Towards Usable VR: An Empirical Study of User Interfaces for Immersive Virtual Environments

Robert W. Lindeman **John L. Sibert** **James K. Hahn**
Institute for Computer Graphics
The George Washington University
801 22nd Street, NW, Washington, DC, 20052
+1-202-994-7181
[gogo | sibert | hahn]@seas.gwu.edu

ABSTRACT
This paper reports empirical results from a study into the use of 2D widgets in 3D immersive virtual environments. Several researchers have proposed the use of 2D interaction techniques in 3D environments, however little empirical work has been done to test the usability of such approaches. We present the results of two experiments conducted on low-level 2D manipulation tasks within an immersive virtual environment. We empirically show that the addition of passive-haptic feedback for use in precise UI manipulation tasks can significantly increase user performance. Furthermore, users prefer interfaces that provide a physical surface, and that allow them to work with interface widgets in the same visual field of view as the objects they are modifying.

Keywords
3D user interfaces, bimanual interaction, virtual environments, virtual reality, passive-haptic feedback

INTRODUCTION
The introduction of Virtual Environment (VE) systems into mainstream computing has not been as rapid as researchers first projected. Indeed, outside of the entertainment industry, most VE systems in use today remain in research labs and universities. One of the reasons for this is that we do not know enough about the nature of user interaction in VEs to create systems which allow people to do real work [16, 13]. This paper presents empirical results from experiments designed to shed some light on effective user interface (UI) techniques for Immersive Virtual Environments (IVEs). An IVE is a virtual world that a user interacts with using devices that block out all elements of the real world that are not part of the experience. We build on recent work in the application of 2D interfaces to 3D worlds in order to identify those aspects which promote usability. We focus here on *symbolic* manipulation, as opposed to *direct* manipulation (see [11] for a good overview of direct manipulation techniques).

Current IVE Interaction Techniques
In order to support symbolic interaction in 3-space, some IVE applications have abandoned desktop interface devices for more freeform interface methods. Glove interfaces allow the user to interact with the environment using gestural commands [4, 9, 8, 18] or menus "floating" in space [11, 7, 5, 12, 19, 6]. The latter use either the user's finger or some sort of laser-pointer, combined with a physical button-click, to manipulate widgets. With these types of interfaces, however, it is difficult to perform precise movements, such as dragging a slider to a specified location, or selecting from a pick list. Part of the difficulty in performing these tasks comes from the fact that the user is pointing in free space, without the aid of anything to steady the hands [11].

Feiner et al describe an approach for using 2D windows in 3D worlds [7]. The system they describe is implemented for an augmented reality system, however the idea can be applied to immersive environments as well. Feiner et al identify three different types of windows, differentiated by what the window is fixed to. World-fixed windows (called *surround-fixed* windows in [7]) have an absolute, fixed position in the VE. As the user moves or looks around, the world-fixed windows go out of, or come into, view, as if they were fixed in space. The second type of window is a view-fixed window (*display-fixed* in [7]). These windows move along with the user as they look around within the VE. They remain at a fixed location, relative to the user's viewpoint, and may be suitable for manipulating system-wide attributes, such as the rendering method to use for objects (Phong, Gouraud, wireframe, etc). The third type of window is an object-fixed window (*world-fixed* in [7]). Each object-fixed window is fixed, relative to a specific object in the VE. If the object moves, the window moves along with it. These may be used to display and manipulate object attributes, such as to display the current velocity of an airplane, or to turn on a virtual lamp. We will use the terms world-fixed, view-fixed, and object-fixed for the remainder of this paper in the manner just defined.

Deering uses hybrid 2D/3D menu widgets organized in a disk layout [6]. The disk is parallel to the view plane, and the user selects items with a 3-button, 6-degree of freedom (DOF) wand held in the dominant hand. When invoked, the menu pops up in a fixed position relative to the current position of the tip of the wand. Similarly, Sowizral [15] and Wloka et al [20] use menus that pop-up in the same location relative to a 6-DOF mouse, then use the mouse buttons to cycle through menu entries. Each of these methods, however, provides limited user precision because of a lack of physical support for manipulations.

To counter this, some researchers have introduced the use of "pen-and-tablet" interfaces [1, 2, 11, 3, 10]. These approaches register an object-fixed window with a physical prop held in the non-dominant hand. We call these *hand-held windows*. Users interact with them using either a finger, or a stylus held in the dominant hand. These interfaces combine the power of 2D window interfaces with the necessary freedom provided by 3D interfaces.

