 3.13: Snapshot of a pull point. A pull point defines a position and attitude for a well path. The center of the sphere defines the position, while the arrow defines the attitude.

A well slider is represented by a sphere. The well slider has one degree of freedom, as it can only be moved along a well path (see figure 3.15). The well slider has many uses in the IDP. In context of this study the two well sliders are used to define the edit region in the pull point edit method (discussed in section 3.1.1). One slider marks the start of the edit region, while the other slider marks the end of the edit region.

Chapter 4

The Experiment

The planning of a new well path through the existing wells of a mature oilfield is a real-world task that requires spatial understanding of a complex three-dimensional environment and the precise placement of objects within that environment. The immersive drilling planner (IDP) is capable of visualizing a mature oilfield and editing a new path within that oilfield, on both a desktop environment and in immersive virtual environment (IVE). Although the user interface is different in the two environments, the scene and the dynamics of the scene are identical. This provides a testbed that can be used to evaluate the added value of immersion on a spatially complex real-world problem. This chapter describes an experiment designed to compare an IVE with a stereoscopic desktop environment in the performance and correctness of a well path editing task.

4.1 Participants

Nineteen unpaid participants were recruited from the staff and students at the University of Colorado at Boulder, employees at Raytheon Systems Corporation, and employees at Seraut Inc, a local software firm. The participants received no tangible benefit from participation in the study. Two participants could not complete the experiment due to hardware failures; the data from these two incomplete runs are not included in the results. Participants were organized into counterbalanced experimental blocks of four. After disregarding the two incomplete runs, the remaining seventeen participants complete four experimental blocks. The fifth experimental block is incomplete, containing only the last run, and has been excluded. Demographics of the remaining sixteen participants are presented in table 4.1.

Subject ID	Handedness	Gender	IVE Experience	Weekly Computer Usage (hrs)
s00	Right	Female	Little	50
s01	Left	Male	None	40
s02	Right	Male	None	60
s03	Right	Male	None	50
s04	Right	Male	None	70
s05	Right	Male	None	50
s06	Right	Male	None	30
s07	Right	Male	None	50
s08	Right	Male	None	25
s09	Left	Male	None	50
s10	Right	Male	None	20
s11	Right	Female	None	20
s12	Right	Male	None	60
s13	Right	Male	Experienced	40
s14	Right	Male	Experienced	50
s15	Right	Female	None	5

Table 4.1: Spreadsheet of participant demographics.

4.2 Apparatus

The IVE used for this study was a 12' x12' x10' Mechdyne MD Flex™ located on the University of Colorado campus at the B.P. Center for Visualization. The IVE consisted of four screens each with a resolution of 1024x768, an InterSense VET 900 tracker, a tracked six degree of freedom wand, and a tracked pair of

CrystalEyes active stereo LCD shuttered glasses. The IVE is described in more detail in section 2.1.1.

The desktop used for this experiment was located on the University of Colorado campus at the B.P. Center for Visualization. The desktop equipment consisted of a 21 inch SGI monitor, a three-button mouse, an SGI keyboard, and a untracked pair of CrystalEyes^{TM} active stereo glasses. The images on the monitor, like the images on the screens in the immersive treatment, were stereoscopic and constrained to a resolution of 1024x768. The desktop equipment is described in detail in Chapter 2.2.

4.3 Experimental Design

The experiment consisted of four separate logged experimental tasks (denoted Task01, Task02, Task03, and Task04) and a training task (denoted Task00). Each participant performed the training task and two experimental tasks on the desktop and the training task and the two experimental tasks in the IVE. Participants were given a time limit of ten minutes to complete each task. The runs were counterbalanced in four run experimental blocks to adjust for learning effects (see Table 4.2).

	1st Treatment				2nd Treatment			
Subject ID	**Environment**	1st Task	**2nd Task**	**3rd Task**	**Environment**	**4th Task**	**5th Task**	**6th Task**
s00	IVE	Task00	Task01	Task02	Desktop	Task00	Task03	Task04
s01	Desktop	Task00	Task01	Task02	IVE	Task00	Task03	Task04
s02	Desktop	Task00	Task03	Task04	IVE	Task00	Task01	Task02
s03	IVE	Task00	Task03	Task04	Desktop	Task00	Task01	Task02
s04	IVE	Task00	Task01	Task02	Desktop	Task00	Task03	Task04
s05	Desktop	Task00	Task01	Task02	IVE	Task00	Task03	Task04
s06	Desktop	Task00	Task03	Task04	IVE	Task00	Task01	Task02
s07	IVE	Task00	Task03	Task04	Desktop	Task00	Task01	Task02
s08	IVE	Task00	Task01	Task02	Desktop	Task00	Task03	Task04
s09	Desktop	Task00	Task01	Task02	IVE	Task00	Task03	Task04
s10	Desktop	Task00	Task03	Task04	IVE	Task00	Task01	Task02
s11	IVE	Task00	Task03	Task04	Desktop	Task00	Task01	Task02
s12	IVE	Task00	Task01	Task02	Desktop	Task00	Task03	Task04
s13	Desktop	Task00	Task01	Task02	IVE	Task00	Task03	Task04
s14	Desktop	Task00	Task03	Task04	IVE	Task00	Task01	Task02
s15	IVE	Task00	Task03	Task04	Desktop	Task00	Task01	Task02

Table 4.2: Spreadsheet of experimental design. Treatments and tasks were counterbalanced to adjust for possible learning effects. Participants were grouped into one of four experimental blocks.

The independent variable was the environment: the head-tracked stereoscopic IVE verses the stereoscopic desktop environment. The dependent variables were the time to complete the task, the correctness of the final well path, and an evaluation of the degree of cybersickness experienced by the user.

4.3.1 Tasks

The experimental tasks in this study involved editing the path of a new well in a mature field. The same dataset was used to construct the virtual mature field (see figure 4.1) for all the experimental tasks in this study. Ninety well logs were used to construct the corresponding ninety well path uncertainty surfaces. A Landsat image of the field was rendered above these uncertainty surfaces. A roughly horizontal surface, representing a geological property of the field' s reservoir, was rendered toward the lower extents of the uncertainty surfaces.[1]

---

[1] The dataset used to construct the experimental task, including well logs, Landsat imagery, and geological horizons, were donated by British Petroleum.

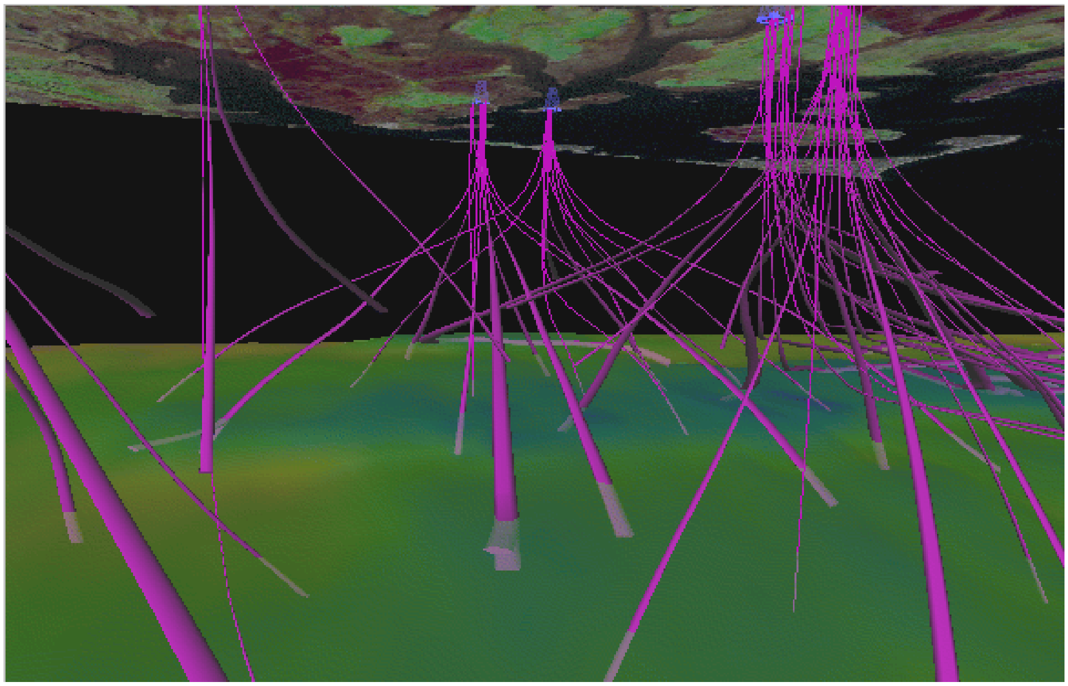

Figure 4.1: Snapshot of the mature oilfield dataset used for all experimental tasks in the study. The virtual oilfield was constructed from a well log dataset donated by British Petroleum, consisting of ninety well logs. Each well log was used to construct an uncertainty surface for the well, shown in purple.

The objective of each task was to edit the new path so that its uncertainty surface did not intersect the uncertainty surface of any existing well while not exceeding a goal complexity value. The path of the new well was edited using the pull point method (see sections 3.1.1 and 3.2) which allows the participants to edit a region of the well. The participants could define an edit region by dragging two well sliders up and down the original path of the new well. The participant could then change the path within the edit region by moving or rotating the pull point. As the pull point is moved, the edited path' s white uncertainty surface is updated in real time (see figures 4.2 and 4.3).

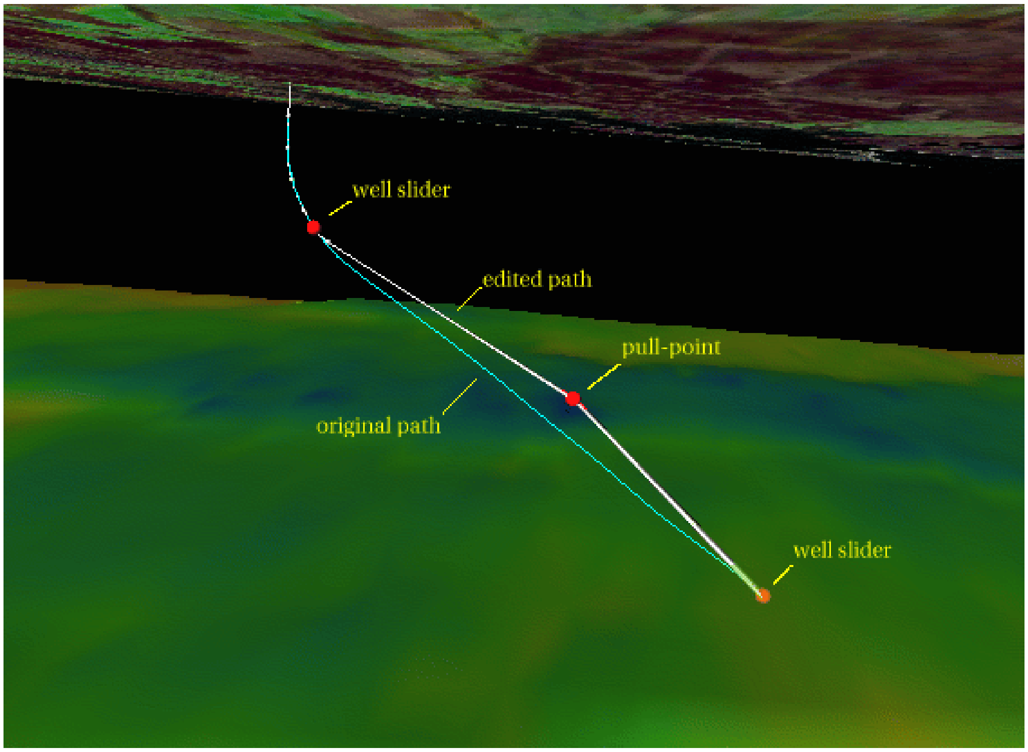

Figure 4.2: Annotated illustration of an editable well. An editable well consists of an original path (cyan), the uncertainty surface of the edited path (white), two well sliders (topmost and bottommost red spheres), and a pull point (center red sphere). The edited path could be modified by interacting with any of the red spheres. The two well sliders defined the region of the original well to be edited. The pull point defined a position and orientation that the edited path had to pass through.

