Hand-Held Windows: Towards Effective 2D Interaction in Immersive Virtual Environments

Robert W. Lindeman John L. Sibert James K. Hahn
Institute for Computer Graphics
The George Washington University, Washington, DC
[gogo | sibert | hahn]@seas.gwu.edu

Abstract

The study of human-computer interaction within immersive virtual environments requires us to balance what we have learned from the design and use of desktop interfaces with novel approaches to allow us to work effectively in three dimensions. While some researchers have called for revolutionary interfaces for these new environments, devoid of two-dimensional (2D) desktop widgets, others have taken a more evolutionary approach. Windowing within immersive virtual environments is an attempt to apply 2D interface techniques to three-dimensional (3D) worlds. 2D techniques are attractive because of their proven acceptance and widespread use on the desktop. With current methods of performing 2D interaction in immersive virtual environments, however, it is difficult for users of 3D worlds to perform precise manipulations, such as dragging sliders, or precisely positioning or orienting objects.

We have developed a testbed designed to take advantage of bimanual interaction, proprioception, and passive-haptic feedback. We present preliminary results from an empirical study of 2D interaction in 3D environments using this system. We use a window registered with a tracked, physical surface, to provide support for precise manipulation of interface widgets displayed in the virtual environment.

1. Introduction

Following the initial enthusiasm about how Virtual Reality was going to radically change the way people interact with computers (and each other), researchers have now started to do some rigorous investigation into the nature of interaction in Immersive Virtual Environments (IVEs). Specifically, what we want to know is, given what we have learned over the past few decades about human-computer interaction in a basically two-dimensional domain, how can we best apply this knowledge to the design of user interfaces for three-dimensional worlds? How can we make the transition from 2D to 3D as painless as possible for users?

If we can provide a framework for users of desktop environments to easily transition to immersive environments, we can harness the experience they (and we, as researchers) have gained in desktop usage to allow people to quickly learn, and become productive within, immersive virtual environments. In our research, we combine recent work in bimanual interaction techniques with research into the use of passive-haptic feedback to provide an interface that allows precise manipulation of user interface (UI) widgets. The Haptic Augmented Reality Paddle (or HARP) system we have developed builds on recent work done in the use of tracked tablets for performing 2D tasks within 3D environments. The windows used within these systems move along with the user, do not unnecessarily occupy precious display real estate, and provide physical support for the work surface. Using this approach, we can effectively bring the desktop GUI into the IVE, without losing the hand support afforded by the mouse on the desktop.

In this paper, we first present related work and provide an overview of our system. Next, we describe some empirical research conducted using our system, and show how the results provide insight into interaction techniques in IVEs. Finally, we present areas for future work.

2. Related Work

Desktop systems typically use a combination of a keyboard and mouse to allow the user to interact with some kind of Window, Icon, Menu, Pointer (WIMP) interface. After a short learning period, users can become extremely proficient; able to perform precise, controlled movements, such as dragging sliders, or resizing windows. As computer interaction moves from 2D to 3D, we would like to take advantage of the physiological and psychological abilities of users and design a functionally equivalent but stylistically different interface for IVEs.

1087-8270/99 $10.00 © 1999 IEEE

Within these environments, where the user wears a Head-Mounted Display (HMD), use of a keyboard and mouse is sometimes not practical, because the user cannot see them. Also, the application might require the user to move around in physical space, which necessitates carrying the keyboard and mouse around. Finally, mapping 2D interaction devices and interface methodologies into 3D worlds can be sub-optimal and cumbersome for the user. Movement and manipulation in 3-space requires new approaches which allow users to perform tasks in a natural and effective way.

Some IVE applications have, therefore, 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, 19] or menus "floating" in space [14, 4, 5, 16, 20, 6, 15, 21]. The latter use either the user's finger or some sort of laser-pointer, combined with a physical button-click, to manipulate widgets. Using 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.

