

Immersive Well-Path Editing: Investigating the Added Value of Immersion

Kenny Gruchalla
BP Center for Visualization
Computer Science Department
University of Colorado at Boulder
gruchall@colorado.edu

Abstract

The benefits of immersive visualization are primarily anecdotal; there have been few controlled user 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. 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™-like IVE. Fifteen of the participants completed well-path editing tasks faster in the IVE than in the desktop environment. The increased speed was complimented by a statistically significant increase in correct solutions in the IVE. The results suggest that an IVE allows for faster and more accurate problem solving in a complex three-dimensional domain.

1. Introduction

There is a common assumption that immersive virtual environments provide an improved interface to view and interact with three dimensional structures over more traditional desktop graphics workstations [1]. 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. This interface would seemingly provide a more natural and intuitive means for viewing and interacting with three-dimensional virtual worlds in a variety of industrial settings. 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 that the benefits

are primarily anecdotal. The goal of this research was 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. The planning of a new oil well-path through the existing wells of a mature oilfield is such a task. It requires spatial understanding of a complex three-dimensional environment and the precise placement of objects within that environment. The Immersive Drilling Planner (IDP) is a software application capable of visualizing a mature oilfield and editing a new path within that oilfield on both a desktop environment and in an 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 paper describes an experiment designed to compare an IVE with a stereoscopic desktop environment in the performance and correctness of a well-path editing task.

2. Related Work

Most human performance virtual environment studies have focused on comparing various navigation and manipulation techniques within the same virtual environment. Only a few studies have attempted to compare IVEs with traditional desktop environments.

Ruddle, Payne, and Jones designed a virtual building walk-through experiment to compare a helmet-mounted display with a desktop monitor display [2]. 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.

Pausch, Proffitt, and Williams conducted a user study comparing a search task between a head-tracked helmet-mounted display and stationary helmet-mounted display [3]. Participants were placed in the center of a virtual room and instructed to search for a camouflaged target. The study showed that when a target was present there was no significant performance improvement in the immersed environment. However, when the target was not present participants in the immersed environment were able to make the conclusion substantially faster than the participants using the stationary display. The study also found a positive transfer of training effect from the immersive environment to the stationary display, and a negative transfer of training effect from the stationary display to the head-tracked environment.

Arns, Cook, and Cruz-Neira conducted a user study comparing statistical data analysis on a desktop and an IVE [4]. The experiment compared both identification and interaction tasks. 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 versus 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.

3. Immersive Drilling Planner

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 was built on top of the CAVELIB™ and Open Inventor libraries. 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 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 were 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.

3.1. Well Planning Background

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* [5]. 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 through straight hole drilling [5].

Oilfields exploited by directional drilling can quickly become a tortuous underground labyrinth of wells, creating a very complex spatial domain (see Figure 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 [5]. 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.

The IDP represent a well-path by its *uncertainty surface*, this surface forms a volume that is known to contain well-path. 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 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. These uncertainties can be visualized as a elliptical volume perpendicular to the well-path. Accumulating the errors at each point along the well enables an uncertainty



Figure 1. Snapshot of a virtual oilfield constructed from an actual well log dataset. Mature oilfields can be very complex three-dimensional structures.

surface to be constructed [5].

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 Design

The three-dimensional user interface is a critical component of an immersive virtual environment’s usability. Bowman [6] 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.

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 [7]. The IDP implements a combination of two well known techniques: *physical navigation* and *pointing*. Physical navigation 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 [8]. However, an oilfield scaled to fit wholly within the physical boundaries of the IVE would be unusably small. Therefore the pointing technique was used to help overcome the limitations of physical navigation. In this technique, the direction of motion depends upon the current orientation of the user’s hand or hand held device [7]. User studies have suggested that the pointing technique is well suited for general-purpose applications that require speed and accuracy [9]. Using a combination of these 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 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 rotates the scene around the user.

Interaction with a virtual object involves selecting, positioning and rotating the object in the virtual environment. The IDP implements a variation of 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 2). 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.

3.3. Desktop Design

The IVE version of the IDP could be run directly on a desktop workstation using the CAVELIB™ simulator. However, the simulator was designed as a tool to test immersive applications, not as a production desktop interface. Therefore, the Open Inventor examiner viewer



Figure 2. Photograph of a IDP user interacting with the virtual oilfield inside the IVE using the ray-casting technique.

is used by the desktop implementation of the IDP as the front-end user interface. 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.