There are many advantages to these approaches. First, hand-held windows move along with the user, so they are always within reach. Second, they do not clutter the user's view, unless explicitly moved there by the user. Hand-held windows also take advantage of the proprioceptive sense, because they reside close to the non-dominant hand. Finally, some systems using hand-held windows have incorporated a lightweight, physical surface that the user carries around, increasing precision [2, 3, 17]. Storing the physical surface when not in use can be an issue with these systems, and increased arm fatigue may degrade performance during prolonged use. Most of the previous work in this field has called for detailed study into how these interfaces can most effectively be designed to enhance user performance.

UI Interaction Decomposition

In order to better study UI interaction techniques, we can decompose user interaction into basic motions, using what Shneiderman calls *Widget-Level* decomposition [14]. This approach looks at the widgets that are defined in the system, and bases decomposition on their manipulation. The testbed we have designed provides 2D widgets for testing typical UI tasks, such as drag-and-drop and button presses. We define (at least) two distinct types of actions based on these widgets: *discrete* (or open-loop) actions and *continuous* (closed-loop) actions. Discrete actions involve ballistic selection operations, such as clicking a toolbar icon, double clicking a filename, or positioning an input cursor. Continuous actions include dragging sliders, using drag-and-drop to move a file, or accessing a cascading pull-down menu.

We have designed a number of empirical studies of user performance and preference on tasks which focus on these basic motions. The results of two of these studies are presented here, and can be used to suggest how designers can develop general IVE interfaces that allow users to work efficiently.

Motivation

Recent work in designing interfaces for immersive virtual environments attempts to apply 2D techniques to 3D worlds. However, there is a dearth of empirical study into how best to implement these interfaces; indeed, most designs seem to arise from simple intuition. As has been done for desktop systems, we need to rigorously explore the different characteristics that make up these interfaces, in order to elicit optimal user performance. Our work hopes to define and compare the characteristics that may be used to improve IVE interfaces.

The Haptic Augmented Reality Paddle (or HARP) system is a testbed we have designed to take advantage of bimanual interaction, proprioception, and passive-haptic feedback (Figure 1). This system allows us to compare many characteristics that may be helpful for IVE user interfaces. The HARP system uses a 2D window, called the work surface, for displaying interface widgets. The user selects widgets using the index finger of the dominant hand, as in [1], or a stylus, as in [3]. The work surface can be world fixed [7], or hand held [17, 20]. Finally, the work surface can be registered with a physical surface [2], or not [1]. Unlike others, our system does not provide support for a specific application, but rather serves as a testbed for comparing low-level interaction tasks. Our research attempts to provide some guidelines for designers of IVE interfaces.

Figure 1: The HARP System

EXPERIMENTAL METHOD

This section describes the experimental design used in the first empirical studies conducted with the HARP system testbed. These experiments were designed to compare interfaces that combine the presence or absence of passive-haptic feedback (i.e. a physical work surface) with hand-held and world-fixed windows. We use quantitative measures of proficiency, such as mean task completion time and mean accuracy, as well as qualitative measures, such as user preference, to compare the interfaces. Two experiments, one involving an open-loop task, and one involving a closed-loop task, were administered. In the interest of space, we present them together.

Experimental Design

These experiments were designed using a 2×2 within-subjects approach, with each axis representing one independent variable. The first independent variable was whether the technique used hand-held (H) or world-fixed (W) windows. The second independent variable was the presence (P) or absence (N) of passive haptic feedback.

Four different interaction techniques (treatments) were implemented which combine these two independent variables into a 2×2 matrix, as shown in Table 1.

	Hand-Held (H)	**World-Fixed (W)**
Passive Haptics (P)	HP Treatment	WP Treatment
No Haptics (N)	HN Treatment	WN Treatment

Table 1: 2×2 Design

Each quadrant is defined as:

HP = Hand-Held Window, with Passive-Haptics.
WP = World-Fixed Window, with Passive-Haptics.
HN = Hand-Held Window, No Haptics.
WN = World-Fixed Window, No Haptics.

For the **HP** treatment, subjects held a paddle-like object in the non-dominant hand (Figure 2), with the work surface defined to be the face of the paddle. The rectangular work surface measured 23cm $\times$ 17cm (W $\times$ H). The paddle handle radius was 2.8cm, and the handle length was 12.5cm. Subjects could hold the paddle in any position that felt comfortable, but that allowed them to accomplish the tasks quickly and accurately. Subjects were presented with a visual avatar of the paddle that matched exactly the physical paddle in dimension (Figure 3). For the **WP** treatment, a panel with the same dimensions as the work surface of the **HP** treatment was mounted on a rigid, floor-standing mounting frame in front of the dominant-hand side of the body of the subject. The panel was mounted on a rigid Styrofoam box attached to the surface of the mounting frame. When the subjects explored the panel with their hands, they was supposed to get the impression that it was "floating" in space in front of them. This matched the visual feedback, which was an avatar of the panel floating in front of the subject.