A further issue with the floating windows interfaces comes from the inherent problems of mapping a 2D interface into a 3D world. One of the reasons the mouse is so effective on the desktop, is that it is a 2D input device used to manipulate 2D (or 2.5D) widgets on a 2D display. Once we move these widgets to 3-space, the mouse is no longer tractable as an input device. 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, Wloka et al 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 [21]. These hand-relative window placement approaches strike a balance between incorporating the advantages of 2D window interfaces, and providing the necessary freedom for working in 3-space.

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, 3, 14]. These approaches register interface windows with a prop held in the non-dominant hand, and allow the user to interact with them using either her finger, or a stylus held in the dominant hand. One important aspect of these interfaces is their asymmetric use of the hands.

2.1. The Asymmetry of the Hands

Current physiology and psychology literature has advocated a move away from the traditional view that people are either right or left handed [10]. Instead, Guiard observed that most tasks we do in real life are accomplished using two hands, but that each hand assumes a different role. In discussing two hands as a *kinematic chain*, Guiard describes several relationships between the hands with regard to coordinated action [11]. First, the role of the non-dominant hand (ND) is not only to provide stability to the object acted upon by the dominant hand (D), but also to provide a reference frame for work done by D. Second, ND has a much coarser resolution of motion than D, and D can, therefore, successfully carryout actions requiring more precision. Third, ND actions have temporal precedence over D actions; the frame of reference must be set (ND) before precise actions are undertaken (D).

2.2. Proprioception

Recent work by Mine et al uses body-relative motions as an interaction technique, which takes advantage of the proprioceptive sense [14]. People have the ability to gauge movements of their hands relative to their own bodies. They describe three types of motion.

Direct Manipulation allows (possibly distant) objects to be manipulated as though they were in the user's hands. The technique automatically scales the world down to bring the grabbed object within the normal range of the hands. Manipulation then takes place in the familiar space of the user's own body. Releasing the object automatically returns the world to its former size.

Physical Mnemonics are 3D body-relative widgets. The researchers describe a pull-down menu which "hides" directly out of view above the user's head, and is accessed using a gesture. Another example is a scaling widget that allows the hands to work in concert to scale an object by moving the hands apart.

Gesture Commands are recognized by the system as having a specific meaning, such as the "over-the-shoulder deletion" gesture. To delete an object, a user simply throws it over her shoulder.

This approach shows the possibilities of working within arm's reach, and supports the notion of combining direct and symbolic manipulation into a single framework. Mine et al, however, point out that one of the major problems with current IVE interaction techniques is the lack of haptic feedback for precise movements.

2.3. Passive-Haptic Feedback

Passive-haptic "devices" are physical objects which provide feedback to the user simply by their shape, texture, or other inherent properties. In contrast to *active*-haptic feedback systems, the feedback provided by passive-haptic feedback devices is not controlled by a computer. These objects can be either rigid or deformable. Hinckley et al [12] introduce the use of passive-haptic "props" as interaction devices, and report rapid mastery of the interface. Their system was designed for the visualization of neurosurgical volumetric data. The user holds a doll's head in his non-dominant hand, and a small, rectangular piece of Plexiglas in his dominant hand. Both objects are tracked, and their positions and orientations are reported to the computer, which updates the visual display accordingly. The doll's head is used to control the orientation of the volumetric dataset on the display, and the Plexiglas is used to control a cutting plane through the data. This allows the user to easily and intuitively explore the dataset, because of the direct mapping of hand movements to virtual object movements.

3. Direct Versus Symbolic Manipulation

We can make a distinction between the types of interaction that IVE interfaces typically support. *Direct Manipulation* is a way of allowing the user to "grab" an object, manipulate it, and then release it. For example, a user might use a fist gesture to begin a manipulation, then rotate his fist, applying a one-to-one mapping of hand rotation to object rotation, and then open his hand to release the object. This is the type of manipulation used in [12]. A similar sequence could also be used for object translation and scaling. With *Symbolic Manipulation*, users manipulate UI widgets which in turn affect change on objects in the IVE. An example of this type of interaction is the use of RGB-sliders to control the color of an object. The user positions the individual sliders, changing the amount each component color contributes to the overall object color.

