The Haptic Hand: Providing User Interface Feedback with the Non-Dominant Hand in Virtual Environments

Luv Kohli Mary Whitton

Department of Computer Science
The University of North Carolina at Chapel Hill

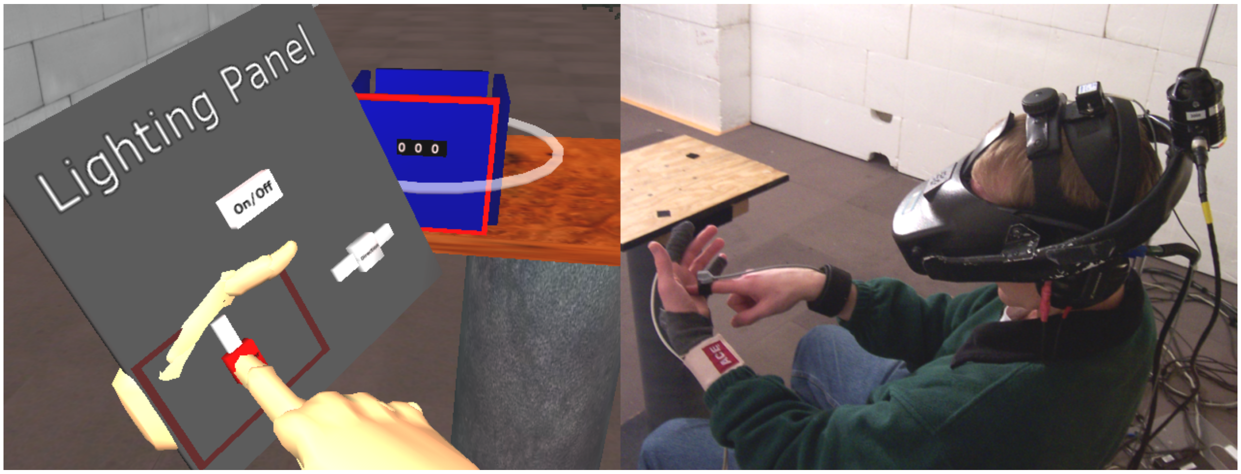

Figure 1 - A user interacts with a virtual interface panel by touching his non-dominant hand.

Abstract

We present a user interface for virtual environments that utilizes the non-dominant hand to provide haptic feedback to the dominant hand while it interacts with widgets on a virtual control panel. We believe this technique improves on existing prop-based methods of providing haptic feedback. To gauge the interface's effectiveness, we performed a usability study. We do not present a formal comparison with prior techniques here. The goal of this study was to determine the feasibility of using the non-dominant hand for haptic feedback, and to obtain subjective data about usability. The results demonstrated that the interface allowed users to perform precision tasks. These results have convinced us that this technique has potential and warrants further development.

Key words: 3D user interfaces, bimanual interaction, haptics, virtual environments, virtual reality

1 Introduction

Virtual objects do not naturally provide haptic (touch) feedback because they do not actually exist in the real world. Haptic feedback is important for interacting with the real world, and Insko, et al. showed that in virtual environments (VEs), passive haptics increase a user's sense of presence—the user's feeling that she is actually in the VE—as well as the user's spatial memory of the VE [10].

Haptic feedback is especially important for fine manipulation of real objects. Without the sense of touch, it is difficult to interact precisely with objects in VEs because there is nothing to steady the user's hands [13].

Guiard studied the distribution of work between the dominant and the non-dominant hands and classified tasks as *unimanual* (e.g., one-handed throwing), *bimanual symmetric* (identical actions performed by each hand), and *bimanual asymmetric* (both hands perform closely coordinated but different actions) [6]. Studies have shown that two-handed interaction techniques can provide significant advantages over one-handed techniques [1, 3]. These observations have been used in the development of user interfaces for VEs [5, 8, 9, 11]. The general idea is to provide a user interface that is held and/or coarsely adjusted by the non-dominant hand (NDH), and manipulated more finely by the dominant hand (DH). However, the lack of haptic feedback makes this fine manipulation difficult in VEs, much as it does in the real world.

In the case of virtual user interfaces that contain widgets (e.g., buttons and sliders), earlier work has shown that using physical props for haptic feedback helps the user interact with VEs more effectively. A common physical prop is a tablet or paddle which is held in the NDH [4, 11, 12, 14, 15]. In the VE, the geometric model of the paddle presents the interface panel with its widgets to the user. The user can move the paddle as necessary and interact with the widgets using the DH.