Figure 2: The Physical Paddle

Before the experiment began, each subject was asked at which height the panel should be mounted, and this remained fixed for the duration of the experiment. Each subject was free to move the chair to a comfortable location before each task. For the **HN** treatment, the subjects held only the handle of the paddle in the non-dominant hand (no physical paddle head), while being presented with a full paddle avatar. Again, subjects were free to hold the paddle in any position that allowed them to work quickly and accurately. The **WN** treatment was exactly the same as **WP**, except that there was no physical panel mounted in front of the subject.

Using a diagram-balanced Latin squares approach, four different orderings of the treatments were defined, and subjects were assigned at random to one of the four orderings. We had each subject perform 20 trials on two separate tasks for each treatment. Four different random orderings for the 20 trials were used. The subjects were seated during the entire experiment.

Each subject performed two tasks (experiments) using the treatments. Task one was a docking task. Subjects were presented with a colored shape on the work surface, and had to slide it to a black outline of the same shape in a different location on the work surface, and then release it (Figure 3). Subjects could repeatedly adjust the location of the shape until they were satisfied with its proximity to the outline shape, and then move on to the next trial by pressing a "Continue" button, displayed in the center at the lower edge of the work surface. This task was designed to test the component UI action of "Drag-and-Drop," which is a continuous task. The trials were a mix between horizontal, vertical, and diagonal movements.

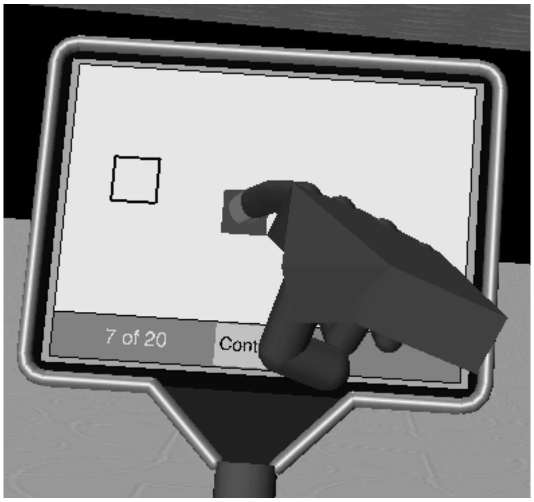

Figure 3: The Docking Task

The second task was a shape selection task. For this task, a signpost was displayed in the VE (Figure 4), upon which one shape was chosen at random to be displayed. For the right-handed subjects, the signpost was positioned in front and to the left of the subject. For the left-handed subjects, it was positioned in front and to the right of the subject. In addition, four shapes were arranged horizontally on the work surface, one of which matched the shape and color of the one on the signpost. The subject had to select the shape that matched the one on the signpost, and then press the "Continue" button to move on to the next trial. The subject could change the selection before moving to the next trial. This task was designed to test the component UI action of "Button Press," which is a discrete task.

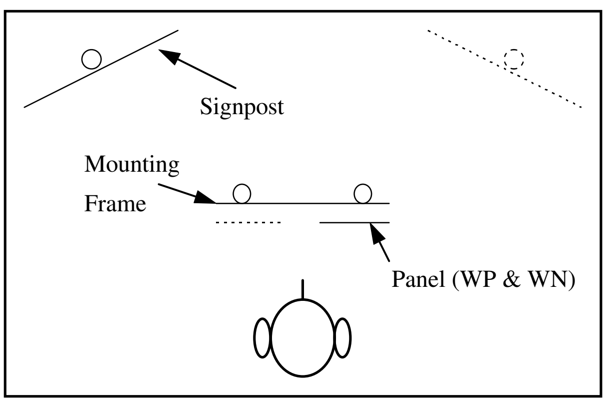

Figure 4: Overhead View of Physical Layout (dashed lines denote object positions for left-handed subjects)

The subject was required to press the "Continue" button after each trial for several reasons. First, this provided a clear distinction for when the trial was over. Subjects had to actively signal that they were through with the trial, so mistakes could be avoided because they could make adjustments before continuing on to the next trial. Second, this forced the user to return to a known "home" position prior to each trial, eliminating timing differences that could have arisen because of trial order. If the target location for one trial was close to the start position of the next trial, and subjects were not required to begin the trial at a home position, then they could acquire the shape for the next trial more quickly than for trials where the target and start position for successive trials were further apart. Finally, this gave a clear cut event which signaled the end of one trial and the start of the next, which is necessary for timing purposes.

Shape Manipulation

Five different shapes were used for these experiments: a circle, a square, a diamond, a triangle, and a five-pointed star. In addition, each shape could appear in any one of three colors: red, green, or blue. The bounding box used for intersection testing was the same for all shapes, so the only difference was their shape in the VE; each one was as easy to select as every other one.