Intuitively, direct manipulation in IVEs seems more prone to instabilities in user motion (inexactness of user movement, fatigue), and therefore should not allow the user to precisely position the manipulated object. Direct manipulation should, however, be faster than symbolic manipulation, because of the directness of the mapping of hand movement to object movement. We can postulate that different tasks require differing levels of abstraction in order to provide optimal user efficiency. Tasks which have analogues in the real world, such as moving an object from one place to another, lend themselves well to direct manipulation. Tasks which require a high degree of precision, or which have no analogue in the real world, might possibly be best accomplished using symbolic techniques. Some researchers combine direct and symbolic techniques within a single system, depending on the nature of the manipulation [18, 15, 3, 14]. We focus on symbolic interaction techniques in this paper, and ground the symbolic manipulation widgets within 2D, hand-held windows.

4. Hand-Held Windows

Feiner et al introduce the notion of 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 he looks 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.

As discussed above, there has been much work lately in the area of bimanual interaction techniques. Two-handed interaction approaches suggest a class of special-purpose, object-fixed windows: hand-held windows. These windows are fixed relative to an object held in the non-dominant hand of the user, and provide many advantages. First, like view-fixed windows, hand-held windows move along with the user, so they are always within reach. Second, unlike view-fixed windows, 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 surface that the user carries around, and upon which UI widgets are drawn and manipulated [2, 3, 18]. This should provide the passive-haptic feedback necessary to carry out precise movements in IVEs.

These are the concepts incorporated into the testbed system we have developed, which takes advantage of bimanual interaction, proprioception, and passive-haptic feedback, using hand-held windows to allow users to perform efficient 2D interaction in IVEs.

5. The HARP System

The HARP system is used for symbolic manipulation, and is based on the three major characteristics described above: bimanual interaction, proprioception, and passive-haptic feedback.

5.1. System Description

The main feature of the HARP system is the use of a passive physical object, shaped like a paddle, as a surface for user interaction. The paddle head has a rectangular shape, with approximately the same dimensions as a common laptop screen (30cm diagonal), and a paddle grip that is roughly the same size as a Ping-Pong paddle handle.

Figure 1: Interacting with the HARP System

The user holds the paddle in his non-dominant hand, and uses the index finger of his dominant hand as a pointer (Figure 1). The VE contains a paddle avatar that matches the dimensions of the real paddle exactly (Figure 3). 2D UI widgets are drawn on the face of the virtual paddle. In addition, a model of a human hand in a pointing gesture is used to represent the actual dominant hand of the user.

One 6-DOF tracker is placed on the paddle, one on the index finger of the user's dominant hand, and one on the user's head. As the user moves the paddle through real space, the paddle avatar matches the real motion of the paddle. Similarly, movement of the pointing hand is matched by the pointing-hand avatar. The user's head motions are tracked so that in the visual image presented to the user, the paddle avatar and pointer avatar are registered with the actual paddle and dominant hand. Thus, because the avatars are registered with their real-world analogues, when the virtual hand touches the surface of the virtual paddle, the real hand contacts the real paddle.

5.2. Haptic Augmented Reality

The term Augmented Reality is typically used to describe a system where computer generated images are combined with real world images, in order to add information to the real world view [13]. We suggest that the use of real world *objects* in the haptic domain parallels the use of real world *images* in the visual domain, enhancing the user's perception of the real world. By holding a physical object in hand, the user is presented with more stimuli, providing higher fidelity. Also, because the virtual objects and real objects are registered, the stimuli are multimodal and complementary, providing enhanced feedback.