It is impractical to carry a physical paddle in cases where both hands are required to perform some task in a VE [12]. The question arises of what to do with the paddle when the NDH is needed to interact with another object in the environment. It is possible to hang the paddle down by one's side, but it can be cumbersome or tiring to carry it while interacting with the environment.

Another issue is that using multiple props for different tasks increases the number of objects that must be tracked, and there must be a sensible place to store these props if the user needs to switch between them often.

Exploiting Proprioception

Proprioception allows us to sense the position, orientation, and movement of our limbs, joints, and muscles. Research suggests that, independent of visual feedback, proprioceptive cues give humans a good sense of where their hands are relative to one another [7]. This idea has been used in the development of user interfaces for immersive VEs [13]. We propose exploiting proprioceptive cues by using the NDH itself to provide haptic feedback for a portable 2D user interface in an immersive VE (see Figure 1).

Implementation of this idea raises several issues. The principal issue is that properly tracking the user's hands is difficult. This is particularly important since the relative position of the two hands is critical for our technique. Because people have varying hand sizes and shapes, trackers cannot be mounted in exactly the same way on every person. This variability necessitates additional calibration steps. Another issue is that it is difficult to provide an accurate representation of the user's hands in a general-purpose VE without explicit per-person calibration and measurement. A third issue is that the surface of the hand is non-planar. Dragging an object along the surface of the hand may not feel smooth enough to satisfyingly emulate a planar surface.

In this paper, we present a system that uses the NDH as a haptic surface for a user interface in an immersive VE. We conducted an initial usability study to determine the feasibility of the technique. Results from the study are sufficiently positive to warrant further development.

2 Our Idea: The Haptic Hand

We have created a system in which a user interface is presented to the user on the plane fitted to the palm of the NDH. The interface panel contains widgets that are used to manipulate and modify objects in the VE. The idea is that touching the NDH palm will give both hands haptic feedback when manipulating these widgets with the DH. Although the work reported here does not include a direct comparison with other techniques, it is our hypothesis (to be tested in future work) that this method will be easier to use than a system with no haptic feedback and less cumbersome to use than systems that use physical paddles. Our goal for this work is to demonstrate feasibility and collect subjective, qualitative data about usability.

2.1 The Interface Panel

To test using the NDH to provide haptic feedback, we included widgets that control both discrete and continuous variables in our interface panel: buttons (discrete) and sliders (continuous) (Figure 2). In its initial state, the panel floats in the VE; it can easily be initialized to appear in a location conveniently within reach of the user.

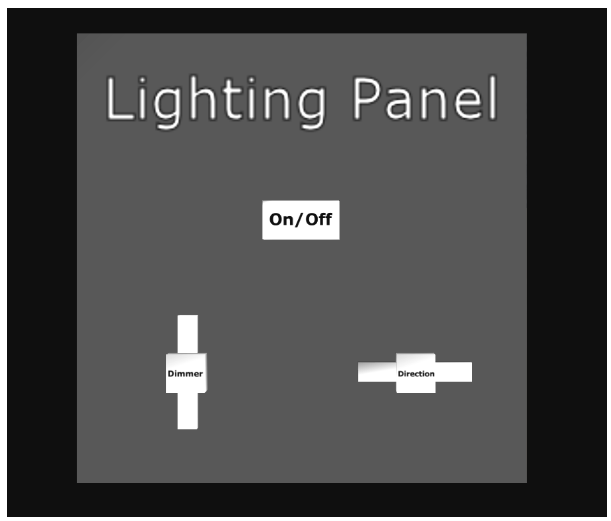

Figure 2 - A sample interface panel with buttons and sliders

2.2 Interacting with the VE with Hand Gestures

Our system assumes that, at a minimum, there are virtual representations (avatars) of the user's two hands. The NDH model must support at least two poses: open and closed. The DH model must support a pose with a pointing index finger. In our system with minimal tracking, these are the only supported poses of the hands. The user interacts with the interface panel through the use of two simple hand gestures: opening and closing the NDH. These gestures are used for grabbing, holding, and releasing the interface panel, activating widgets, and manipulating other objects in the environment.