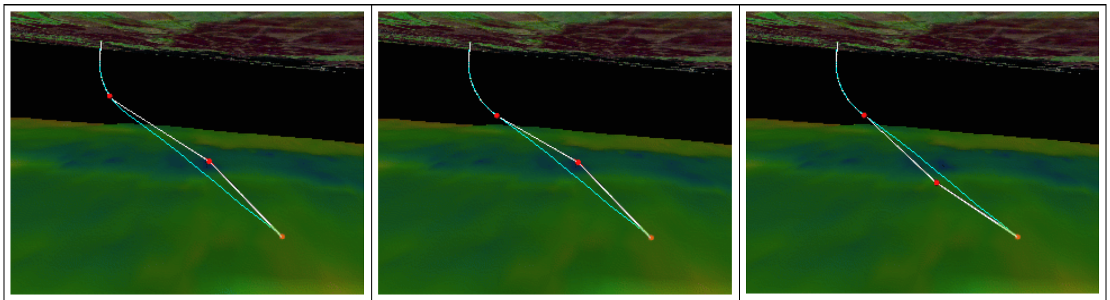

Figure 4.3: Three snapshots of a well edit. The topmost well slider is dragged down the original path between frames 1 and 2. The position of the pull point is moved between frames 2 and 3.

The test application begins by presenting the participant with a two-dimensional start dialog. This dialog provides the time allotted for the task, the goal complexity of the new well, and a start button (see figure 4.4). When the start button is pressed, the dialog is closed and the test application begins a timed log of the user' s actions. All changes to the user' s viewpoint (i.e., head and camera motion) and all interactions (i.e., mouse and wand movements and button presses) are logged. The log can be played back allowing the participant' s actions and the final position of the new well to be scrutinized after the test. Once the allotted time has been reached the test application terminates.

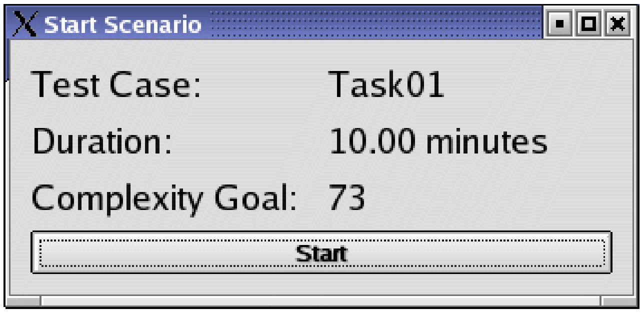

Test Case: Task01
Duration: 10.00 minutes
Complexity Goal: 73

Figure 4.4: Snapshot of the start dialog presented to participant at the beginning of each task. The dialog is closed and the IDP testbed begins logging user interactions when the user presses the start button.

The participant begins at a fixed starting position outside of the virtual field, then navigates through the field to the new well. Then, through a series of well slider and pull point movements, the participant can edit the path of the new well. A three-dimensional text readout above the pull point provides the user with complexity value feedback. Once the participant believes that the new path' s uncertainty surface does not intersect the uncertainty surface of any existing well and that the new path has a complexity value at or below the goal complexity, the task is complete and the test is ended.

4.3.1.1 Training Task (Task00)

In addition to the experimental tasks, a training task was designed. The training task consisted of a grossly simplified field with a single existing well path intersected by the new well path (see figure 4.5). The training task was provided to the participants as a medium for learning the interfaces of the two environments and for exploring the dynamics of well path manipulation. A separate training task was provided for each of the two environments. Participants participated in the training task prior to the experimental tasks.

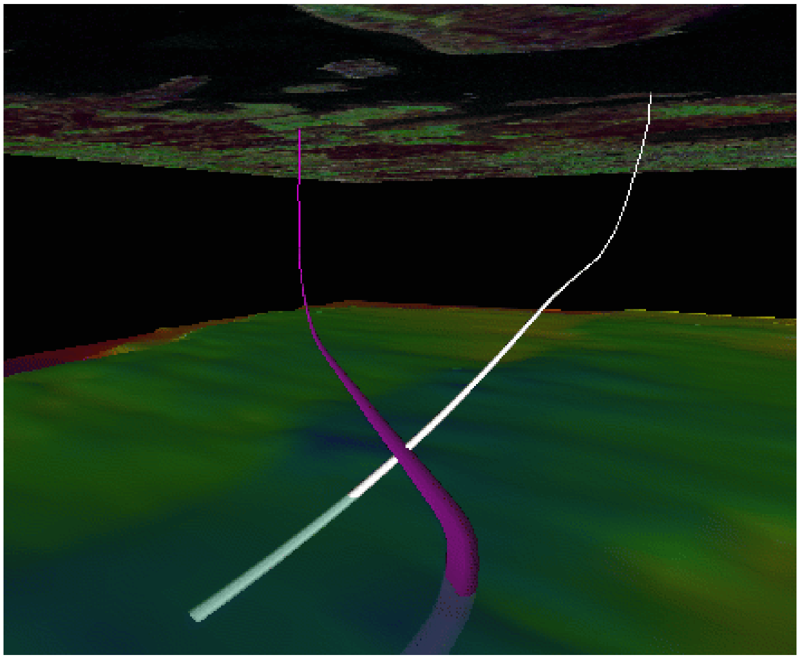

Figure 4.5: Snapshot of the training task, Task00. The training task consisted of an existing well, shown in purple, and one editable well, shown in white. The training was run in both environments, prior to the experimental tasks in that environment.

4.3.1.2 Experimental Tasks (Task01, Task02, Task03, and Task04)

The mature field dataset used for all four logged tasks was identical; the four tasks varied only in the layout and position of the new well and the goal complexity value.

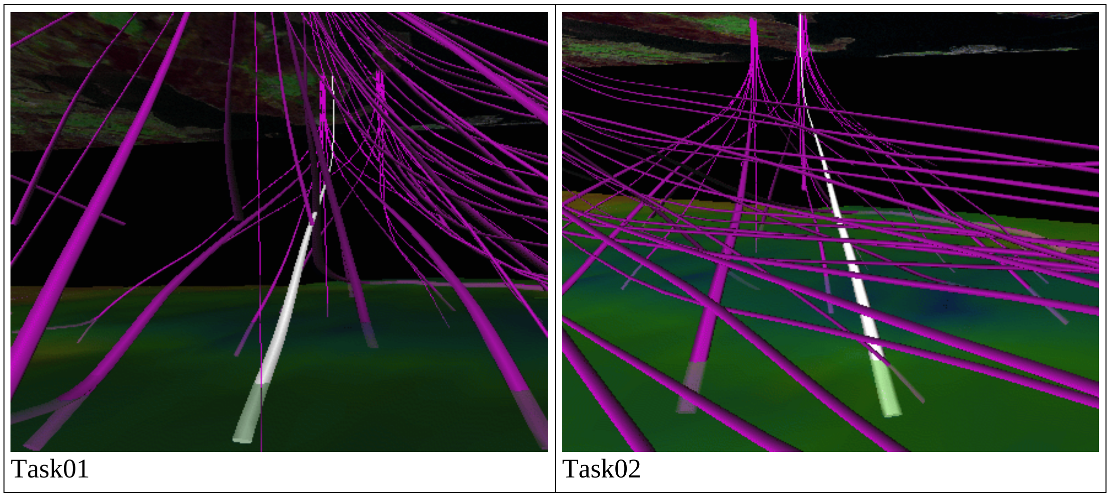

Task01 Task02

Figure 4.6: Snapshots of Task01 and Task02. Existing well uncertainty surfaces are shown in purple; the new editable well uncertainty surface is shown in white. Participants were asked to move the editable well so that the uncertainty surface of the editable well did not intersect the uncertainty surface of any existing well. Task01 and Task02 were always run as a group in the same environment. Task01 was always run before Task02.

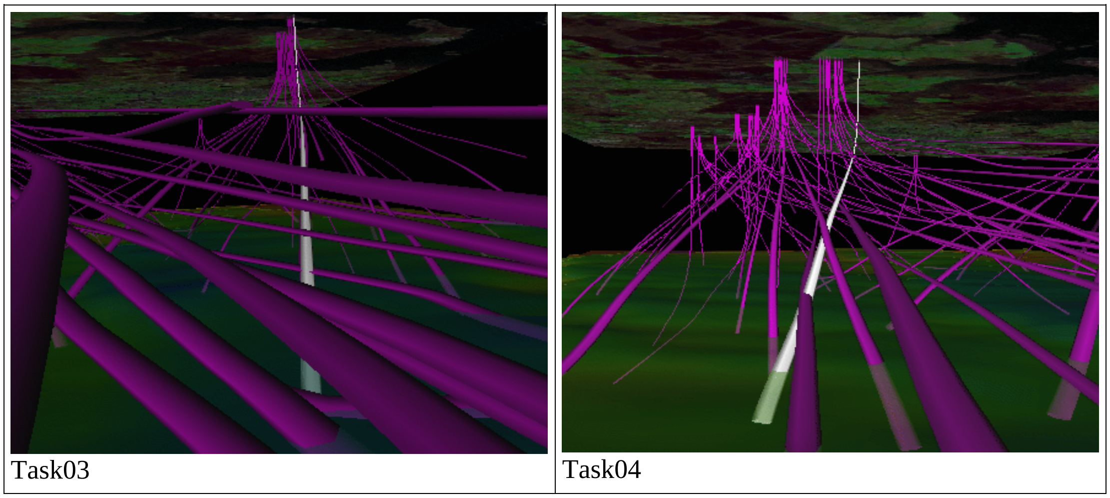

Task03 Task04

Figure 4.7: Snapshots of Task03 and Task04. Existing well uncertainty surfaces are shown in purple; the new editable well uncertainty surface is shown in white. Participants were asked to move the editable well so that the uncertainty surface of the editable well did not intersect the uncertainty surface of any existing well. Task03 and Task04 were always run as a group in the same environment. Task03 was always run before Task04.

4.3.2 Performance Measures

The IDP maintians a timed log of the participant' s interactions with the virtual environment; the time to complete the task can be derived from the log. The final well path can be reconstructed from the log to evaluate the correctness of the participant' s solution. Any final well path whose uncertainty surface did not intersect with any existing uncertainty surface and whose complexity did not exceed the task' s goal complexity value was considered to be correct.

4.3.2.1 Cybersickness

An evaluation of cybersickness is a secondary result of this study. Some IVE users experience symptoms that parallel those of classical motion sickness. This type of sickness, *cybersickness* or *simulator sickness,* is different from motion sickness in that the user is stationary but has a sense of motion through moving visual imagery. The causes are not completely known; however, *sensory conflict theory* is the most widely accepted explanation. Sensory conflict theory holds that inconsistent sensory information about one' s motion and orientation can cause ill effects. That is, images projected in the IVE can be inconsistent with the orientation and motion detected by the user' s inner ear [LaViola 2000].

Common symptoms of cybersickness can include: fatigue, eyestrain, blurred vision, headache, pallor (paleness of skin), sweating, dryness of mouth, disorientation, vertigo (dizziness), and ataxia (lack of coordination). Less common symptoms of cybersickness include nausea and vomiting [LaViola 2000]. Symptoms of cybersickness are a common phenomenon; published estimates suggest that as many as 60% of users experience some adverse effects in a virtual environment, and as many as 20% experience moderate to severe dizziness and nausea [Potel 1998]. There are published accounts of the symptoms lasting several hours after exposure. However, most published cybersickness studies are not representative of typical industrial and academic use of an IVE [Potel 1998, Lewis 1997].