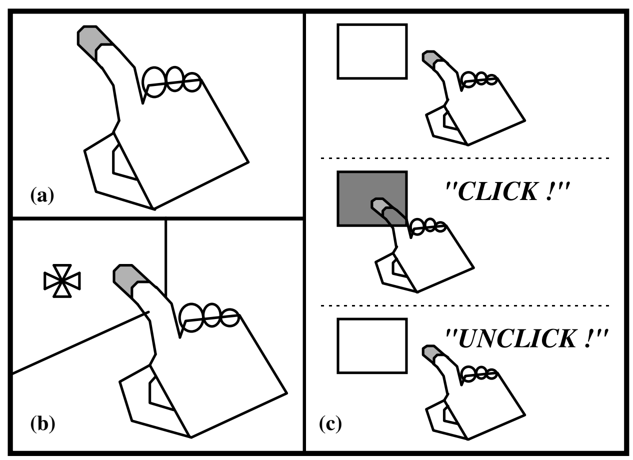

Figure 2: Manipulation Cues: (a) Yellow Fingertip; (b) Red Drop-Cursor; (c) Widget Highlighting and Audible Feedback

5.3. Additional Feedback

In addition to visual and haptic feedback, the HARP system provides other cues for the user. First, the tip of the index finger of the dominant-hand avatar is colored yellow (Figure 2a). Second, in order to simulate a shadow of the dominant hand, a red drop-cursor, which follows the movement of the fingertip in relation to the plane of the paddle surface, is displayed on the paddle (Figure 2b). The location of the drop-cursor is determined by dropping a perpendicular from the fingertip to the work surface, and drawing the cursor centered at that location. When the fingertip is not in the space directly in front of the work surface, no cursor is displayed. To help the subjects gauge when the fingertip is intersecting UI widgets, each widget becomes highlighted, by changing to a different color, and an audible CLICK! sound is output to the headphones worn by the user (Figure 2c). When the user releases the widget, it returns to its normal color, and a different UNCLICK! sound is triggered.

6. Empirical Study

We have designed several studies comparing the presence or absence of passive-haptic feedback, and bimanual versus unimanual interaction. This section describes the experimental design used in the first of these studies using the HARP system interface. Because of the dearth of empirical data on virtual environment interface techniques, we are conducting rigorous, basic research for use in the design of future immersive VE interfaces. We use measures of proficiency, such as mean task completion time and mean accuracy, to compare the interfaces.

6.1. UI Interaction Decomposition

Using what Shneiderman calls *Widget-Level* decomposition [17], we can look at the widgets that are defined in the system, and define types of actions based on possible manipulation of these widgets. In our current system, we have implemented buttons and sliders that can be configured to represent some typical UI widgets, requiring actions such as "drag-and-drop" and button presses. We define two distinct types of actions based on these widgets: *discrete actions* and *continuous actions*. Discrete (open-loop) actions involve a ballistic selection operation, such as clicking a toolbar icon, double clicking a filename, or positioning an input cursor. Continuous (closed-loop) actions include dragging sliders, using drag-and-drop to move a file, or accessing a cascading pull-down menu. We include tasks using each of these types of actions in our experimental design.

6.2. Experimental Design

This experiment was designed using a 2×2 within-subjects approach, with each axis representing one independent variable. Each subject performed two tasks using four interface techniques. The first independent variable was whether the technique used hand-held (**H**) or world-fixed (**W**) windows. The second independent variable was the presence or absence of passive-haptic feedback (Passive Haptics, or **P**, and No Haptics, or **N**, respectively). Four different interaction techniques (treatments) were implemented which combine these two independent variables into a 2×2 matrix (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, the subject held a paddle-like object in her non-dominant hand, 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. The subject could hold the paddle in any position that felt comfortable, but that allowed her to accomplish the tasks quickly and effectively. The subject was presented with a visual avatar of the paddle that matched exactly the physical paddle in dimension. 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 subject explored the panel with her hands, she was supposed to get the impression that it was "floating" in space in front of her. This matched the visual feedback, which was an avatar of the panel floating in front of the subject. Before the experiment began, the subject was asked at which height the panel should be mounted, and this remained fixed for the duration of the experiment. The subject was free to move the chair to a comfortable location before each task. For the **HN** treatment, the subject held only the handle of the paddle in her non-dominant hand (no physical paddle head), while being presented with a full paddle avatar. Again, the subject was free to hold the paddle in any position that allowed her 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 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. The subjects were seated during the entire experiment.