2.3 Widget Activation

One of the issues with using the NDH for haptic feedback as opposed to a fixed-size (large) physical paddle is that there is a limited amount of real estate available on the palm. Unless tracking is very precise, this implies that many widgets cannot be active simultaneously. Consequently, we need a way to select which widget is currently active when the interface panel is larger than the surface of the palm.

In our initial implementation, which proved unsatisfactory, the panel was placed in the VE at some location near the user, and widget activation was performed by moving the open NDH behind the desired widget. Once activated, the user could manipulate the widget with the DH. Closing the NDH while intersecting the panel allowed the user to move and rotate the panel as desired to position it more comfortably. We quickly realized, however, that this particular implementation suffered from precisely the same problem that we were intending to tackle. Since there was no physical manifestation of the interface panel, there was no way for the user to steady the NDH. This made it extremely difficult and tiresome to keep a widget active, since the NDH would drift in and out of the widget's activation area.

Given this information, we took a new approach in which the panel is snapped to the NDH so that the panel and NDH move and rotate together in the VE. This way, an activated widget remains active until explicitly deactivated by the user, instead of being affected by the hand's natural tendency to drift when nothing steadies it. To achieve this, we reversed the effect the two hand gestures have on the panel and widgets (Figure 3a-e).

Initially, the panel is not attached to the NDH, and the two are not intersecting. To select which widget to activate, the user must close her NDH and move it so that it intersects the panel. At this point, the widget that is closest to the NDH is highlighted via a surrounding red rectangle. If the NDH is opened while a widget is highlighted, the panel snaps to the NDH such that the highlighted widget is centered on and oriented with the palm. The highlighted widget changes color to indicate that it has been activated and can be manipulated with the DH. The panel now moves and rotates with the NDH.

To activate a different widget the user closes the NDH, selects the desired widget, and opens the NDH, causing the new widget to snap to the center of the palm. The user can let go of the panel by closing the NDH and moving it away from the panel.

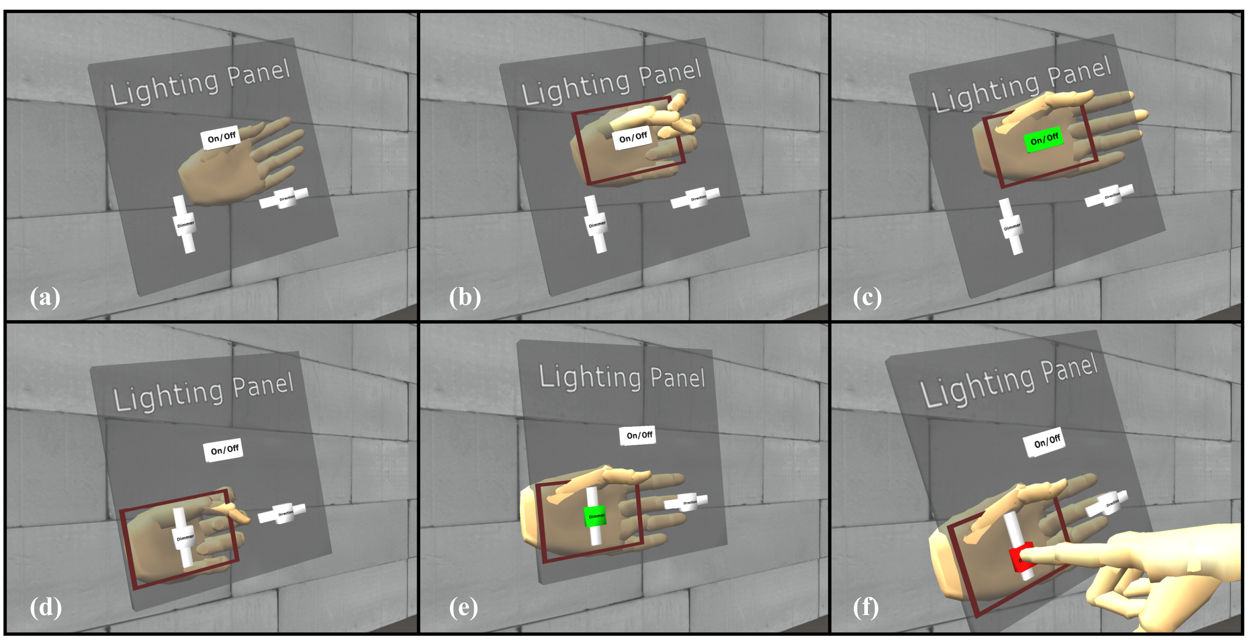

Figure 3 - Interface panel manipulation: (a) the panel is not attached to the NDH; (b) the NDH is closed and intersects the panel to highlight the desired widget; (c) the NDH is opened to snap the desired widget to the center of the palm; (d) the NDH is closed and moved to a new desired widget; (e) the NDH is opened and the new widget is snapped and activated; (f) the DH is used to manipulate the active widget.