Currently, most published cybersickness data originates from one of two types of studies: military simulator studies and IVE experiments designed to elicit and isolate the causes of cybersickness [So 1999, Takahashi, Kennedy 1997]. In the military experiments, participants are immersed in a dynamic, motion-intensive virtual environment, often remaining immersed for several hours. This is not representative of most academic and industrial IVE use. Presumably, a study designed to isolate the causes of cybersickness would have a higher incident rate of cybersickness than would a study on the casual use of an IVE.

To measure cybersickness in this study, the Simulator Sickness Questionnaire (SSQ), developed by Kennedy et al. [Kennedy 1993], was used. The SSQ is used as the standard measure of simulator sickness in many virtual environment studies. It breaks cybersickness into three components: nausea (nausea, stomach awareness, increased salivation, and burping), oculomotor (eyestrain, difficulty focusing, blurred vision, and headache), and disorientation (dizziness and vertigo). The components can be combined to compute a total SSQ score. The SSQ was administered to participants in this study immediately before and after the immersive treatment and before and after the desktop treatment. Participants were asked report the degree to which they experience each cybersickness symptom as: none, slight, moderate, or severe. A copy of the SSQ can be found in Appendix A, *Experimental Scripts and Questionnaires*.

4.4 Experimental Procedure

The experimental procedure[2] was conducted individually, one participant at a time. Participants were greeted at the B.P. Center for Visualization and given a brief tour of the facilities and a brief explanation of the experiment. Participants were then asked to read and sign the *Subject Informed Consent Form,* and fill out a SSQ. The initial SSQ provided a cybersickness baseline. Depending on the participant' s position in the experimental block, the participant would sit at the desktop or enter the IVE. While the experimenter read from a script explaining the environment' s interface and the objective of the tasks, the participant explored the training task. The participant was encouraged to explore the environment' s interface and the dynamics of the well path editing until they felt comfortable or until the ten minute time limit was reached. After completing the training task, the participants then performed the two logged experimental tasks as assigned per their position in the experimental block. Then, after completing a second SSQ, the participant would perform the training task and two logged experimental tasks on the other environment. Again, while performing the training task, the experimenter would read from a script

---

[2] The experimental protocol was approved by an expedited review by the University of Colorado Human Research Committee, under protocol #1202.17 "Immersive Path Planning."

describing the user interface in that environment. After completing the second treatment, the participant was then asked to complete a final SSQ and a post-experiment questionnaire. A copy of the experimental scripts and questionnaires can be found in Appendix A, *Experimental Scripts and Questionnaires.*

4.5 Complications

The first eight participants were run without incident. However between the eighth and ninth participants, the machine room at the B.P. Center for Visualization suffered a minor flood. The flood lead to the failure and replacement of several low-level hardware components and the eventual upgrade of the SGI Origin 3800 computer' s operating system. After the flood, the tracker exhibited a higher level of noise in its reporting of both position and orientation. The additional noise was slight, but perceivable by users of the IVE. The frame rate of the system was also slightly reduced. The average frame rate during an immersive test before the flood averaged 45 frames per second, while averaging only 40 frames per second after the hardware and software repairs. The reduction in frame rate was not immediately perceivable and was only deduced by examining the test logs. The system repairs did not have any apparent effect on the desktop version of the IDP.

4.6 Pilot Testing

Before conducting the experiments, a series of pilot tests were run. Three graduate students at the B.P. Center for Visualization who were experienced in both IVE use and well planning were used as pilot testers. The pilot tests served several purposes, including the refinement of the experimental procedure so it would run smoothly for all participants, demonstration that the experimental tasks were not too simple or too difficult, and demonstration that the tasks were of approximately equal difficultly. The pilot tests were also used to demonstrate the effectiveness of the tasks to isolate the differences between users and experimental treatments.

The original concept for the experimental tasks was a single well whose path would need to be moved to avoid a collision with one or two geological hazards, such as a salt dome (see figure 4.10). Early pilot tests suggested that collision avoidance with a one or two large geological hazards was relatively easy, and there was only a very slight difference between the two treatments. After a series of additional pilot tests, collision avoidance with existing wells in a mature field was found to show a much more substantial difference between the two treatments.

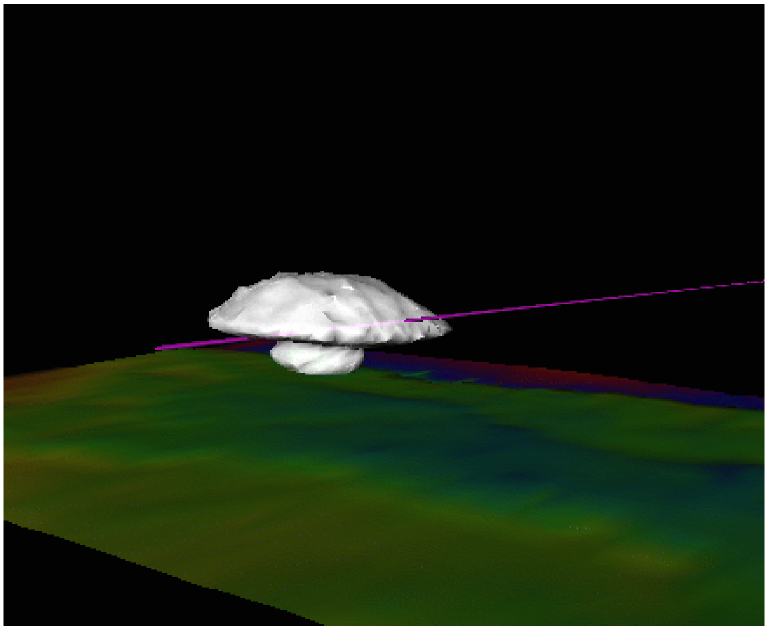

Figure 4.8: Snapshot from an early pilot test depicting a well intersecting a model of a salt dome.

Chapter 5

Results

The data from the four completed experimental blocks, consisting of sixteen participant runs, are presented here.

5.1 Objective Measures

As described at length in the previous chapter, each participant was asked to plan the path of four oil wells. Two well paths were planned on the desktop workstation with a stereoscopic display, and two well paths were planned in the IVE. The objective of each task was to edit the new path so that the uncertainty surface was not intersecting with the uncertainty surface of any existing well, while not exceeding a goal complexity value. Each task was timed and limited to a maximum of ten minutes. The solution of each task was evaluated after the experiment for correctness. The solution was deemed correct if, and only if, the new well path' s complexity value was at or below the goal complexity value and the uncertainty surface of the new well did not intersect with any uncertainty surfaces of existing wells. Table 5.1 presents the data for Task01 and Task02. Table 5.2 presents the data for Task03 and Task04.

Comparing the number of correct solutions within the participants shows a significant difference between the two environments (see Figure 5.1). Of the sixteen participants, nine had more correct solutions in the IVE, one had more correct solutions in the desktop environment, and six had the same number of correct solutions in the two environments. The sign test shows a statistically significant difference at the 0.05 significance level.

Comparing the total solution time taken to complete two tasks in the IVE with the two tasks in the desktop environment provides a more significant result (see Figure 5.2). Of the sixteen participants, only one participant[3] took more time in the IVE. The sign test shows this to be statistically significant at the 0.001 significance level.

---

[3] The one participant that completed the desktop tasks more quickly indicated that he was an expert user of Open Inventor viewer interfaces.

Subject ID	Task01				Taks02			
	Time (s)	Correct	Environment	Treatment	Time (s)	Correct	Environment	Treatment
s00	353	1	ive	1st	313	1	ive	1st
s01	266	1	desktop	1st	600	0	desktop	1st
s02	434	0	desktop	2nd	306	1	desktop	2nd
s03	600	0	ive	2nd	330	1	ive	2nd
s04	281	1	ive	1st	249	0	ive	1st
s05	460	1	desktop	1st	572	1	desktop	1st
s06	333	0	desktop	2nd	277	0	desktop	2nd
s07	360	1	ive	2nd	574	1	ive	2nd
s08	336	1	ive	1st	347	1	ive	1st
s09	375	1	desktop	1st	222	0	desktop	1st
s10	423	1	desktop	2nd	206	0	desktop	2nd
s11	142	1	ive	2nd	247	1	ive	2nd
s12	127	1	ive	1st	383	1	ive	1st
s13	445	1	desktop	1st	262	0	desktop	1st
s14	243	1	desktop	2nd	270	1	desktop	2nd
s15	221	1	ive	2nd	377	1	ive	2nd

Table 5.1: Spreadsheet of data from Task01 and Task02. The *Time* is the number of seconds taken by the participant to complete the task. The *Correct* field indicates whether or not the final solution was correct (1 = correct, 0 = incorrect). The *Environment* field indicates the environment in which the task was performed. The *Treatment* field indicates whether the given environment was the first or second treatment in the run. Task01 was always performed before Task02. Both tasks were limited to a maximum of ten minutes.

Subject ID	Task03				Task04			
	Time (s)	Correct	Environment	Treatment	Time (s)	Correct	Environment	Treatment
s00	600	0	desktop	2nd	196	1	desktop	2nd
s01	352	1	ive	2nd	191	1	ive	2nd
s02	327	1	ive	1st	398	1	ive	1st
s03	600	0	desktop	1st	520	1	desktop	1st
s04	452	1	desktop	2nd	323	1	desktop	2nd
s05	499	1	ive	2nd	491	1	ive	2nd
s06	333	1	ive	1st	208	0	ive	1st
s07	600	0	desktop	1st	600	0	desktop	1st
s08	594	1	desktop	2nd	148	1	desktop	2nd
s09	62	1	ive	2nd	135	0	ive	2nd
s10	206	1	ive	1st	385	1	ive	1st
s11	457	1	desktop	1st	214	0	desktop	1st
s12	383	1	desktop	2nd	376	0	desktop	2nd
s13	183	0	ive	2nd	154	1	ive	2nd
s14	126	1	ive	1st	407	1	ive	1st
s15	349	0	desktop	1st	600	0	desktop	1st

Table 5.2: Spreadsheet of data from Task03 and Task04. The *Time* is the number of seconds taken by the participant to complete the task. The *Correct* field indicates whether or not the final solution was correct (1 = correct, 0 = incorrect). The *Environment* field indicates the environment in which the task was performed. The *Treatment* field indicates whether the given environment was the first or second treatment in the run. Task03 was always performed before Task04. Both tasks were limited to a maximum of ten minutes.

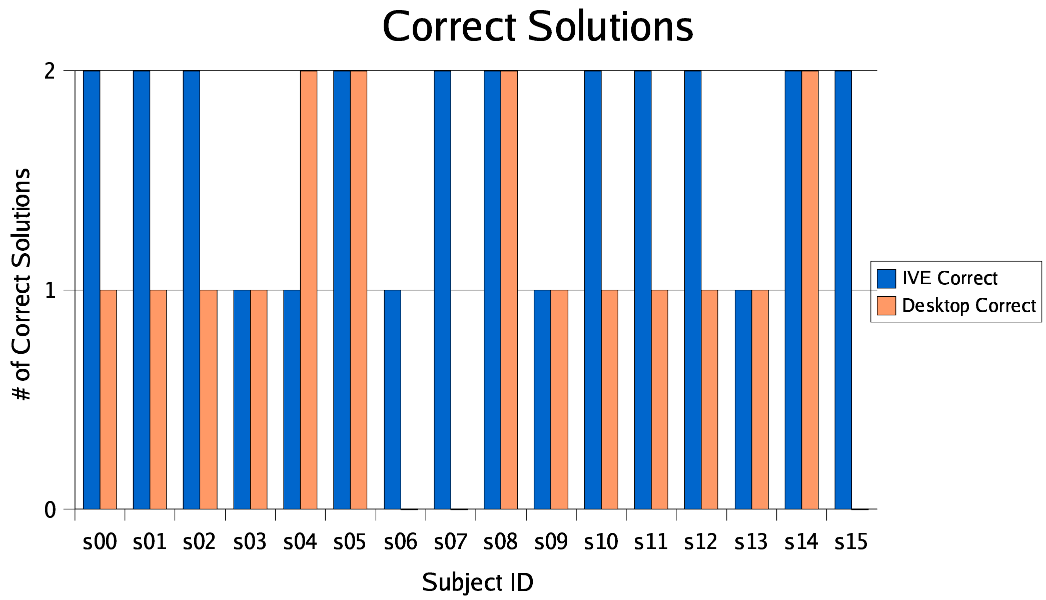

Figure 5.1: Graph illustrating the number of correct solutions for each participant in each environment. Nine participants had more correct solutions in the IVE, one participant had more correct solutions on the desktop, and six participants had an equal number of correct solutions in the two environments. The sign test indicates this is a significant result ($p < 0.05$).