Subjects selected shapes simply by moving the fingertip of their dominant-hand index finger to intersect the shape. A shape was released by moving the finger away from the shape, so that the fingertip no longer intersected it. For movable shapes (docking task), this required the subject to lift (or push) the fingertip so that it no longer intersected the virtual work surface, as moving the finger tip along the plane of the work surface translated the shape along with the fingertip. For immovable objects (selection task), the subjects were free to move the fingertip in any direction in order to release the object. Once the fingertip left the bounding box of the shape, the shape was considered released.

System Characteristics

The HARP software was running on a two-processor SiliconGraphics (SGI) Onyx workstation equipped with a RealityEngine^2 graphics subsystem, two 75MHz MIPS R8000 processors, 64 megabytes of RAM, and 4 megabytes of texture RAM. Because of a lack of audio support on the Onyx, audio feedback software (see below) was run on an SGI Indy workstation, and communicated with the HARP system over Ethernet. The video came from the Onyx, while the audio came from the Indy. We used a Virtual I/O i-glasses HMD to display the video and audio, with a Logitech ultrasonic tracker mounted on the front to track 6-DOF head motion. For the index-finger and paddle, we used an Ascension Flock-of-Birds magnetic tracker. The mounting stand for the panel was constructed using only wood and PVC tubing, so as to avoid introducing noise to the magnetic trackers. The work space was calibrated once, and the computed values were used for all subsequent runs of the software. All the software ran in one Unix thread. A minimum of 11 frames per second (FPS) and a maximum of 16 FPS were maintained throughout the tests, with the average being 14 FPS.

Subject Demographics

A total of 32 unpaid subjects were selected on a first-come, first-served basis, in response to a call for subjects. Most of the subjects were college students (20), either undergraduate (8) or graduate (12). The rest (12) were not students. The mean age of the subjects was 27 years, 5 months. In all, 30 of the subjects reported they used a computer with a mouse at least 10 hours per week, with 22 reporting computer usage exceeding 30 hours per week. Three subjects reported that they used their left hand for writing. 15 of the subjects were female and 17 were male. 19 subjects said they had experienced some kind of "Virtual Reality" before. Each subject passed a test for colorblindness. 15 subjects reported having suffered from motion sickness at some time in their lives, when asked prior to the experiment.

Protocol

Each subject signed an "Informed Consent for Human Subjects" form, and was given a copy to keep. Before beginning the actual experiment, demographic information was collected. The user was then fitted with the dominant-hand index finger tracker, and asked to adjust it so that it fit snugly. The user then chose between two different heights for the mounting position of the world-fixed work surface. Six subjects chose to use the higher mounting location of the panel (103cm above the floor) and 26 chose the lower position (94cm). The subjects were free to move the chair forward or back during the experiment. The chair surface was 46cm from the floor. Each subject was read a general introduction to the experiment, explaining what the user would see in the virtual environment, which techniques they could use to manipulate the shapes in the environment, how the paddle and dominant-hand avatars mimicked the motions of the subject's hands, and how the HMD worked.

After fitting the subject with the HMD, the software was started, the visuals would appear, and the audio emitted two sounds. The subjects were asked if they heard the sounds at the start of each task. To help subjects orient themselves, they were asked to look at certain virtual objects placed in specific locations within the VE. A blue cube was stationed on the ground plane of the VE at approximately 45° to the left of, and 3 meters away from, the subject. A green cone was placed at 45° to the right of, and 3 meters away from, the subject. Subjects were told that if they turned their head to the left, they should see a blue cube, and the same for the green cone to the right.

The subjects' location within the VE was such that they were in the center of a horizontal plane, texture-mapped with a beige, repeating pattern. Above the subject was a sky plane, which was texture-mapped with a blue sky and clouds. The subject was told to look up to see the blue sky, and to look down to see the patterned ground. This sequence of having the subject look left, right, up, and down was done before each task during the experiment, in order to orient the user each time.

At the beginning of the first task, the subject was instructed to move their dominant hand into the field of view, and that they would see the hand avatar (Figure 5). After moving their hand around for a few moments to get used to the mapping of hand movements to avatar movements, for the **H** treatments they were then asked to hold out their non-dominant hand, into which the paddle was placed, and they were given a few moments to get used to its movement. For the **W** treatments, it was pointed out that the panel in front of them was the panel that had been described in the introduction.