Task one was a docking task, where 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 release it (Figure 3). The subject could repeatedly adjust the location of the shape until he was satisfied with its proximity to the outline shape. After the subject was satisfied that the shape was close enough, he selected a "Continue" button, displayed in the center at the lower edge of the work surface, and was then presented with the next trial. This task was designed to test the component UI action of "Drag-and-Drop," which is a *continuous* task.

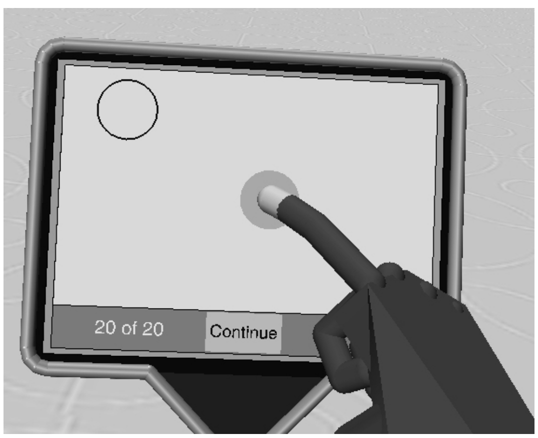

Figure 3: The Docking Task

The second task was a shape selection task. For this task, a signpost was displayed in the VE, 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 his 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.

6.3. Subjects

A total of 32 subjects were selected on a first-come, first-served basis, in response to a call for subjects. The mean age of the subjects was 27years, 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. All of the subjects were tested for colorblindness.

6.4. Shape Manipulation

Subjects selected shapes simply by moving the tip of their dominant-hand index finger to intersect the shape. A shape was released by moving the finger outside the bounding-box, so that the fingertip no longer intersected it. For the docking task, the user was required to lift the finger perpendicular to the work surface, as any movement parallel to the surface would translate the shape. For the selection task, moving the finger in any direction away from the shape would release it. Once the fingertip left the region of influence around the shape, the shape was considered released.

6.5. Preliminary Results

The primary measures of performance were mean task completion time and accuracy. For each subject, the data recorded over the 20 trials of a given task under a given treatment were averaged, and the mean scores used as data-points in statistical analyses. A total of 32 subjects performed each task under each treatment, giving a total of 32 data-points for time and 32 data-points for accuracy for each treatment. For the docking task, trial time was measured as the time from the moment the shape/outline pair appeared on the work surface to the time the "Continue" button was pressed. Accuracy was measured as the distance the shape was from the outline shape at the end of the trial. Shorter distances were better. For the selection task, trial time was measured as the time from the moment the shape appeared on the signpost to the time the "Continue" button was pressed. Accuracy was measured as whether or not the correct shape was selected at the end of the trial.

A factorial analysis of variance (ANOVA) was performed on the trial time and accuracy data for each of the two tasks. Table 2 shows the means (in seconds) for trial time for the docking task. Table 3 shows the means (in cm) for accuracy on the docking task. The means report the average distance the shape was from the target location at the end of the trial. Lower numbers are better, with a distance of 0cm meaning that the shape was placed exactly on the target. Table 4 shows means (in seconds) for trial time for the selection task. The means for the percent correctly selected on the selection task was 99% for each treatment, so no table is given. This result is due to the task being trivial.

In general, trial time on both tasks was significantly faster when passive-haptic feedback was present. Subjects performed 75% slower on the docking task, and 20% slower on the selection task when passive-haptic feedback was not present. Docking task accuracy was also significantly better with passive haptics, with accuracy dropping by over 60% when passive-haptic feedback was not present. For the docking task, there were mixed results in terms of the hand-held versus world-fixed treatments, with the hand-held treatments being significantly slower, but significantly more accurate than the world-fixed treatments. For trial time on the selection task, the hand-held treatments were significantly faster (15%). There was a slight interaction effect.