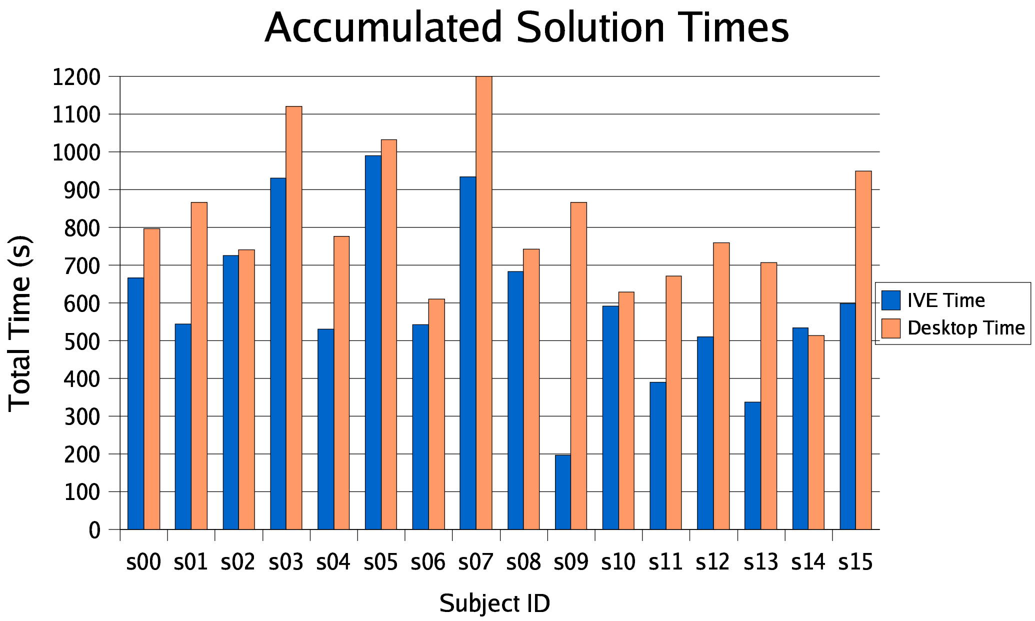

Figure 5.2: Graph illustrating the total accumulated solution time for each participant in each environment. Fifteen participants took more time to complete two tasks on the desktop, and only one participant took more time to complete the two IVE tasks. The sign test indicates this is a significant result ($p < 0.001$).

5.1.1 Analysis of Variance of Time, Correctness, Environment, and Order

Results were analyzed using a repeated measures ANOVA, with two between-subjects factors. The analysis of variance of solution time is presented in Table 5.3, and the analysis of variance of correctness is presented in Table 5.4. The repeated measures were the two tasks in each environment. The treatment order and the environment were used as the between subject factors. There was a significant effect of the environment in both solution time, F(1,28) = 6.468, p = 0.017, and correctness, F(1,28) = 7.986, p = 0.009. The analysis revealed no significance for treatment order, or an interaction between treatment order and environment.

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Intercept	7842100.141	1	7842100.141	358.528	0.000
Environment	141470.016	1	141470.016	6.468	0.017
Treatment Order	6142.641	1	6142.641	0.281	0.600
Environment X Treatment Order	13894.516	1	13894.516	0.635	0.432
Error	612445.188	28	21873.042		

Table 5.3: Spreadsheet of results from a repeated measures ANOVA of solution time, with between-subject factors environment and treatment order.

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Intercept	31.641	1	31.641	199.648	0.000
Environment	1.266	1	1.266	7.986	0.009
Treatment Order	0.141	1	0.141	0.887	0.354
Environment X Treatment Order	0.016	1	0.016	0.099	0.756
Error	4.438	28	0.158		

Table 5.4: Spreadsheet of results of a repeated measures AVOVA of correctness, with between-subject factors environment and treatment order.

5.1.2 Analysis of Variance of Experimental Blocks.

Results were also analyzed by a repeated measures ANOVA on the experimental blocks. Each experimental block consisted of four participants, and contains a complete counterbalance of task and treatment orders. A repeated measures

ANOVA was conducted treating each experimental block as a “super-subject.” The analysis of variance for solution time is presented in Table 5.5, and the analysis of variance for correctness is presented in Table 5.6. The effect for the environment was significant for both solution time, F(1,3) = 103.406, p = 0.002, and solution correctness, F(1,3) = 75.0, p = 0.003. There was an indication of a treatment order effect, F(1,3) = 7.686, p = 0.69, but the effect is not significant at the 0.05 level. There were no significant interactions between task, treatment order, and the environment.

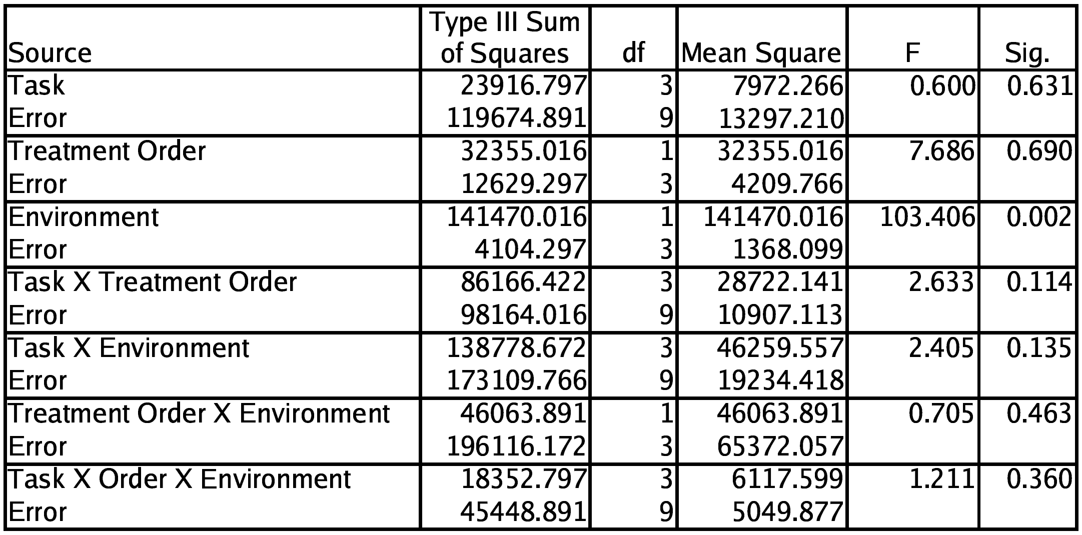

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Task	23916.797	3	7972.266	0.600	0.631
Error	119674.891	9	13297.210		
Treatment Order	32355.016	1	32355.016	7.686	0.690
Error	12629.297	3	4209.766		
Environment	141470.016	1	141470.016	103.406	0.002
Error	4104.297	3	1368.099		
Task X Treatment Order	86166.422	3	28722.141	2.633	0.114
Error	98164.016	9	10907.113		
Task X Environment	138778.672	3	46259.557	2.405	0.135
Error	173109.766	9	19234.418		
Treatment Order X Environment	46063.891	1	46063.891	0.705	0.463
Error	196116.172	3	65372.057		
Task X Order X Environment	18352.797	3	6117.599	1.211	0.360
Error	45448.891	9	5049.877		

Table 5.5: Spreadsheet of within-subject effects for repeated measures ANOVA of time for experimental blocks.

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Task	0.313	3	0.014	0.517	0.681
Error	1.812	9	0.201		
Treatment Order	0.250	1	0.250	1.200	0.353
Error	0.625	3	0.208		
Environment	1.563	1	1.563	75.000	0.003
Error	0.006	3	0.002		
Task X Treatment Order	1.375	3	0.458	3.300	0.720
Error	1.250	9	0.139		
Task X Environment	0.563	3	0.188	1.286	0.337
Error	1.313	9	0.146		
Treatment Order X Environment	0.250	1	0.250	2.000	0.252
Error	0.375	3	0.125		
Task X Order X Environment	0.625	3	0.208	0.833	0.509
Error	2.250	9	0.250		

Table 5.6: Spreadsheet of within-subjects effects for repeated measure ANOVA of correctness for experimental blocks.

5.1.3 Task Analysis

An analysis of mean solution times and number of correct solutions per task illustrate differences between the tasks (see Figures 5.3 and 5.4). On average, the solution time in the IVE was approximately 23% faster than in the desktop environment for Task01. The number of correct solutions for Task01 were similar in the two environments, with seven correct solutions on IVE and six correct solutions in the desktop environment. The mean solution times for Task02 were also nearly equal, with the desktop just 4% faster than the IVE. However, the shorter mean solution time on the desktop for Task02 was offset by a decrease in correctness. Only three correct solutions were found on the desktop for Task02 compared to seven correct solutions in the IVE. Task03 had the largest difference in mean solution times between the two environments. On average, the Task03 solutions were found approximately 93% faster in the IVE than in the desktop environment. The increased speed in the IVE did not correspond to a decrease in correctness. There were seven correct Task03 solutions in the IVE and only four correct Task03 solutions in the desktop environment. On average, the solution time in the IVE was approximately 26% faster than in the desktop environment for Task04. There were six correct solutions on IVE and four correct solutions in the desktop environment.

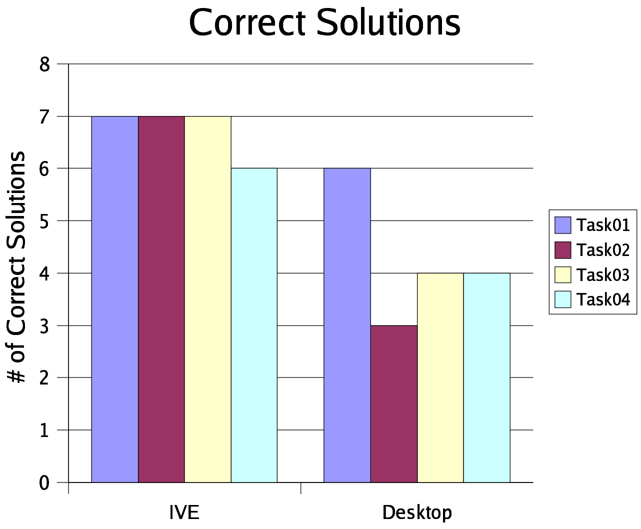

Figure 5.3: Graph of the number of correct solutions by task.

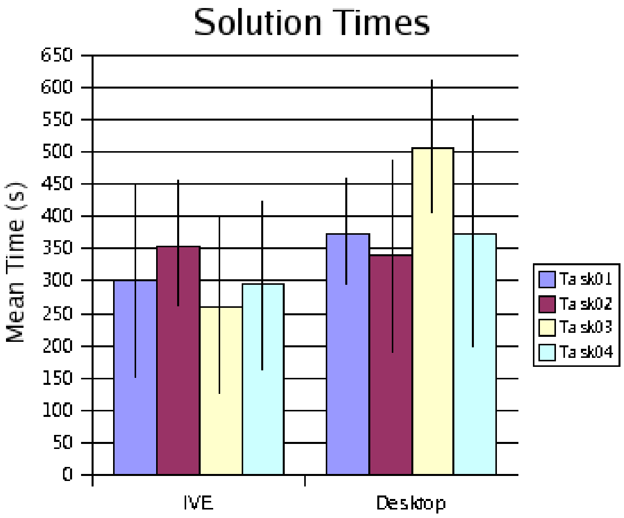

Figure 5.4: Graph of the mean solution time by task. Error bars show standard deviation.