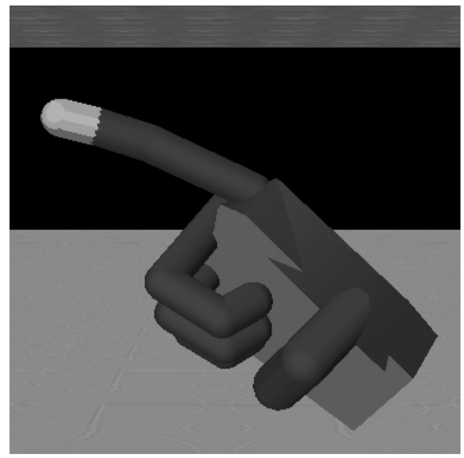

Figure 5: The Dominant-Hand Avatar

The work surface displayed the message, 'To begin the first trial, press the "Begin" button.' Subjects were asked to press the "Begin" button on the work surface by touching it with their finger. Five practice trials were then given, during which subjects were read a verbal description of the task they had to perform within the IVE. Each subject was coached as to how best to manipulate the shapes for each specific treatment.

After the practice trials, the subject was asked to take a brief rest, and was told that when ready, 20 more trials would be given, and would be scored in terms of both time and accuracy. It was made clear to the subjects that neither time nor accuracy was more important than the other, and that they should try to strike a balance between the two. Trial time for both tasks was measured as the total time between successive presses of the "Continue" button. Accuracy for the docking task was measured by how close the center of the shape was placed to the center of the target position, and for the selection task, accuracy was simply whether the correct shape was selected from among the four choices. After each treatment, the HMD was removed, the paddle was taken away (for **H**), and the subject was allowed to relax as long as they wanted to before beginning the next treatment.

Additional Feedback

In addition to visual and (in some cases) haptic feedback, the HARP system provided other cues for the subject, regardless of treatment. First, the tip of the index finger of the dominant-hand avatar was colored yellow (Figure 6a). Second, in order to simulate a shadow of the dominant hand, a red drop-cursor, which followed the movement of the fingertip in relation to the plane of the paddle surface, was displayed on the work surface (Figure 6b). The location of the drop-cursor was determined by dropping a perpendicular from the fingertip to the work surface, and drawing the cursor centered at that location. When the fingertip was not in the space directly in front of the work surface, no cursor was displayed. To help the subjects gauge when the fingertip was intersecting UI widgets, each widget became highlighted, and an audible CLICK! sound was output to the headphones worn by the subject (Figure 6c). When the user released the widget, it returned to its normal color, and a different UNCLICK! sound was triggered.

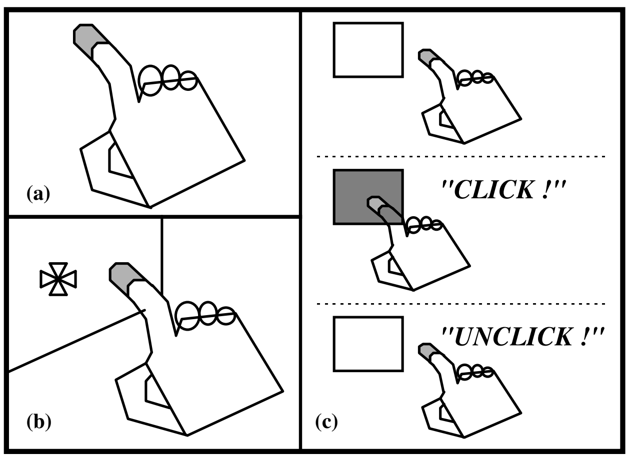

Figure 6: Manipulation Cues:
(a) Yellow Fingertip; (b) Red Drop-Cursor,
(c) Widget Highlighting and Audio Feedback

Data Collection

Qualitative data was collected for each treatment using a questionnaire. Four questions, arranged on Likert scales, were administered to gather data on perceived ease-of-use, arm fatigue, eye fatigue, and motion sickness, respectively. The questionnaire was administered after each treatment.

Quantitative data was collected by the software for each trial of each task. This data varied for the two tasks. For the docking task, the start position, target position, and final position of the shapes were recorded. In addition, the total trial time and the number of times the subject selected and released the shape for each trial was recorded. For the selection task, the total trial time, number of selections made for each trial, the correct answer, and the answer given by the subject were recorded.

Results

In order to produce an overall measure of subject preference for the four treatments, we have computed a composite value from the qualitative data. This measure is computed by averaging each of the Likert values from the four questions posed after each treatment. Because "positive" responses for the four characteristics were given higher numbers, on a scale between one and five, the average of the ease-of-use, arm fatigue, eye fatigue, and motion sickness questions gives us an overall measure of preference. A score of 1 would signify a lower preference than a score of 5. Table 2 shows the mean values for each question, as well as the composite value, for all four treatments ($N = 32$, standard-deviations in parentheses).