2.4 Widget Manipulation

Once a widget has been activated, it is manipulated by bringing the DH index fingertip in contact with the widget on the NDH palm. An active widget will once again change color when it is being manipulated (Figure 3f). Buttons are pressed and released using discrete, ballistic motions, while sliders are changed by continuous dragging along the surface of the NDH.

2.5 Calibration

To snap the interface panel to the palm of the user's NDH, we must know where the surface of the palm is. We track the user's hands with Polhemus Fastrak magnetic trackers. Since we need to keep the surface of the palm unobstructed for haptic feedback, we mount a tracker rigidly on the back of the user's NDH using stretchable sports bandages (Figure 4). Because different people's hands are different in shape and size, the tracker will mount differently relative to the palm surface for each person. To determine where the user's palm is relative to the NDH tracker, we place the user's hand flat in a fixed pose on a pedestal. We then determine the transform between the NDH tracker and the fixed pose and apply this transform to all subsequent tracker readings.

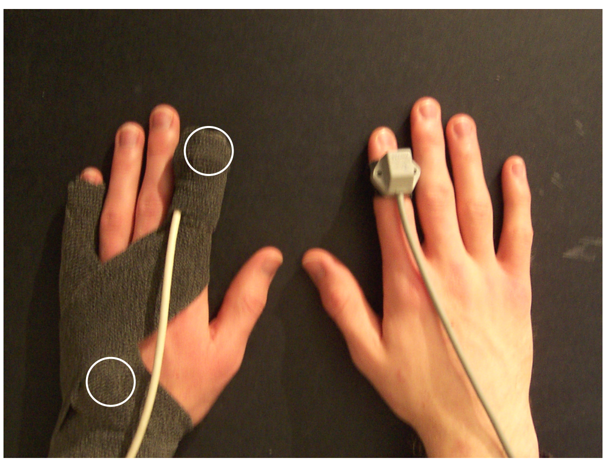

Figure 4 - The NDH (left) and DH (right) with trackers attached. The white circles indicate where the NDH trackers are located underneath the sports bandage: one on the fingertip to monitor whether the palm is open or closed, and one on the back of the hand to determine its pose.

The only other calibration we do is set the DH index finger tracker to 2cm from the index fingertip.

2.6 Manipulation Issues

Since the interface panel has been locked to the user's NDH palm, we determine when the user is manipulating a widget by detecting collisions between the user's virtual DH fingertip and the widget. This technique usually works well. However, since the surface of the palm is not truly planar, and the user can make any number of small motions that affect the shape of the palm, the approximated plane and the surface of the NDH palm do not always coincide. As a result, we may not always know when the DH index fingertip is in physical contact with the NDH palm. One consequence of this is that lifting the DH finger off the NDH palm does not immediately deactivate widgets. This works well for buttons because they require only ballistic input and have two discrete states: pressed or released. Sliders, however, are continuous input devices; lifting the DH finger off of the NDH palm may leave the slider still active until the virtual finger and slider are no longer colliding. In this situation, the user once again has no way to steady the DH finger, and the slider value changes from the intended value very easily.

To alleviate this problem, we implemented a slider locking mechanism. If the user holds the slider's value for at least 200ms, the value is locked. If the user removes the virtual finger from the slider within the next 500ms, the value of the slider snaps back to the stored locked value. If the virtual finger does not leave the slider within those 500ms, then the slider value will not snap back. During those 500ms, the slider behaves normally (i.e., values are updated). These time thresholds were chosen by trial and error, but they seem to work acceptably in practice. This slider locking mechanism is similar to techniques proposed for solving the same problem for touch-sensitive tablets. In particular, Buxton, et al. discuss keeping a short FIFO queue of tracking samples [2]. When the user lifts her finger off of the tablet, the oldest sample in the queue is used, and the queue is cleared. The length of this queue is determined through experimentation.

The problem of precise collision detection could also be solved using a hardware contact switch or pressure sensor on the DH fingertip to tell whether the fingertip and palm are physically in contact.