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 handle box manipulator is used to position interactive objects in the desktop version of the IDP. This manipulator draws a bounding box around the interactive object. The manipulator responds to click-and-drag mouse events by translating the interactive object it surrounds. It also provides scaling functionality, which is not used in the IDP. A trackball manipulator is used to rotate interactive objects in the desktop version of the IDP. This manipulator wraps the interactive object with three circular stripes. It 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.

4. Experimental Design

The experiment consisted of four separate logged experimental tasks (denoted Task01, Task02, Task03, and Task04) and a training task (denoted Task00). Each par-

ticipant performed the training task and two experimental tasks on the desktop and the training task and the other 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 1). The independent variable was the environment: the head-tracked stereoscopic IVE versus the stereoscopic desktop environment. The dependent variables were the time to complete the task and the correctness of the final well-path.

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 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. 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 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 manipulated, the edited path's uncertainty surface is updated in real-time.

4.1. Participants

Nineteen unpaid participants were recruited from the staff and students at the University of Colorado at Boulder and from employees of several Colorado software firms. 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.

4.2. Apparatus

The IVE used for this study is located on the University of Colorado campus at the B.P. Center for Visualization.

Table 1. Counter-balanced experimental design

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

The IVE at the B.P. Center for Visualization is a Mechdyne™MD Flex™, which is a configurable large screen projection-based system. In closed configuration, the MD Flex™ is a 12'x12'x10' theater, resembling a CAVE™. The MD Flex™ can be re-configured to a 36'x12'x10' open configuration or presentation mode. 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™ consists of four walls: three rear-projected screens measuring 12'x10' which form the right, back, and left walls of the IVE, the fourth wall is the 12'x12' floor which is projected from above.

The four display screens were driven by one Silicon Graphics Incorporated (SGI) Origin 3800 computer with four SGI Infinite Reality3 graphics pipes. Each pipe feeds a Barco™909 projector. The 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 was created inside the IVE through active stereo projection. Participants wore infrared CrystalEyes™LCD shutter glasses to view the stereoscopic images. The sole interaction device used in this study was a wired InterSense™wand. 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. An InterSense™VET 900 tracking system tracked the position and orientation of the shutter glasses and the wand.

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. Unlike those in most homes and offices, the desktop interface in this study was connected to a SGI Origin 3800 (the same machine used to drive the IVE). 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™LCD shutter glasses. Unlike the immersive environment, the desktop environment did not include head

tracking.

4.3. Procedure

The experimental procedure (approved by an expedited review by the University of Colorado Human Research Committee) 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 a Subject Informed Consent Form. 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 the participant would perform the training task and two logged experimental tasks in the other environment. Again, while performing the training task, the experimenter would read from a script describing the user interface in that environment. After completing the second treatment, the participant was then asked to complete a post-experiment questionnaire.

Each task would begin 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. When the start button was pressed, the dialog would be closed and the test application began 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) were logged.

The participant began at a fixed starting position outside of the virtual field, then would navigate through the field to the new well. Then, through a series of well slider and pull point movements, the participant would edit the path of the new well. A three-dimensional text readout above the pull point provided the user with complexity

value feedback. Once the participant believed that the new path's uncertainty surface did not intersect the uncertainty surface of any existing well and that the new path had a complexity value at or below the goal complexity, the participant was instructed to complete the task by closing the application. If the allotted time was reached, the test application would terminate automatically.

4.4. Performance Measures

There were two performance measures per task: the time to complete the task and the correctness of the final well-path. The IDP maintained a timed log of the participant's interactions with the virtual environment; the time to complete the task was derived from the log. The final well-path was 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.

5. Results

Comparing the number of correct solutions within the participants shows a significant difference between the two environments (see Figure 3). 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 4). Of the sixteen participants, only one participant took more time in the IVE. The sign test shows this to be statistically significant at the 0.001 significance level.

The results were also analyzed using an ANOVA. The environment (IVE versus desktop) was treated as a within-subject factor, and the environment order and the task order were used as two between-subjects factors. The environment order represents whether the first treatment was in the IVE or on the desktop. The task order represents which pair of tasks (Task01 and Task02 or Task03 and Task04) occurred in the first treatment. The ANOVA of total time spent on each pair of tasks (see Table 2) shows a highly significant effect $F(1,12)=54.740$, $p=0.000$ of the environment, however an interaction effect $F(1,12)=12.519$, $p=0.004$ between the environment and the environment order is also present. A paired samples t-test (see Table 3) shows that a significant effect of the environment does exist. The ANOVA of the number of




Figure 3. Graph illustrating the number of correct solutions for each participant in each environment.