	Preference Factor				
	Ease of Use	**Arm Fatigue**	**Eye Fatigue**	**Motion Sickness**	**Composite Value**
HP	4.53 (0.671)	4.13 (0.942)	4.28 (0.851)	4.91 (0.296)	4.46 (0.445)
WP	4.28 (0.851)	3.19 (1.091)	4.25 (0.842)	4.84 (0.448)	4.14 (0.466)
HN	2.88 (1.070)	3.44 (0.914)	4.03 (0.999)	4.72 (0.683)	3.76 (0.638)
WN	2.84 (1.019)	2.25 (0.803)	3.78 (1.039)	4.75 (0.622)	3.41 (0.487)

Table 2: Mean Subject Preference by Treatment
(standard-deviations in parentheses)

In terms of the main effects, a within-subjects, 2×2 analysis of variance (ANOVA) for the composite value of subject preference showed that subjects significantly preferred **H** over **W** ($f = 23.02$; $df = 1/31$; $p < 0.001$) and **P** over **N** ($f = 86.178$; $df = 1/31$; $p < 0.001$). There was no interaction effect ($f = 0.09$; $df = 1/31$; $p > 0.760$).

An analysis of the mean trial completion time measure (in seconds) for the docking task yields the following descriptive statistics ($N = 32$ for all treatments): **HP** ($m = 6.71$; $sd = 2.582$), **WP** ($m = 6.60$; $sd = 3.284$), **HN** ($m = 12.78$; $sd = 5.832$), and **WN** ($m = 10.98$; $sd = 4.495$). In terms of the main effects, a within-subjects, 2×2 ANOVA showed that subjects performed significantly faster on **W** than **H** ($f = 6.63$; $df = 1/31$; $p < 0.050$) and on **P** than **N** ($f = 67.61$; $df = 1/31$; $p < 0.001$). There was also an interaction effect ($f = 9.42$; $df = 1/31$; $p < 0.005$). Figure 7 shows a box plot of the trial times by treatment, where the boxes represent the middle 50% of the values, the thick line represents the median, and the whiskers represent lines to the highest and lowest values. A comparison of the means tells us that subjects had particular manipulation problems with hand-held windows when no passive-haptic feedback was present. We postulate that precisely coordinating two movable objects (i.e. the paddle and the dominant hand) in free-space proved challenging for subjects, but that when either the user only had to manipulate the dominant hand (**W**), or when the physical work surface constrained the motion of the dominant hand (**P**), subjects were more capable of working efficiently.

Recall that accuracy on the docking task was measured as the distance the center of the shape was from the center of the target location at the end of the trial; lower values being better. Looking at the mean trial accuracy measure (in centimeters) we see (N = 32 for all treatments): **HP** (m = 0.15; sd = 0.079), **WP** (m = 0.17; sd = 0.072), **HN** (m = 0.25; sd = 0.194), and **WN** (m = 0.28; sd = 0.205). In terms of the main effects, a within-subjects, 2 × 2 ANOVA showed that subjects were significantly more accurate on **H** than **W** (f = 6.16; df = 1/31; p < 0.050) and on **P** than **N** (f = 17.87; df = 1/31; p < 0.001). There was no interaction effect (f = 0.29; df = 1/31; p > 0.590). Figure 8 shows a box plot of the accuracy by treatment. Intuitively, it comes as little surprise that **P** allowed subjects to be more precise than **N**, as **P** provided physical support for fine motor control. Also, because **H** allowed subjects to hold the work surface in any position that was comfortable, subjects seemed to choose positions that allowed them to steady their hands, whereas with **W**, subjects were forced to work in a fixed location, similar to using a touch-screen interface.

For the selection task, the statistics for the trial time measure (in seconds) were as follows (N = 32 for all treatments): **HP** (m = 2.83; sd = 0.791), **WP** (m = 3.49; sd = 0.614), **HN** (m = 3.35; sd = 0.736), and **WN** (m = 4.31; sd = 0.876). In terms of the main effects, a within-subjects, 2 × 2 ANOVA showed that subjects performed significantly faster on **H** than **W** (f = 34.42; df = 1/31; p < 0.001) and **P** than **N** (f = 35.50; df = 1/31; p < 0.001). There was a slight interaction effect (f = 3.96; df = 1/31; p > 0.050). Figure 9 shows a box plot of the trial times by treatment. Recall that the signpost was positioned outside the field of view of the subject when looking straight ahead. Because the subject could position the work surface in line with the signpost for **H**, these treatments were significantly faster than **W**. Accuracy on this task was measured as whether the correct choice was selected by the subject. For this measure, all treatments had a mean score of at least 99%, and no significance was found.