3 Usability Study

We ran a small usability study to test the idea of using the NDH as haptic feedback for user interfaces in immersive VEs. The goal was to determine if the technique is worth pursuing further. As such, we did not run a rigorous study comparing our technique with a physical paddle interface, nor did we make extensive performance measurements. We were more interested in finding out how difficult it is to learn to use the interface, and whether the technique allows both discrete and continuous interaction.

3.1 Materials and Methods

Participants

Eight people participated in this study (6 males, 2 females, 7 right-handed, 1 left-handed) and were paid $5 each for their time. Participants were a mix of graduate and undergraduate students of various majors. 3 participants had never been in a VE, 3 had been in a VE occasionally, and 2 had been in a VE many times.

Equipment

In addition to the Polhemus Fastrak magnetic hand trackers already mentioned, participants wore a Virtual Research Systems V8 head-mounted display (HMD). The head was tracked by a 3rdTech HiBall Wide Area Tracker. Participants were immersed in the VE for approximately 15 minutes and were seated the entire time.

3.2 Study Design

Our study consisted of a training session and a matching task. Before entering the VE, participants were shown the environment and interface panel on a desktop display and were told how to interact with the panel and with objects in the environment. The training session allowed participants to familiarize themselves with these interactions while immersed in the VE. Once participants felt comfortable, they were given a matching task to perform.

Training Session

The interface panel used for the training session controlled a red light in the environment (Figure 2). A button labeled “On/Off” toggled whether the light was on or off, and sliders labeled “Dimmer” and “Direction” controlled the intensity of the light and the lateral direction it came from, respectively. The training environment also included a four-sided object (essentially a cube without the top and bottom faces) sitting atop a pedestal in front of the participant. Each face of the cube was selectable by the DH index fingertip. The cube was surrounded by a ring that served as a handle to rotate the cube about its vertical axis. Intersecting the ring with a closed NDH and moving around the ring with the intersecting hand rotated the cube. This action required the participant to let go of the interface panel.

Matching Task

The participant’s goal in the matching task was to make two cubes look alike (see Figure 5). The cube on the right was the reference cube; each face was a different shade of blue, and had a different number from 0 to 100 on it. The interface panel for the matching task had two sliders (“Number” and “Color”) and two buttons (“Reset Face” and “Done”). The “Number” and “Color” sliders adjusted the number and color of the currently selected face, respectively. The “Reset Face” button reset the currently selected face to a default shade of blue and the number 0. The “Done” button checked whether the two cubes matched. Participants received audio cues indicating whether their attempt was successful upon pressing the “Done” button.

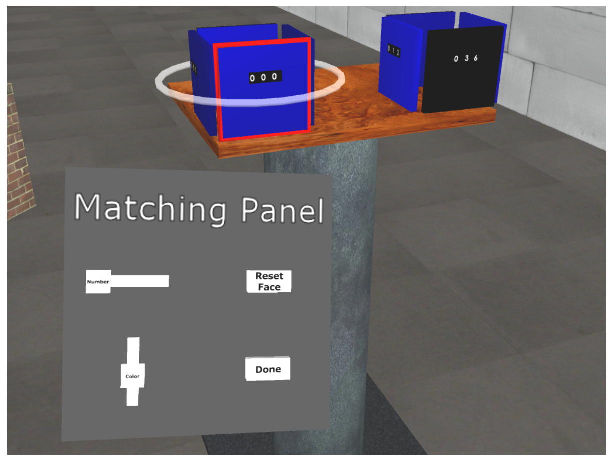

Figure 5 - The interface panel and environment for the matching task

Participants were told that the numbers on the two cubes had to be exactly the same, but that the colors only had to be “close.” A fairly large threshold was used to determine if colors were matched. We did not want there to be potential frustration from performing a perceptual task (such as color matching). However, we did want to have both horizontal and vertical sliders to use both dimensions of the NDH palm surface plane.

When rotating the left cube, the right cube rotated in synchronism, so the participant always saw the two corresponding faces that were to be matched. The user had to rotate the left cube several times to get to all four sides. This task required users to discard and reacquire the panel several times and, consequently, allowed us to query the participants on the difficulty of this action.

Measures

After participants had completed both the training and matching tasks, they were orally interviewed with several open-ended questions. These questions were designed to elicit comments on fatigue, ease of learning to use the panel, ease of manipulating widgets, and ease of discarding and reacquiring the panel as needed for their task.