Comparing the number of correct solutions and the mean solution times between treatment order does not show any significant learning effects (see figures 5.5, 5.6, 5.7, and 5.8).

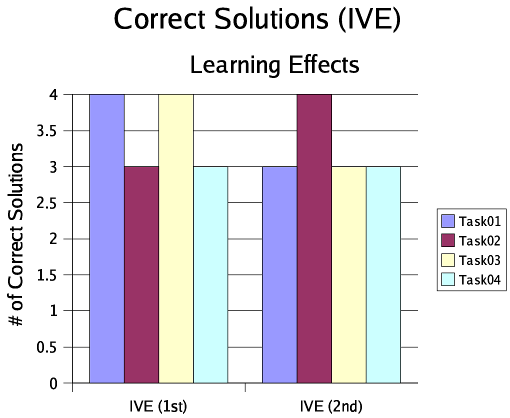

Figure 5.5: Graph of the comparison of correct solutions in the IVE between first and second treatments.

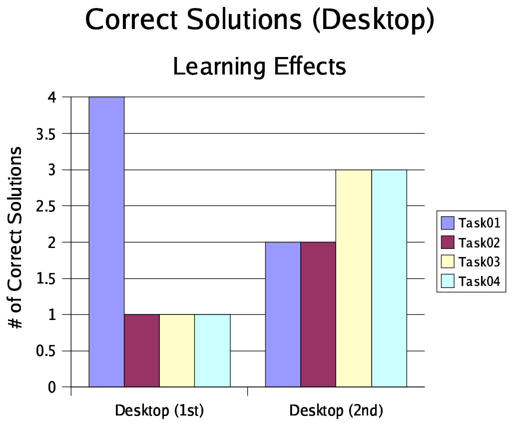

Figure 5.6: Graph of the comparison of correct solutions in the desktop environment between first and second treatments.

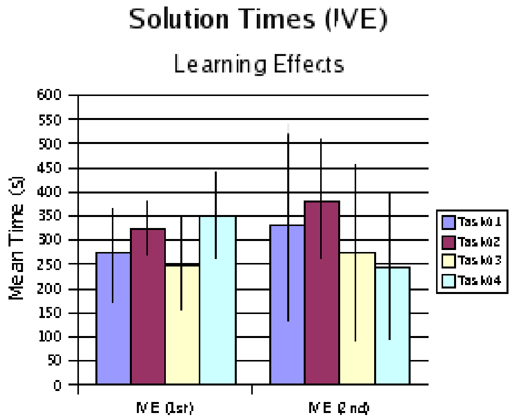

Figure 5.7: Graph of the comparison of mean solution times in the IVE between first and second treatments. Error bars represent standard deviation.

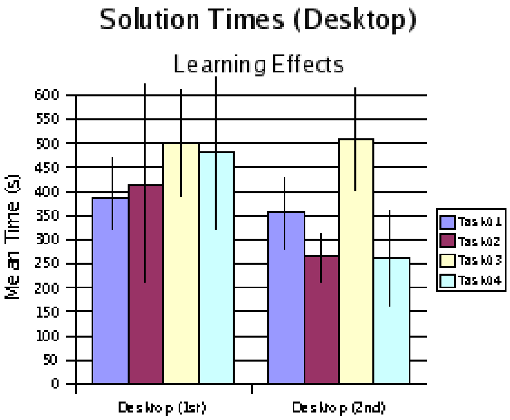

Figure 5.8: Graph of the comparison of solution times in the desktop environment between first and second treatments. Error bars represent standard deviation.

5.1.3.1 Analysis of Variance of Task Times and Correctness

Results from individual tasks were analyzed using a one-way ANOVAs, with the environment and treatment order as between-group factors. In the analysis of task solution times, the effect for the environment was significant only for Task03, F(1, 15) = 15.027, p = 0.002. In the analysis of task correctness, the effect for the environment was only significant for Task02, F(1,15) = 5.091, p = 0.041. There were no significant effects for task order. The effects for the environment are presented in Table 5.7 and Table 5.8. The effects for the task order are presented in Table 5.9 and Table 5.10.

		Sum of Squares	df	Mean Square	F	Sig.
Solution time for Task01	Between Groups	19530.063	1	19530.063	1.314	0.271
	Within Groups	208073.875	14	14862.420		
	Total	227603.938	15			
Solution time for Task02	Between Groups	689.063	1	689.063	0.040	0.845
	Within Groups	243941.875	14	17424.420		
	Total	244630.938	15			
Solution time for Task03	Between Groups	236925.563	1	236925.563	15.027	0.002
	Within Groups	220725.875	14	15766.134		
	Total	457651.438	15			
Solution time for Task04	Between Groups	23104.000	1	23104.000	0.879	0.364
	Within Groups	367929.750	14	26280.696		
	Total	391033.750	15			

Table 5.7: Spreadsheet of results from one-way ANOVAs of solution time for each task with the environment as the between-subjects factor. Only Task03 shows a significant effect.

		Sum of Squares	df	Mean Square	F	Sig.
Correct solutions for Task01	Between Groups	0.000	1	0.000	0.000	1.000
	Within Groups	1.750	14	0.125		
	Total	1.750	15			
Correct solutions for Task02	Between Groups	1.000	1	1.000	5.091	0.041
	Within Groups	2.750	14	0.196		
	Total	3.750	15			
Correct solutions for Task03	Between Groups	0.563	1	0.563	2.739	0.120
	Within Groups	2.875	14	0.205		
	Total	3.438	15			
Correct solutions for Task04	Between Groups	0.250	1	0.250	1.000	0.334
	Within Groups	3.500	14	0.250		
	Total	3.750	15			

Table 5.8: Spreadsheet of results from one-way ANOVAs of correctness for each task with the environment as the between-subjects factor. Only Task02 shows a significant effect.

		Sum of Squares	df	Mean Square	F	Sig.
Solution time for Task01	Between Groups	7182.563	1	7182.563	0.456	0.510
	Within Groups	220421.375	14	15744.384		
	Total	227603.938	15			
Solution time for Task02	Between Groups	43368.063	1	43368.063	3.017	0.104
	Within Groups	201262.875	14	14375.920		
	Total	244630.938	15			
Solution time for Task03	Between Groups	410.063	1	410.063	0.013	0.912
	Within Groups	457241.375	14	32660.098		
	Total	457651.438	15			
Solution time for Task04	Between Groups	13456.000	1	13456.000	0.499	0.492
	Within Groups	377577.750	14	26969.839		
	Total	391033.750	15			

Table 5.9: Spreadsheet of results from one-way ANOVAs of solution time for each task with the task order as the between-subjects factor.

		Sum of Squares	df	Mean Square	F	Sig.
Correct solutions for Task01	Between Groups	0.000	1	0.000	0.000	1.000
	Within Groups	1.750	14	0.125		
	Total	1.750	15			
Correct solutions for Task02	Between Groups	0.250	1	0.250	1.000	0.334
	Within Groups	3.500	14	0.250		
	Total	3.750	15			
Correct solutions for Task03	Between Groups	0.563	1	0.563	2.739	0.120
	Within Groups	2.875	14	0.205		
	Total	3.438	15			
Correct solutions for Task04	Between Groups	0.250	1	0.250	1.000	0.334
	Within Groups	3.500	14	0.250		
	Total	3.750	15			

Table 5.10: Spreadsheet of results from one-way ANOVAs of correctness for each task with the task order as the between-subjects factor.

5.2 Subjective Measures

Following the experiment, each participant was asked to fill out a post-experiment questionnaire. A copy of the post-experiment questionnaire can be found in Appendix A, *Experimental Instructions and Questionnaires.* 100% of the participants indicated that they felt the IVE was a more intuitive interface for understanding and interacting with the complex three-dimensional geometry presented in the four tasks. Several participants described having more confidence in the correctness of their solutions in the IVE. Participants' complete comments can be found in Appendix C, *Participants' Remarks.*

5.2.1 Cybersickness Results

To measure cybersickness, the Simulator Sickness Questionnaire (SSQ) was used. The SSQ was administered to participants in this study immediately before and after both the immersive treatment and the desktop treatment. A copy of the SSQ can be found in Appendix A, *Experimental Instructions and Questionnaires*.

Complete SSQ results can be found in Appendix B, *Experimental Data.*

Only 25% of participants reported an increase in cybersickness symptoms after their immersive treatment, and in all cases the severity of the symptoms where reported as "Slight." 38% of participants reported an increase in symptoms after the desktop treatment. The total sickness score following the immersive treatment ranged from 0.0 to 22.44. The total sickness score following the desktop treatment ranged from 0.0 to 18.7. Figure 5.9 shows the mean sickness scores for runs consisting of a immersive treatment followed by a desktop treatment. Figure 5.10 shows the mean sickness scores for runs consisting of a desktop treatment followed by a immersive treatment.

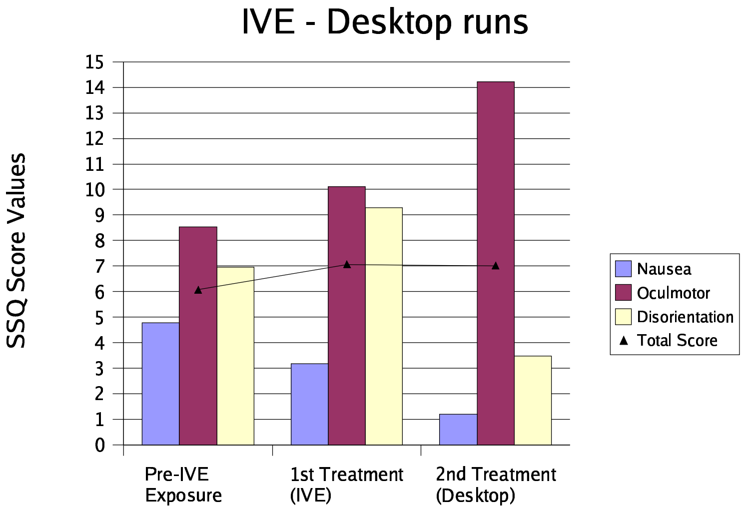

Figure 5.9: Graph of the mean SSQ scores for runs consisting of an immersive treatment followed by a desktop treatment.

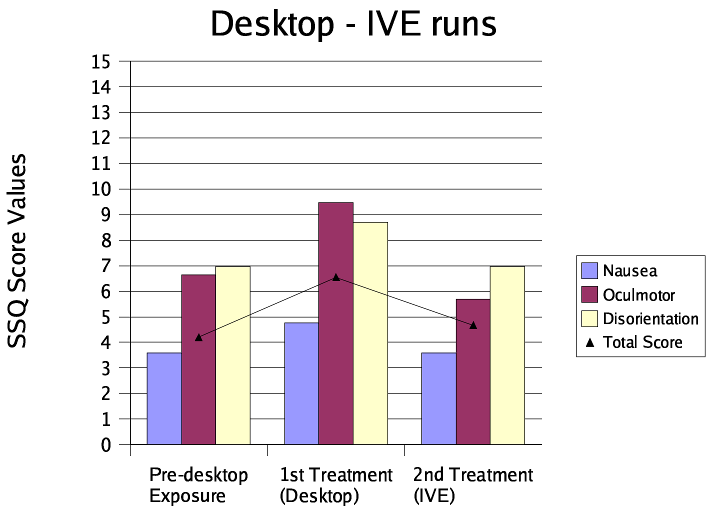

Figure 5.10: Graph of mean SSQ scores for runs consisting of a desktop treatment followed by an immersive treatment.

5.2.1.1 Analysis of Variance of SSQ deltas

The deltas of total SSQ scores between treatments were analyzed using a repeated measures ANOVA. The ANOVA is presented in Table 5.11. There was no significant effect found for the environment, treatment order, or interaction between the environment and treatment order.