Discussion

The addition of passive-haptic feedback to IVE interfaces can significantly decrease the time necessary to perform UI tasks. Subjects performed 44% faster on the docking task, and 17% faster on the selection task when passive-haptic feedback was present. Interfaces which implement a 2D pen-and-tablet metaphor within 3D worlds can provide better support for both precise and ballistic actions by registering a physical surface with the virtual tablet. Docking task accuracy was significantly better, increasing by over 38%, when passive-haptic feedback was present. Furthermore, we have shown that hand-held windows

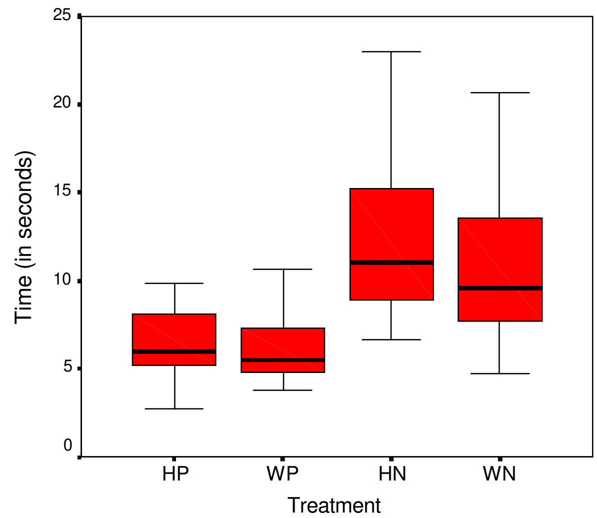

Figure 7: Docking Task Trial Time × Treatment

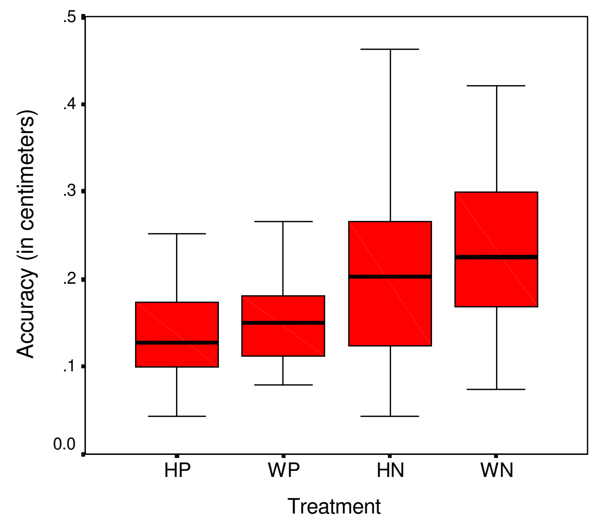

Figure 8: Docking Task Accuracy × Treatment

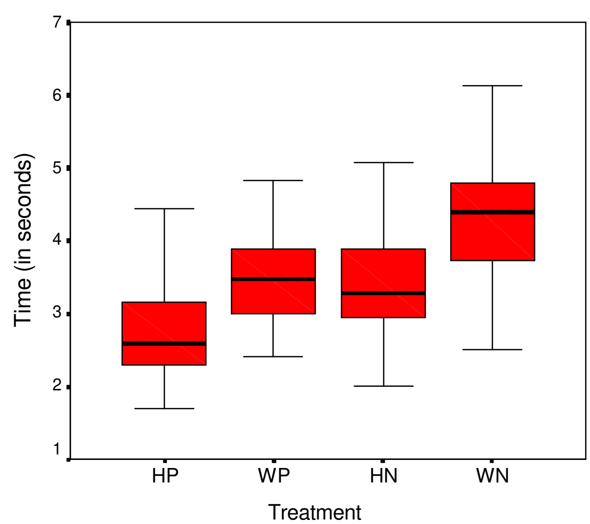

Figure 9: Selection Task Trial Time × Treatment

provide the freedom of movement necessary for working effectively in IVEs. On the selection task, which required looking around the IVE, subjects performed 21% faster using hand-fixed as opposed to world-fixed windows.

These quantitative findings are in line with our qualitative results. Users prefer interfaces that allow them to work efficiently and effectively. The use of passive-haptic feedback, coupled with a hand-held device, can greatly aid interaction in immersive virtual environments.

During our analysis, we found that some learning effects were present. Specifically, for the docking task, **HN** and **WN** trial times improved and accuracy increased over time, while **HP** and **WP** stayed fairly constant. This suggests that the **P** treatments presented subjects with the feedback they expected from real-world experience (i.e. that objects don't simply pass through other objects), and that the **N** treatments required more training. All treatments showed a general improvement trend over time on the selection task.

CONCLUSIONS AND FUTURE WORK

Our results show that the addition of passive-haptic feedback for use in precise UI manipulation tasks can significantly increase user performance. In addition, users prefer interfaces that provide a physical surface, and that allow them to work with UI widgets in the same visual field of view as the objects they are modifying.