Following the oral interview, participants were given a written exit questionnaire with several questions arranged on Likert scales. These questions asked participants to rate ease of use on a scale from 1 (extremely easy) to 7 (extremely difficult) and complemented the open-ended comments received during the interview.

4 Results

Overall, participants reported our interface intuitive to use after becoming familiar with it over the 10-15 minutes they used it for the training session and the matching task.

Using the widgets

Figure 6 shows how participants rated the ease of interacting with our system. All but two participants rated widget activation as moderately easy, easy, or extremely easy. The remaining two rated it neutral. Universally, participants found buttons to be easier to use than sliders. Half of the participants rated sliders to be moderately easy, easy, or extremely easy. Two participants rated sliders neutral, and two rated them moderately difficult. Those who rated sliders neutral had not been told about the slider locking mechanism due to experimenter error, and we observed that the two who rated sliders as moderately difficult did not hold the NDH palm flat when open, leading to the issues discussed in Section 2.6.

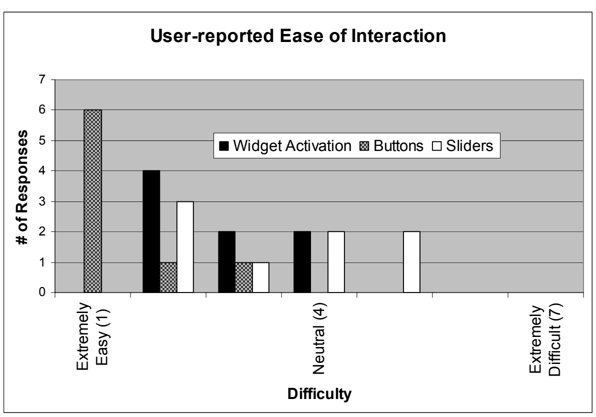

Figure 6 – User-reported ratings of ease of interaction on a scale from 1 (extremely easy) to 7 (extremely difficult)

The numeric ratings are supported by user comments in the interviews.

> *At first I found myself not actually touching my hand, but once I realized it's a lot easier to go all the way down and touch [my NDH], it got a lot better. It made it much quicker…instead of [slowly] moving my hand [(DH finger)] until I got close enough, I could just put my hand [(DH finger)] down and control [the widgets].*
>
> *[Touching the NDH] was fine; ...it made it easier to actually comprehend the contact with the widgets.*
>
> *Selecting the various widgets, that was really easy….*
>
> *[For] buttons ... it was fine, it was good to have a stop, but with the sliders, your hand isn't perfectly smooth, so you can hit something [that restricts your movement].*
>
> *…eventually my [DH] finger ran out of room to go, especially when [trying to] go all the way to 99. …you're all the way over here on the palm, right near the [NDH] fingers, so that gets pretty difficult.*
>
> *During the color section, I would touch the bottom of the [slider], and my finger was too low [and off the bottom of the palm]. I had to go a little bit higher to [use the slider].*

Interacting with the panel and cubes

After some practice, none of the participants had any difficulty discarding or reacquiring the interface panel. All participants felt that moving and rotating the interface panel when it was attached to the NDH was easy. Several participants found rotating the cube to be difficult until they adjusted their chair to a more comfortable position.

> *No trouble getting rid of the panel… [but] the panel was just a little too big to leave in the scene and yet still have a view of the two cubes.*
>
> *Once I got used to the interface, it was pretty simple.*
>
> *I actually liked [moving and rotating the panel]; that was extremely handy. The cube was difficult to rotate. …I found that if I grabbed [the cube] from above as opposed to from under, I could control it better.*
>
> *[Moving and rotating the panel] was just like moving my hand. When I was trying to [rotate the cube] all the way around, I had a tendency to let my palm slide off of the ring.*

Fatigue

A participant whose chair was improperly positioned in the real environment reported fatigue from having to stretch out his arm to rotate the cube. Otherwise, none of the participants reported having felt fatigued as compared to everyday arm and hand usage while in a seated position.

Suggested Improvements

During their interviews, the study participants made suggestions for improvements that we will consider in the next iteration of the system.