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Environment	188.400	1	26.910	0.410	0.544
Error	10.930	7	10.930		
Order	139.440	1	19.920	3.710	0.096
Error	73.870	7	73.870		
Environment X Order	266.860	1	32.410	0.010	0.911
Error	0.440	7	0.440		

Table 5.11: Spreadsheet of results from a repeated measures ANOVA of SSQ total scores deltas between treatments. No significant effects are present.

Chapter 6

Conclusion

The results of this study support the hypothesis that an immersive virtual environment (IVE) allows for faster and more accurate problem solving in a complex interactive spatial domain.

Participants in this study were consistently able to complete well path editing tasks faster in the IVE than in the desktop environment. The total solution time taken by an individual participant to complete two tasks in the IVE was, with one exception, faster than the total solution time taken by the same participant to complete the two tasks in the desktop environment. Fifteen participants had faster solution times in the IVE than in the desktop, leaving a single participant with faster desktop solution times (see Figure 5.2). The sign test indicates this is a statistically significant result
($p < 0.001$).

Participants in this study had more accurate perceptions and judgments in the IVE, as evidenced by the number of correct solutions. Of the sixteen participants, nine participants had more correct solutions in the IVE, one participant had more correct solutions in the desktop environment, and six participants had an equal number of correct solutions in the two environments (see Figure 5.1). The sign test indicates this is a statistically significant result ($p < 0.05$).

Participants' written comments also reflect the added value of immersion. All of the participants indicated that the IVE provided a more intuitive interface for the experimental tasks. Several participants described being more confident in the correctness of their solutions in the IVE.

The incident rate and severity of cybersickness symptoms in this study were far below published accounts. Previous studies have estimated that as many as 60% of IVE users experience cybersickness symptoms [LaViola 2000]. In this study, only 25% of participants indicated an increase of cybersickness symptoms after their immersive treatment. The post-immersive-exposure total sickness scores ranged from 0.0 to 22.44, which is considerably lower than previously published data (19-55) [So 1999, Kennedy 1997]. This discrepancy may be due to the nature of the previous studies. Most published cybersickness data were gathered from either military simulator experiments or experiments specifically designed to induce cybersickness. This study is likely more representative of how immersive technology would be used in most industrial and academic settings. An analysis of variance of the deltas of the total SSQ scores between treatments, shows no significant difference between the increase of cybersickness symptoms in two environments. There is no indication that that cybersickness is any more frequent or severe in the IVE than in the desktop environment (with a stereoscopic display) for this type of task.

6.1 Future Work

The data suggest that IVEs may be more suitable for certain types of problems. Notice in Figures 5.3 and 5.4 that the number of correct solutions and the mean time for Task01 is nearly equivalent for the two environments, while the Task03 solutions have four times more errors in the desktop environment and the mean solution time is significantly slower in the desktop environment. Although the the two tasks were shown to take approximately the same amount of time to solve in the pilot tests, Task01 is less spatially complicated than Task03. A similar phenomenon was observed during the pilot tests. Several initial pilot tests involved spatially simple domains consisting of a few large geological hazards. Moving a well path to avoid the geological hazards did not show any apparent significant differences between the two environments.

These observations imply that the added value of immersion may be correlated to the spatial complexity of the problem. In fact, there may be classes of spatial problems that would benefit from immersion. There have been studies [Boritz 1997, Ruddle 1999] showing that navigation through a three-dimensional world is improved by immersion, but there are no controlled studies which have shown which interactions are improved by immersion. A logical progression of this work would be to identify classes problems that benefit from immersion, by constructing a taxonomy of user interactions that are faster, more precise, and more accurate in an IVE.

6.1.1 Improvements

Although the results of this study indicate the IVE provides a more suitable environment for well editing tasks, the results could likely be further improved through hardware and software enhancements. Several test participants complained of difficulty reading the complexity value in the IVE, and suggested increasing the font size and adding a complexity read-out to the well sliders. Providing users with haptic and/or audio feedback while moving the well might improve results in both environments. The hardware instabilities in the IVE, incurred after the flood, may have had a negative effect on the IVE' s data. Presumably, increasing the frame rate and decreasing the tracker instability would improve the IVE' s usability.

6.2 Summary

This work is one of the first controlled studies designed to evaluate the added value of immersion when interacting with virtual three-dimensional objects. The results of this study indicate that immersive technology can provide an improved interface for solving real-world problems. Not only were the solutions found more quickly in the IVE, but also the solutions were found with far fewer errors. Increasing the speed and accuracy of an industrial problem like oil well planning could save money, time, and potentially lives.

References

Arns, L., Cook, D., Cruz-Neira, C. (1999). The benefits of statistical visualization in an immersive environment. In Proceedings of IEEE Virtual Reality 1999, 88-95.

Barfield, W., Hendrix, C., Bystrom, K. (1997). Visualizing the structure of virutal objects using head tracked stereoscopic displays. Proceedings of the IEEE Virtual Reality Annual International Symposium, 114-120.

Boritz, J., Booth K.S. (1997). A study of interactive 3D point location in a computer simulated virtual environment. Proceedings of the ACM Symposium on Virtual Reality Software and Technology. 181-187.

Bowman, D.A., Koller, D., Hodges, L.F. (1997). Travel in immersive virtual environments: An evaluation of viewpoint motion control techniques. Virtual Reality Annual International Symposium, 1997, 45-52.

Bowman, D.A. (1999). Interaction Techniques for Common Tasks in Immersive Virtual Environments: Design, Evaluation, and Application. Doctoral dissertation, Georgia Institute of Technology.

Bowman, D.A., Kruijff, E., LaViola, J., Poupyrev, I. (2001). An Introduction to 3-D User Interface Design.

Brooks, F.P. (1986). A dynamic graphics system for simulating virtual buildings. Proceedings of the ACM Symposium on Interactive 3D Graphics.

Coninx,K., Van Reeth, F., Flerackers, E. (1997) A hybrid 2D/3D user interface for immersive object modeling. Proceedings of Computer Graphics International 1997, 47-55.

Czernuszenko, M., Pape, D., Sandin, D., DeFanti, T., Dawe, G., and Brown, M. (1997). The immersadesk and infinity wall projection-based virtual reality displays. Computer Graphics 31(2), 46-49.

Cruz-Neira, C., Sandin, D.J., DeFanti, T., Kenyon, R., Hart, J. (1992) The CAVE: Audio visual experience automatic virtual environment, Communications of the ACM 35(6), 65-72.

Forsberg, A.S., Laidlaw, D.H., van Dam, A., Kirkby, R.M., Karniadakis, G.E., Elion, J.L. (2000). Immersive virtual reality for visualizing flow through an artery. Proceedings Visualization 2000, 457-460.

Frohlich., B., Barrass, S., Bjorn, Z., Plate, J., Gobel, M., Exploring geo-scientific data in virtual environments. Proceedings Visualization 1999, 169-173.

Herdon, K., van Dam, A., Gleicher, M. (1994). The challenges of 3D Interaction. SIGCHI Bulletin, 26(4), October, 36-43.

Hyne, N.J. (1984). *Geology for Petroleum Exploration, Drilling, and Production.* New York: McGraw Hill Book Company.

Hyne, N.J. (2001). *Nontechnical Guide to Petroleum Geology, Exploration, Drilling, and Production* (2nd Ed.). Tulsa, Oklahoma: Penn Well Corporation.

Kennedy, R.S., Stanney, K.M. (1993). Simulator Sickness Questionnaire (SSQ): A new method for quantifying simulator sickness. International Journal of Aviation Psychology, 3, 203-220.

Kennedy, R.S., Stanney, K.M. (1997). The psychometrics of cybersickness. Communications of the ACM, 40, 66-68.

LaViola, J.J. (2000). A discussion of cybersickness in virtual environments, ACM SIGCHI Bulletin, 32, 47-56.

Lewis, C.H., Griffin, M.J. (1998). Human factors consideration in clinical applications of virtual reality. In Giuseppe Riva (Ed.), Virtual Reality in Neuro-Psycho-Physiology, Ios Press, Amsterdam, Netherlands.

North, F.K. (1985). *Petroleum Geology.* Boston, MA: Unwin Hyman.

Mine, M. (1995). Virtual environment interaction techniques (University of North Carolina Chapel Hill Computer Science Technical Report TR95-018).

Mizell, D., Jones, S., Jackson, P., Pickett, D. (1995). Is VR better than a workstation? A report on human performance experiments in progress. In M. Gobel. (Ed.), Virtual Environments (pp. 1-7), Springer.

Mizell D.. Jones, S., Slater, M., Spanlang, B., (2000). Comparing immersive virtual reality with other display modes for visualizing complex 3D geometry (University College London Technical Report GR/M866200/01) . Retrieved November 22, 2002, from http://www.cs.ucl.ac.uk/research/vr/Projects/Immersion/Experiment2/index.htm

Pape, D. (1996). CAVE user' s guide. Electronic Visualization Laboratory, University of Illinois at Chicago, Dec. 1996.

Potel, M. (1998). Motion sick in cyberspace. IEEE Computer Graphics and Applications, 18, 16-21.

Poupyrev, I., Billinghurst, M., Weghorst, S. (1996). The Go-Go interaction technique: Non-linear mapping for direct manipulation in VR. Proceedings of the ACM Symposium on User Interface Software and Technology. 79-80.

Poupyrev, I., Billinghurst, M., Weghorst, S., Ichikawa, T. (1997). A framework and testbed for studying manipulation techniques for immersive VR. Proceedings of ACM Symposium on Virtual Reality Software and Technology. 21-28.

Ruddle., R.A., Stephen, J.A., Jones, D.A. (1999). Navigating large-scale virtual environments: What differences occur between helmet-mounted and desk-top displays?. Presence: Teleoperators and Virtual Environments, 8, 157-168.

Slater, M., Linakis, V., Usoh, M., Kooper, R. (1996). Immersion, presence, and performance in virtual environments: An experiment with tri-dimensional chess. In M. G. (Ed.), ACM Virtual Reality Software and Technology (pp. 163-172).

So, R.H.Y., Lo, W.T. (1999). Cybersickness: An experimental study to isolate the effects of rotational scene oscillations. IEEE Proceedings of Virtual Reality 1999, 237-241.

Southren, T. (2000). An introduction to oilplan: Tech-21' s well planning software suite. Tech-21 Solutions Ltd.

Sutherland, I. (1965). The ultimate display. Proceedings of the IFIP Congress, 502-508.

Sutherland, I. (1968). A head-mounted three dimensional display. In Proceedings of the Fall Joint Computer Conference, 757-764.

Takahshi, Y., Murata, A. (2001). Change of equilibrium under the influence of VR experience. Proceedings of 10th IEEE International Worksho on Robot and Human Interactive Communication 2001, 642-647.

van Dam, A., Forsberg, A.S., Laidlaw, D.H., LaViola, J.J., Jr., Simpson, R.M. (2000). Immersive VR for scientific visualization: A progress report. Computer Graphics and Applications, 20(6), 26- 52.

Wernecke, J. (1994). The Inventor Mentor: Programming Object-Oriented 3D Graphics With Open Inventor, Release 2 (1st ed.), New York: Addison-Wesley Pub Co.

Winkler, C., Bosquet, F., Cavin, X., Jean-Claude, P., Design and implementation of an immersive geoscience toolkit (case study). Proceedings Visualization 1999, 429-432.

Yuan, X., Sun, H. (1997). Mechanical assembly with data glove devices. Canadian Conference on Electrical and Computer Engineering, 1, 177-180.

Appendix A

Experimental Instructions and Questionnaires

A.1 Instructions

The following text was referenced and loosely followed by the experimenter while describing the experiment to the participants.

A.1.1 General Well Planning instructions

Existing well paths are shown in purple. The original editable well path is shown in cyan. The original editable well is shown for reference. The edited well path is shown in white. The goal is to move the white well to a position where it does not intersect with any existing (purple) well while not exceeding the goal complexity value.