In our work, we have tried to provide data to help IVE designers produce interfaces that allow users to perform real work. Because of the complexity of user interaction in IVEs, much work still needs to be done. We have shown that the HARP testbed is an effective environment for performing future studies. We will look at ways of improving non-haptic interfaces for those systems where it is impractical to provide passive-haptic feedback. Possible modifications include the use of 3D representations of widgets instead of 2D representations, and the imposition of simulated physical surface constraints by clamping user movement to the virtual surface of the paddle. Also, we would like to explore other component interaction techniques, such as cascading menus, within the HARP testbed.

ACKNOWLEDGEMENTS

This research was supported in part by the Office of Naval Research. Thanks to Jim Templeman for his insight.

REFERENCES

1. Angus, I., Sowizral, H., "VRMosaic: Web Access from within a Virtual Environment," *IEEE Computer Graphics and Applications*, 16, 3, (1996), pp. 6-10.
2. Billinghurst, M., Baldis, S., Matheson, L., Philips, M., "3D Palette: A Virtual Reality Content Creation Tool," *Proc. of VRST '97*, (1997), pp. 155-156.
3. Bowman, D., Wineman, J., Hodges, L., "Exploratory Design of Animal Habitats within an Immersive Virtual Environment," *GA Institute of Technology GVU Technical Report GIT-GVU-98-06*, (1998).
4. Bryson, S., Levit, C., "The Virtual Windtunnel: An Environment for the Exploration of Three-Dimensional Unsteady Flows," *Proc. of Visualization '91*, (1991), pp. 17-24.
5. Cutler, L., Fröhlich, B., Hanrahan, P., "Two-Handed Direct Manipulation on the Responsive Workbench," *1997 Symp. on Interactive 3D Graphics*, Providence, RI, (1997), pp. 107-114.
6. Deering, M., "The HoloSketch VR Sketching System," *Comm. of the ACM*, 39, 5, (1996), pp. 55-61.
7. Feiner, S., MacIntyre, B., Haupt, M., Solomon, E., "Windows on the World: 2D Windows for 3D Augmented Reality," *Proc. of UIST '93*, (1993), pp. 145-155.
8. Fels, S., Hinton, G., "Glove-TalkII: An Adaptive Gesture-to-Format Interface," *Proc. of SIGCHI '95*, (1995), pp. 456-463.
9. Fisher, S., McGreevy, M., Humphries, J., Robinett, W., "Virtual Environment Display System," *1986 Workshop on Interactive 3D Graphics*, Chapel Hill, NC, (1986), pp. 77-87.
10. Fuhrmann, A., Löffelmann, H., Schmalstieg, D., Gervautz, M., "Collaborative Visualization in Augmented Reality," *IEEE Computer Graphics and Applications*, 18, 4, (1998), pp. 54-59.
11. Mine, M., Brooks, F., Séquin, C., "Moving Objects in Space: Exploiting Proprioception in Virtual-Environment Interaction," *Proc. of SIGGRAPH '97*, (1997), pp. 19-26.
12. Mine, M, "ISAAC: A Meta-CAD System for Virtual Environments," *Computer-Aided Design*, 29, 8, (1997), pp. 547-553.
13. Poupyrev, I., Weghorst, S., Billinghurst, M., Ichikawa, T., "A Framework and Testbed for Studying Manipulation Techniques for Immersive VR," *Proc. of VRST '97*, (1997), pp. 21-28.
14. Shneiderman, B., *Designing the User Interface*, (3rd Edition), Addison-Wesley: Reading, MA, (1998).
15. Sowizral, H., "Interacting with Virtual Environments Using Augmented Virtual Tools," *Stereoscopic Displays and Virtual Reality Systems*, Fisher, S., Merrit, J., Bolas, M., Eds. Proc. SPIE 2177, (1994), pp. 409-416.
16. Stanney, K., "Realizing the Full Potential of Virtual Reality: Human Factors Issues that Could Stand in the Way," *Proc. of VRAIS '95*, (1995), pp. 28-34.
17. Stoakley, R., Conway, M., Pausch, R., "Virtual Reality on a WIM: Interactive Worlds in Miniature," *Proc. of SIGCHI '95*, (1995), pp. 265-272.
18. Sturman, D., Zeltzer, D., Pieper, S., "Hands-on Interaction With Virtual Environments," *Proc. of UIST '89*, (1989), pp. 19-24.
19. van Teylingen, R., Ribarsky, W., and van der Mast, C., "Virtual Data Visualizer," *IEEE Transactions on Visualization and Computer Graphics*, 3, 1, (1997), pp. 65-74.
20. Wloka, M., Greenfield, E., "The Virtual Tricorder: A Uniform Interface for Virtual Reality," *Proc. of UIST '95*, (1995), pp. 39-40.