6.6. Discussion

Looking at these results suggests a number of interpretations. First, the tasks were different in at least two significant ways. The selection task only required a ballistic motion to complete the task (similar to throwing a dart), while the docking task required a fairly high degree of precision in order to successfully move the shape from its start position to the target position. Another difference between the tasks is that the docking task did not require the subjects to look at any object other than the work surface, while the selection task required the user to look at the signpost before making a choice from among the shapes on the work surface. This means that the location of the work surface was more significant for the selection task than for the docking task. Because the work surface for the world-fixed window treatments (**WP** & **WN**) was located directly in front of the subject, and because the signpost was positioned slightly off to one side, the subject was forced to turn her head to acquire the target shape, then turn back to make a selection. With the hand-held window treatments (**HP** & **HN**), the user could hold the work surface in the same field of view as the signpost. Thus, in this experiment, the hand-held windows were significantly faster for the tasks where the user had to look around, and passive-haptic feedback was important where precision was required.

Docking Task Trial Time (sec.)	Hand-Held (H)	World-Fixed (W)	Averages
Passive Haptics (P)	$m = 6.71^{**}$ ($sd = 2.58$) $N = 32$	$m = 6.60^{**}$ ($sd = 3.28$) $N = 32$	$m = 6.65^{**}$ ($sd = 2.93$) $N = 64$
No Haptics (N)	$m = 12.78$ ($sd = 5.83$) $N = 32$	$m = 10.98$ ($sd = 4.49$) $N = 32$	$m = 11.88$ ($sd = 5.24$) $N = 64$
Averages	$m = 9.74$ ($sd = 5.42$) $N = 64$	$m = 8.79$ ($sd = 4.48$) $N = 64$	$m = 9.27$ ($sd = 4.98$) $N = 128$
$^{**}p < .01$			

Table 2: Docking Task Means and Standard Deviations for Trial Time

Docking Task Accuracy (cm)	Hand-Held (H)	World-Fixed (W)	Averages
Passive Haptics (P)	$m = 0.15^{*}$ ($sd = 0.08$) $N = 32$	$m = 0.17^{*}$ ($sd = 0.07$) $N = 32$	$m = 0.16^{**}$ ($sd = 0.08$) $N = 64$
No Haptics (N)	$m = 0.25$ ($sd = 0.19$) $N = 32$	$m = 0.28$ ($sd = 0.21$) $N = 32$	$m = 0.26$ ($sd = 0.20$) $N = 64$
Averages	$m = 0.20$ ($sd = 0.16$) $N = 64$	$m = 0.22$ ($sd = 0.16$) $N = 64$	$m = 0.21$ ($sd = 0.16$) $N = 128$
$^{*}p < .05$ $^{**}p < .01$			

Table 3: Docking Task Means and Standard Deviations for Accuracy

Selecting Task Trial Time (sec.)	Hand-Held (H)	World-Fixed (W)	Averages
Passive Haptics (P)	$m = 2.83^{**}$ ($sd = 0.79$) $N = 32$	$m = 3.49^{**}$ ($sd = 0.61$) $N = 32$	$m = 3.16^{**}$ ($sd = 0.78$) $N = 64$
No Haptics (N)	$m = 3.35^{**}$ ($sd = 0.74$) $N = 32$	$m = 4.31$ ($sd = 0.88$) $N = 32$	$m = 3.83$ ($sd = 0.94$) $N = 64$
Averages	$m = 3.09^{**}$ ($sd = 0.80$) $N = 64$	$m = 3.90$ ($sd = 0.85$) $N = 64$	$m = 3.50$ ($sd = 0.92$) $N = 128$
$^{**}p < .01$			

Table 4: Selection Task Means and Standard Deviations for Trial Time

7. Conclusions and Future Work

The results of the study support the utility of hand-held windows as a general interaction tool. Hand-held windows move along with the user, and can be positioned to allow the user to work effectively. They do not clutter the user's view, unless they are currently being accessed, and take advantage of the proprioceptive sense. With the added support provided by passive-haptic feedback, hand-held windows can provide an effective interface for manipulating 2D widgets.

The benefits of hand-held windows are enhanced by the ability to use one window for interacting with many objects within an IVE. Rather than having many windows, we can imagine an environment where the user can select objects, and have any controls they possess appear on the surface of the paddle. Alternatively, we can envision an immersive VRML browser which queries objects in the vicinity of the user, and if they have any controls, they are displayed on the paddle.