Complexity:
Complexity is driven by the shape of the well path. Curvature and sharp bends in the well path will increase the well complexity. Smooth straight wells with no "kinks" will have the lowest complexity values.

Interaction:
There are three editable objects in the scene. Two well sliders, represented by red spheres at each end of the well, defined the edit region. The pull point, represented by the red sphere pierced by the arrow, defines a point the well must pass through and the direction the well must have at that point. The white well path can be moved by adjusting these objects. The edit region can be adjusted by dragging the well sliders up and down the reference well. The pull point can be dragged and rotated independent of the reference well.

A.1.2 Desktop Interface Instructions

The desktop application has two modes: view mode and pick mode.

View mode, represented by the hand cursor, allows users to navigate the scene. The left mouse button rotates the scene, the middle mouse button translates the scene (up, down, right, left), left+middle mouse buttons zoom in and out.

Pick mode, represented by the arrow cursor, allows users to interact with objects in the scene. The left mouse button interacts with the red objects in the scene. Move the cursor over one of the red objects, the object should highlight. By pressing and holding the left mouse button the red object can be dragged.

Users can toggle between view and pick modes using the upper left hand buttons or using the ESC key.

The "Rotx" thumbwheel rotates the scene along the x-axis.
The "Roty" thumbwheel rotates the scene along the y-axis.
The "Zoom" thumbwheel zooms in and out of the scene.

The button with crossed arrows enables translation of the pull point (default).
The button with the curved arrow enables rotation of the pull point.

A complexity read-out is provided on the pull point object and on the lower bar of the viewer.

A.1.3 IVE Interface Instructions

Navigation:
The scene can be navigated by physically moving in the IVE and/or by using the wand. To use the wand to navigate point the wand in toward a destination. Pressing forward on the wand' s joystick will drive you toward the destination. Pressing backward on the wand' s joystick will drive you away from the destination. Pressing right or left on the wand' s joystick will turn you in that direction.

Interaction:
A white ray extends from the end of the wand. When the ray intersects with one of the selectable objects (one of the red spheres) the object will highlight. Pressing and holding the lower left wand button will drag the selected object. Well sliders will be dragged up and down the reference well' s path. The pull point can dragged to any position. To changed the orientation of the pull point, intersect the pull point with ray, press and hold the lower right button and the pull point will follow the orientation of the wand. The lower right button has no effect on the well sliders.

A.2 Simulator Sickness Questionnaire (SSQ)

SSQ

Please circle the degree of which you are currently experiencing the following symptoms:

General discomfort	None	Slight	Moderate	Severe
Fatigue	None	Slight	Moderate	Severe
Headache	None	Slight	Moderate	Severe
Eyestrain	None	Slight	Moderate	Severe
Difficultly focusing	None	Slight	Moderate	Severe
Increased salvation	None	Slight	Moderate	Severe
Sweating	None	Slight	Moderate	Severe
Nausea	None	Slight	Moderate	Severe
Difficultly concentrating	None	Slight	Moderate	Severe
Fullness of head	None	Slight	Moderate	Severe
Blurred vision	None	Slight	Moderate	Severe
Dizzy (eyes open)	None	Slight	Moderate	Severe
Dizzy (eyes closed)	None	Slight	Moderate	Severe
Vertigo	None	Slight	Moderate	Severe
Stomach awareness	None	Slight	Moderate	Severe
Burping	None	Slight	Moderate	Severe

Time:____________________________

A.3 Post-experiment user survey

User Survey

Please tell us about you background. Feel free to add comments to clarify your answers. If you need extra space, you may use the back of the page.

1. Are you:

   a) right-handed   b) left-handed c) ambidextrous

2. How many hours a week do you use a computer?
3. Have you ever used a virtual reality or virtual environment before today? If so please describe.

4. Which platform did you find more intuitive for interacting in three-dimensional space, the desktop computer or the IVE? Why?

5. What are your general impressions of the IVE?

Appendix B

Experimental Data

Subject ID	**Task01**			
	Time (s)	**Correct**	**Environment**	**Treatment**
s00	353	1	ive	1st
s01	266	1	desktop	1st
s02	434	0	desktop	2nd
s03	600	0	ive	2nd
s04	281	1	ive	1st
s05	460	1	desktop	1st
s06	333	0	desktop	2nd
s07	360	1	ive	2nd
s08	336	1	ive	1st
s09	375	1	desktop	1st
s10	423	1	desktop	2nd
s11	142	1	ive	2nd
s12	127	1	ive	1st
s13	445	1	desktop	1st
s14	243	1	desktop	2nd
s15	221	1	ive	2nd

	Task01			
	Time (s)	**Mean (s)**	**Std Dev**	**Correct**
Task	5399	337.44	123.18	13
Task (1st)	2643	330.38	107.18	8
Task (2nd)	2756	344.5	144.61	5
IVE	2420	302.5	150.78	7
IVE (1st)	1097	274.25	102.86	4
IVE (2nd)	1323	330.75	200.85	3
Desktop	2979	372.38	83.6	6
Desktop (1st)	1546	386.5	88.46	4
Desktop (2nd)	1433	358.25	89.16	2

Subject ID	**Task02**			
	Time (s)	**Correct**	**Environment**	**Treatment**
s00	313	1	ive	1st
s01	600	0	desktop	1st
s02	306	1	desktop	2nd
s03	330	1	ive	2nd
s04	249	0	ive	1st
s05	572	1	desktop	1st
s06	277	0	desktop	2nd
s07	574	1	ive	2nd
s08	347	1	ive	1st
s09	222	0	desktop	1st
s10	206	0	desktop	2nd
s11	247	1	ive	2nd
s12	383	1	ive	1st
s13	262	0	desktop	1st
s14	270	1	desktop	2nd
s15	377	1	ive	2nd

	Task02			
	Time (s)	**Mean (s)**	**Std Dev**	**Correct**
Task	5535	345.94	127.71	10
Task (1st)	2948	368.5	144.34	4
Task (2nd)	2587	323.38	113.79	6
IVE	2820	352.5	103.19	7
IVE (1st)	1292	323	57.01	3
IVE (2nd)	1528	382	138.83	4
Desktop	2715	339.38	155.57	3
Desktop (1st)	1656	414	199.61	1
Desktop (2nd)	1059	264.75	42.15	2

Subject ID	**Task03**			
	Time (s)	**Correct**	**Environment**	**Treatment**
s00	600	0	desktop	2nd
s01	352	1	ive	2nd
s02	327	1	ive	1st
s03	600	0	desktop	1st
s04	452	1	desktop	2nd
s05	499	1	ive	2nd
s06	333	1	ive	1st
s07	600	0	desktop	1st
s08	594	1	desktop	2nd
s09	62	1	ive	2nd
s10	206	1	ive	1st
s11	457	1	desktop	1st
s12	383	1	desktop	2nd
s13	183	0	ive	2nd
s14	126	1	ive	1st
s15	349	0	desktop	1st

	Task03			
	Time (s)	**Mean (s)**	**Std Dev**	**Correct**
Task	6123	382.69	174.67	11
Task (1st)	2998	374.75	170.41	5
Task (2nd)	3125	390.63	190.25	6
IVE	2088	261	142.13	7
IVE (1st)	992	248	100.19	4
IVE (2nd)	1096	274	191.43	3
Desktop	4035	504.38	106.45	4
Desktop (1st)	2006	501.5	121.98	1
Desktop (2nd)	2029	507.25	107.42	3

Subject ID	**Task04**			
	Time (s)	**Correct**	**Environment**	**Treatment**
s00	196	1	desktop	2nd
s01	191	1	ive	2nd
s02	398	1	ive	1st
s03	520	1	desktop	1st
s04	323	1	desktop	2nd
s05	491	1	ive	2nd
s06	208	0	ive	1st
s07	600	0	desktop	1st
s08	148	1	desktop	2nd
s09	135	0	ive	2nd
s10	385	1	ive	1st
s11	214	0	desktop	1st
s12	376	0	desktop	2nd
s13	154	1	ive	2nd
s14	407	1	ive	1st
s15	600	0	desktop	1st

	Task04			
	Time (s)	**Mean (s)**	**Std Dev**	**Correct**
Task	5346	334.13	161.46	10
Task (1st)	3332	416.5	153.05	4
Task (2nd)	2014	251.75	130.11	6
IVE	2369	296.13	138.11	6
IVE (1st)	1398	349.5	94.76	3
IVE (2nd)	971	242.75	167.13	3
Desktop	2977	372.13	182.99	4
Desktop (1st)	1934	483.5	183.58	1
Desktop (2nd)	1043	260.75	106.56	3

Subject ID	IVE Correct	Desktop Correct	IVE Time	Desktop Time
s00	2	1	666	796
s01	2	1	543	866
s02	2	1	725	740
s03	1	1	930	1120
s04	1	2	530	775
s05	2	2	990	1032
s06	1	0	541	610
s07	2	0	934	1200
s08	2	2	683	742
s09	1	1	197	866
s10	2	1	591	629
s11	2	1	389	671
s12	2	1	510	759
s13	1	1	337	707
s14	2	2	533	513
s15	2	0	598	949

SSQ Symptom	s00.0	s00.1	s00.2	s01.0	s01.1	s01.2
General Discomfort	0	0	0	0	0	0
Fatigue	0	0	0	0	0	0
Headache	0	1	1	0	0	0
Eyestrain	0	1	1	0	0	0
Difficulty Focusing	0	0	0	1	1	1
Increased Salvation	0	0	0	0	0	0
Sweating	0	0	0	0	0	0
Nausea	0	0	0	0	0	0
Difficulty Concentrating	0	0	0	0	0	0
Fullness of Head	0	0	0	0	0	0
Blurred Vision	0	0	0	0	0	0
Dizzy (eyes open)	0	0	0	0	0	0
Dizzy (eyes closed)	0	0	0	0	0	0
Vertigo	0	0	0	0	0	0
Stomach Awareness	0	0	0	0	0	0
Burping	0	0	0	0	0	0
Nausea	0	0	0	0	0	0
Oculmotor	0	15.16	15.16	7.58	7.58	7.58
Disorientation	0	0	0	13.92	13.92	13.92
Total Score	0	7.48	7.48	3.74	3.74	3.74
Nausea Delta		0	0		0	0
Oculmotor Delta		15.16	0		0	0
Disorentation Delta		0	0		0	0
Total Score Delta		7.48	0		0	0

SSQ Symptom	s02.0	s02.1	s02.2	s03.0	s03.1	s03.2
General Discomfort	0	0	0	0	0	0
Fatigue	1	1	1	1	1	1
Headache	0	0	0	0	0	0
Eyestrain	0	1	1	0	0	1
Difficulty Focusing	0	0	0	1	1	1
Increased Salvation	0	0	0	0	0	0
Sweating	0	0	0	0	0	0
Nausea	0	0	0	0	0	0
Difficulty Concentrating	0	0	0	0	0	0
Fullness of Head	0	0	0	0	0	0
Blurred Vision	0	0	0	1	1	1
Dizzy (eyes open)	0	0	0	0	0	0
Dizzy (eyes closed)	0	0	0	0	0	0
Vertigo	0	0	0	0	0	0
Stomach Awareness	0	0	0	1	1	1
Burping	0	0	0	0	0	0
Nausea	0	0	0	9.54	9.54	9.54
Oculmotor	7.58	15.16	15.16	22.74	22.74	30.32
Disorientation	0	0	0	27.84	27.84	27.84
Total Score	3.74	7.48	7.48	14.96	14.96	18.7
Nausea Delta		0	0		0	0
Oculmotor Delta		7.58	0		0	7.58
Disorentation Delta		0	0		0	0
Total Score Delta		3.74	0		0	3.74