Figure 4. Graph illustrating the total accumulated solution time for each participant in each environment.

correct solutions in each pair of tasks (see Table 4) shows a significant effect $F(1,12)=10.714$, $p=0.007$ of environment.

An analysis of individual tasks is difficult since the tasks were not fully crossed, that is, Task02 always followed Task01 and Task04 always followed Task03. However, it is clear that there are differences between the tasks (see Figures 5 and 6). 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

Table 2. ANOVA of solution time

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Environment	282940.031	1	282940.031	54.74	.000
Environment x pair order	12285.281	1	12285.281	2.377	.149
Environment x environment order	64710.031	1	64710.031	12.519	.004
Environment x pair order x environment order	8944.531	1	8944.531	1.73	.213
Error	62025.625	12	5168.802		

Table 3. Paired Samples T-test : Immersed Time - Desktop Time

Environment Order	Paired Differences					t	df	Sig (2-tailed)
	Mean	Std. Dev.	Std. Err. Mean	95% Confidence Interval Lower	Upper			
IVE first	-98.128	101.566	35.909	-183.036	-13.214	-2.733	7.000	.029
Desktop first	-278.000	116.068	41.036	-375.035	-180.965	-6.775	7.000	0.000

Table 4. ANOVA of number of correct solutions

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Environment	3.125	1	3.125	10.714	.007
Environment x pair order	1.125	1	1.125	3.857	.073
Environment x environment order	0.125	1	0.125	.429	.525
Environment x pair order x environment order	0.125	1	0.125	.429	.525
Error	3.500	12	.292		




Figure 5. Graph of the number of correct solutions by task.




Figure 6. Graph of the mean solution time by task. Error bars show standard deviation.

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.

Participants' written comments reflect the added value of immersion. In a post-experiment questionnaire, all sixteen 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.

6. Conclusions

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 accumulated solution time. This speed difference was shown to be statistically significant.

Participants in this study had more accurate percep-

tions 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. This was also shown to be statistically significant.

The data suggest that IVEs may be more suitable for certain types of problems. Notice in Figures 5 and 6 that the number of correct solutions and the mean time for Task01 are 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. The two tasks are similar, but Task01 is less spatially complicated than Task03 as there are fewer wells in the immediate vicinity of the new Task01 well. A similar phenomenon was observed during the pilot tests. Several initial pilot tests involved spatially simple domains and failed to show a significant difference between the two environments. These observations imply that the added value of immersion may be correlated to the spatial complexity of the problem. Clearly, there may be classes of spatial problems that would benefit from immersion.

This study showed that oil well-path planning can be improved by immersion, but does not address why it was improved. Were the benefits in the IVE a result of a better understanding of the data? Were the benefits a result of more natural navigation? Were the benefits a result of more natural interaction with the virtual objects? Or were the benefits a result of some combination of the three? There have been studies [2] showing that navigation through a three-dimensional world is improved by immersion, but there are no controlled studies showing which types of interactions are improved by immersion. A logical progression of this work would be to identify classes of problems that benefit from immersion, by constructing a taxonomy of user interactions that are faster, more precise, and more accurate in an IVE and by constructing a taxonomy of spatial situations in which human understanding is improved in an IVE.

This work is a controlled study 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. Generally, not only were 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.

7. Acknowledgments

The author would like to thank all the research participants for volunteering their time, Clayton Lewis for his guidance, Bill Oliver for his assistance with the statistical analysis, Jonathan Marbach for his countless contributions to this work, and the staff at the BP Center for Visualization for their support.

References

- [1] 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.
- [2] Ruddle, R.A., Payne S.J., 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.
- [3] Pausch, R., Proffitt, D., Williams, G. (1997). Quantifying Immersion in Virtual Reality. *Proceedings of the 24th annual conference on computer graphics and interactive techniques*, 13-18.
- [4] Arns, L., Cook, D., Cruz-Neira, C. (1999). The benefits of statistical visualization in an immersive environment. *Proceedings of IEEE Virtual Reality 1999*, 88-95.
- [5] North, F.K. (1985). *Petroleum Geology*. Boston, MA: Unwin Hyman.
- [6] Bowman, D.A. (1999). *Interaction Techniques for Common Tasks in Immersive Virtual Environments: Design, Evaluation, and Application*. Doctoral dissertation, Georgia Institute of Technology.
- [7] Mine, M. (1995). *Virtual environment interaction techniques* (University of North Carolina Chapel Hill Computer Science Technical Report TR95-018).
- [8] Usuh, M., Slater, M. (1995). An exploration of immersive virtual environments. *Endeavor*, 19, 1, 34-38.
- [9] 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.