In future studies, we will look at more types of actions, as well as different window types. Using the HARP system as a testbed, it would be interesting to compare view-fixed and hand-held windows on a pull-down menu type of task. In order to construct a more representative task, we will incorporate these techniques into a larger system, and gather real-world data.

Acknowledgements. This research was supported in part by the Office of Naval Research. We would also like to thank all the members of the Graphics and User-Interface groups at GW for all their support and feedback. Special thanks to Jim Templeman and Linda Sibert for their insight and comments.

8. 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," *University of Washington, Seattle, HIT Lab Technical Report R-97-23*, (1997).
3. Bowman, D., Wineman, J., Hodges, L., "Exploratory Design of Animal Habitats within an Immersive Virtual Environment," *Georgia 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," *Proceedings of Visualization '91*, (1991), pp. 17-24.
5. Cutler, L., Fröhlich, B., Hanrahan, P., "Two-Handed Direct Manipulation on the Responsive Workbench," *1997 Symposium on Interactive 3D Graphics*, Providence, RI, (1997), pp. 107-114.
6. Deering, M., "The HoloSketch VR Sketching System," *Communications 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," *Proceedings of UIST '93 ACM Symposium on User Interface Software and Technology*, (1993), pp. 145-155.
8. Fels, S., Hinton, G., "Glove-TalkII: An Adaptive Gesture-to-Format Interface," *Proceedings of ACM CHI '95 Conference on Human Factors in Computing Systems*, (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. Guiard, Y., "Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model," *Journal of Motor Behavior*, 19, 4, (1987), pp. 486-517.
11. Guiard, Y., "The Kinematic Chain as a Model for Human Asymmetrical Bimanual Cooperation," *Cognition and Action in Skilled Behaviour*, Colley, A., Beech, J., Editors, Elsevier Science Publishers B.V., North-Holland, (1988), pp. 205-228.
12. Hinckley, K., Pausch, R., Goble, J., Kassell, N., "Passive Real-World Interface Props for Neurosurgical Visualization," *Proceedings of the ACM CHI '94 Conference on Human Factors in Computing Systems*, (1994), pp. 452-458.
13. Milgram, P., Drascic, D., Grodski, J., Restogi, A., Zhai, S., Zhou, C., "Merging Real and Virtual Worlds," *Proceedings of IMAGINA '95*, (1995).
14. Mine, M., Brooks, F., Séquin, C., "Moving Objects in Space: Exploiting Proprioception in Virtual-Environment Interaction," *Proceedings of the ACM SIGGRAPH '97 Conference on Computer Graphics*, (1997), pp. 19-26.
15. Mine, M, "ISAAC: A Meta-CAD System for Virtual Environments," *Computer-Aided Design*, 29, 8, (1997), pp. 547-553.
16. Poston, T., Serra, L., "Dexterous Virtual Work," *Communications of the ACM*, 39, 5, (1996), pp. 37-45.
17. Shneiderman, B., *Designing the User Interface*, (Third Edition), Addison-Wesley: Reading, Massachusetts, (1998), pp. 60-61.
18. Stoakley, R., Conway, M., Pausch, R., "Virtual Reality on a WIM: Interactive Worlds in Miniature," *Proceedings of the ACM CHI '95 Conference on Human Factors in Computing Systems*, (1995), pp. 265-272.
19. Sturman, D., Zeltzer, D., Pieper, S., "Hands-on Interaction With Virtual Environments," *Proceedings of UIST '89 ACM Symposium on User Interface Software and Technology*, (1989), pp. 19-24.
20. van Teylingen, R., Ribarsky, W., van der Mast, C., "Virtual Data Visualizer," *IEEE Transactions on Visualization and Computer Graphics*, 3, 1, (1997), pp. 65-74.
21. Wloka, M., Greenfield, E., "The Virtual Tricorder: A Uniform Interface for Virtual Reality," *Proceedings of UIST '95 ACM Symposium on User Interface Software and Technology*, (1995), pp. 39-40.