SSQ Symptom	s04.0	s04.1	s04.2	s05.0	s05.1	s05.2
General Discomfort	0	0	0	0	0	0
Fatigue	1	0	0	0	0	0
Headache	0	0	0	0	0	0
Eyestrain	0	1	1	0	1	1
Difficulty Focusing	0	0	0	0	0	0
Increased Salvation	0	0	0	0	0	0
Sweating	1	0	0	0	0	0
Nausea	0	0	0	0	0	0
Difficulty Concentrating	0	0	0	0	0	0
Fullness of Head	0	0	0	0	0	0
Blurred Vision	0	0	0	0	0	0
Dizzy (eyes open)	0	0	0	0	0	0
Dizzy (eyes closed)	0	0	0	0	0	0
Vertigo	0	0	0	0	0	0
Stomach Awareness	0	0	0	0	0	0
Burping	0	0	0	0	0	0
Nausea	9.54	0	0	0	0	0
Oculmotor	7.58	7.58	7.58	0	7.58	7.58
Disorientation	0	0	0	0	0	0
Total Score	7.48	3.74	3.74	0	3.74	3.74
Nausea Delta		-9.54	0		0	0
Oculmotor Delta		0	0		7.58	0
Disorentation Delta		0	0		0	0
Total Score Delta		-3.74	0		3.74	0

SSQ Symptom	s06.0	s06.1	s06.2	s07.0	s07.1	s07.2
General Discomfort	1	1	0	0	0	0
Fatigue	0	0	2	0	0	0
Headache	0	0	0	0	0	0
Eyestrain	1	1	1	0	1	0
Difficulty Focusing	1	1	1	0	0	0
Increased Salvation	1	0	0	0	0	0
Sweating	0	0	0	0	0	0
Nausea	0	0	0	0	0	0
Difficulty Concentrating	0	0	0	0	0	0
Fullness of Head	1	1	0	0	0	0
Blurred Vision	1	1	1	0	0	0
Dizzy (eyes open)	0	0	0	0	0	0
Dizzy (eyes closed)	1	1	0	0	0	0
Vertigo	0	0	0	0	0	0
Stomach Awareness	0	0	0	0	0	0
Burping	0	0	0	0	0	0
Nausea	19.08	9.54	0	0	0	0
Oculmotor	30.32	30.32	37.9	0	7.58	0
Disorientation	55.68	55.68	27.84	0	0	0
Total Score	26.18	22.44	18.7	0	3.74	0
Nausea Delta		-9.54	-9.54		0	0
Oculmotor Delta		0	7.58		7.58	-7.58
Disorentation Delta		0	-27.84		0	0
Total Score Delta		-3.74	-3.74		3.74	-3.74

SSQ Symptom	s08.0	s08.1	s08.2	s09.0	s09.1	s09.2
General Discomfort	0	0	0	0	0	0
Fatigue	0	0	0	0	0	0
Headache	0	0	0	0	0	0
Eyestrain	0	0	0	0	1	0
Difficulty Focusing	0	0	0	0	0	0
Increased Salvation	0	0	0	0	0	0
Sweating	0	0	0	0	0	0
Nausea	0	0	0	0	0	0
Difficulty Concentrating	0	0	0	0	0	0
Fullness of Head	0	0	0	0	0	0
Blurred Vision	0	0	0	0	0	0
Dizzy (eyes open)	0	1	0	0	0	0
Dizzy (eyes closed)	0	1	0	0	0	0
Vertigo	0	0	0	0	0	0
Stomach Awareness	0	0	0	0	0	0
Burping	0	0	0	0	0	0
Nausea	0	0	0	0	0	0
Oculmotor	0	0	0	0	7.58	0
Disorientation	0	27.84	0	0	0	0
Total Score	0	7.48	0	0	3.74	0
Nausea Delta		0	0		0	0
Oculmotor Delta		0	0		7.58	-7.58
Disorentation Delta		27.84	-27.84		0	0
Total Score Delta		7.48	-7.48		3.74	-3.74

SSQ Symptom	s10.0	s10.1	s10.2	s11.0	s11.1	s11.2
General Discomfort	0	0	0	0	0	0
Fatigue	0	0	1	0	0	0
Headache	0	0	0	0	0	0
Eyestrain	0	0	0	0	0	0
Difficulty Focusing	0	0	0	0	0	0
Increased Salvation	0	0	0	0	0	0
Sweating	0	0	0	0	0	0
Nausea	0	0	0	0	0	0
Difficulty Concentrating	0	0	0	0	0	0
Fullness of Head	0	0	0	0	0	0
Blurred Vision	0	0	0	0	0	0
Dizzy (eyes open)	0	0	0	0	0	0
Dizzy (eyes closed)	0	0	0	0	0	0
Vertigo	0	0	0	0	0	0
Stomach Awareness	0	0	0	0	0	0
Burping	0	0	0	0	0	0
Nausea	0	0	0	0	0	0
Oculmotor	0	0	7.58	0	0	0
Disorientation	0	0	0	0	0	0
Total Score	0	0	3.74	0	0	0
Nausea Delta		0	0		0	0
Oculmotor Delta		0	7.58		0	0
Disorentation Delta		0	0		0	0
Total Score Delta		0	3.74		0	0

SSQ Symptom	s12.0	s12.1	s12.2	s13.0	s13.1	s13.2
General Discomfort	1	1	1	1	0	0
Fatigue	0	0	0	0	0	0
Headache	1	1	1	0	0	0
Eyestrain	1	1	1	1	0	0
Difficulty Focusing	0	0	0	1	0	0
Increased Salvation	0	0	0	0	0	0
Sweating	0	0	0	0	0	0
Nausea	0	0	0	0	0	0
Difficulty Concentrating	0	0	0	0	0	0
Fullness of Head	0	0	0	0	0	0
Blurred Vision	0	0	0	0	0	0
Dizzy (eyes open)	0	0	0	0	0	0
Dizzy (eyes closed)	0	0	0	0	0	0
Vertigo	0	0	0	0	0	0
Stomach Awareness	0	0	0	1	1	1
Burping	0	0	0	0	0	0
Nausea	9.54	9.54	9.54	19.08	9.54	9.54
Oculmotor	22.74	22.74	22.74	22.74	0	0
Disorientation	0	0	0	13.92	0	0
Total Score	11.22	11.22	11.22	14.96	3.74	3.74
Nausea Delta		0	0		-9.54	0
Oculmotor Delta		0	0		-22.74	0
Disorentation Delta		0	0		-13.92	0
Total Score Delta		0	0		-11.22	0

SSQ Symptom	s14.0	s14.1	s14.2	s15.0	s15.1	s15.2
General Discomfort	0	0	0	0	0	0
Fatigue	0	0	0	0	1	0
Headache	0	0	0	0	0	0
Eyestrain	0	0	1	0	0	0
Difficulty Focusing	0	0	0	0	0	0
Increased Salvation	0	0	0	0	1	1
Sweating	0	0	0	0	0	0
Nausea	0	0	0	0	0	0
Difficulty Concentrating	0	0	0	0	1	0
Fullness of Head	0	0	0	0	1	1
Blurred Vision	0	0	0	0	1	0
Dizzy (eyes open)	0	0	0	0	0	0
Dizzy (eyes closed)	0	0	0	0	0	0
Vertigo	0	0	0	0	0	0
Stomach Awareness	0	0	0	0	0	0
Burping	0	0	0	0	0	0
Nausea	0	0	0	0	19.08	9.54
Oculmotor	0	0	7.58	0	22.74	0
Disorientation	0	0	0	0	27.84	13.92
Total Score	0	0	3.74	0	18.7	7.48
Nausea Delta		0	0		19.08	-9.54
Oculmotor Delta		0	7.58		22.74	-22.74
Disorentation Delta		0	0		27.84	-13.92
Total Score Delta		0	3.74		18.7	-11.22

Appendix C

Participants' Remarks

IVE, easier to navigate and visually inspect that there were no intersections between pipes.

IVE was more intuitive. Controls were much easier to use an more responsive. Motions and controls were more obvious and natural feeling. Desktop switching (mouse clicking) between modes was inconvenient and distracting. Keeping track of the mouse was hard. IVE became very natural to interact with. Desktop had better detailed visuals for precision positioning.

IVE Ease of movement -forward and backward – more intuitive. Ease of getting to the perspective that you want was key.

The IVE! I could crouch down, see around things, etc rather than tediously move wheels and the desktop display. It Translated my body movements so it was much smoother working around the wells.

IVE. Being able to walk around and looking at the model at different angles was easier/faster than truing reorient the view.

IVE. East to just turn head to point of interest instead of translating the 3D motion in to the appropriate combo (always more than one) of the desktop controls required to get there. Even being very comfortable with the mouse, I found movement in 3D with the wand much more intuitive.

IVE, with the IVE I was the frame of reference as opposed to the arbitrary point of reference with the desktop. Also, the non-modal nature of the IVE let concentrate on the problem' s solution.

IVE. Much more intuitive interface, much easier to adapt to the controls. Feels more natural moving around with the scene. Less eye strain.

IVE, because I could physically move around the objects to see the intersections.

IVE. Better interface for navigation within 3D environment: Pan, Zoom, Rotate. Much more intuitive. Also better spatial orientation, accompanied by more 3D space to work within. Desktop had good 3D, little sharper, but less room, 3D on a 2D surface, cumbersome navigation when close.

IVE. Much easier and more intuitive to navigate. Easier to view "big" picture and examine the full environment.

IVE. It was much easier to fine tune my position in the 3d space to get the goal complexity value. With the desktop environment I felt like I was fighting the controls.

The IVE, because I could zoom it in to a larger degree and the walls gave you a better impression of right and left than the computer's screen did. In the IVE movements are possible with the wand and body, on the desktop you only have the mouse.

IVE. It maintained vertical orientation better (I lost "up" in the desktop a few times) and far away problems/features could more easily be seen and evaluated by turning head, stepping, leaning, or squating in the IVE. The one advantage of the desktop was the ability to easily orbit the scene, surveying general features and dimensions very quickly. This is more an interface issue than a display issue.

IVE. Easier to get into picture and maneuver. Able to see all angles and intersections at one time, as well as in close to specific area.

It's a very cool idea. Of course it has some minor visual glitches that need to be addressed, but all in all it's a much better way to solve problems involving a 3 dimensional space.

The feeling of being inside the picture is very strong and authentic. The wand responses well to what the users wants to do or move. All in all: great invention and impressive experience.

I think that it is a good way to visualize 3D info/data since you can move around within it.

Fantastic experience. Only takes a short time to become convinced that the visuals work like a "real" environment.

I can see where it'd be quite useful – especially with mods to the UI overtime with user testing. Going "wireless" would be optima, with four walls instead of three.

The coolest thing I've ever seen a computer do! Very intuitive, very usable.

Pretty neat technology. Easy to walk and look around. Glasses are a little uncomfortable. At times it was hard to see the complexity number when moving the two end points.

Enjoyed the IVE, more comfortable in it than on the desktop. Controls in the IVE more intuitive thus could focus on the problem more, (rather than focusing and thinking/translating the tools, mouse and three scroll bars, plus view vs edit mode) in the desktop. Felt more sure of checking for intersections in the IVE, to check for intersections in the desktop took much more time. In IVE I felt I could see the full problem, could focus in on one spot, but then just look (with my head) to the other end of the well, very easy and quick (to go parts, whole and vise versa). But in the desktop took multiple steps to focus on part, had to multiple steps to whole, the multiple steps to get back to part to continue with tweaks.

Fantastic. A wonderful experience. Can see how the IVE can be/is helpful to full picture.

Very cool, natural way to view information. Sometimes I forgot that I could move around and relied on the wand too much.

Very interesting concept, with many many potential uses. A great interface for interacting with spatial data of all kinds – not just geographic and geological, but micro-scale as well. Medical?

Impressive and easy to use.

IVE much more realistic, and more intuitive to solving a 3D spatial problem.

Very positive experience.

I'm impressed – this is what VR is supposed to be – clearly the technology has come of age.