> *I liked the locking behavior on the sliders. It was just a little bit odd to see [a slider] snap back, but it was better to have it snap back than for my finger to move the slider off the originally intended position. Some visual feedback that you've locked the slider might be nice though.*

> *The sliders were a little too sensitive, I thought…. If it had worked more consistently with my hand I think it would have been a little easier – maybe I didn't hold my hand flat enough.*

> *I think the size of the widget panel was somewhat restrictive…. If the widget panel was a cube on your hand that you could rotate…then you wouldn't have half your view blocked by this big panel. I wanted to be able to look at the cubes clearly while manipulating the panel, and I felt a little bit like it was one or the other [because of the low field of view].*

5 Conclusions and Future Work

We have presented a user interface for VEs which makes use of the NDH for haptic feedback. Our first round of usability tests indicates that the technique is viable and warrants further development.

There are several directions for future research. A more performance- and precision-oriented study directly comparing our technique with systems using physical paddles would provide valuable information on the class of tasks suitable for the technique.

A contact switch or pressure sensor on the DH fingertip would give us better information about when the fingertip and NDH palm are in contact, allowing us to avoid using the slider locking mechanism. Additionally, more precise hand-size calibration techniques may significantly increase the system's effectiveness. For example, given more data on a user's hand size, we could scale the size of the widgets such that the user's DH finger will not move off the surface of the NDH palm (e.g. when setting extreme slider values).

Current head-mounted display technology for immersive VEs restricts the field of view considerably. When the interface panel is large, the restricted field of view often causes the panel to obstruct other objects in the environment. This makes it necessary for the user to look back and forth between the panel and the objects they are manipulating with widgets (the same is true for physical paddle interfaces). Research addressing similar problems for desktop user interfaces may be applicable to immersive VEs as well [16]. We have several ideas:

- When a widget is activated, do not display any of the rest of the interface panel – only display the hand and currently snapped widget. Upon widget deactivation, the entire panel will be visible again.
- Exploit transparency effects to give a view of both the panel and the environment.
- Map the interface to a cylinder (or truncated cone) instead, and place it around the user's non-dominant forearm.

This last idea about the cylinder has several potential advantages. Since the forearm provides a larger surface than the palm, all of the widgets can be active simultaneously. The length of the forearm will also allow for longer dragging operations (e.g. sliders with much longer tracks). Additionally, the forearm (essentially) is a single rigid body, so it can be tracked with a single rigidly mounted tracker, calibrated to determine the forearm's primary axis of rotation.

Acknowledgements

Support for this research was provided by the Office of Naval Research and the NIH National Institute for Biomedical Imaging and Bioengineering. The hand models used in our system are © 2005 Immersion Corporation.

We would like to thank Chris Oates, Chris VanderKnyff, Jeff Feasel, Jason Jerald, Gennette Gill, Fred Brooks, and Rob Lindeman for their helpful ideas and advice. We would also particularly like to thank Elise London for her help in recruiting participants, and Eric Burns for invaluable discussion and for his work on the video.

References

[1] R. Balakrishnan and G. Kurtenbach. Exploring Bimanual Camera Control and Object Manipulation in 3D Graphics Interfaces. In *Proceedings of ACM Conference on Human*

Factors in Computer Systems (CHI), pp. 56-63, 1999.

[2] W. Buxton, R. Hill and P. Rowley. Issues and Techniques in Touch-Sensitive Tablet Input. *Computer Graphics*, 19(3): 215-224, 1985.

[3] W. Buxton and B. Myers. A Study in Two-Handed Input. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI)*, pp. 321-326, 1986.

[4] S. Coquillart and G. Wesche. The Virtual Palette and the Virtual Remote Control Panel: A Device and an Interaction Paradigm for the Responsive Workbench. In *Proceedings of IEEE Virtual Reality (VR)*, pp. 213-217, 1999.

[5] L. Cutler, B. Fröhlich and P. Hanrahan. Two-Handed Direct Manipulation on the Responsive Workbench. In *Proceedings of ACM Symposium on Interactive 3D Graphics (I3D)*, pp. 107-114, 1997.

[6] Y. Guiard. Symmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model. *The Journal of Motor Behavior*, 19(4): 486-517, 1987.

[7] K. Hinckley, R. Pausch, D. Proffitt and N. Kassell. Attention and Visual Feedback: The Bimanual Frame-of-Reference. In *Proceedings of ACM Symposium on Interactive 3D Graphics (I3D)*, pp. 121-126, 1997.

[8] K. Hinckley, R. Pausch, D. Proffitt and N. Kassell. Two-Handed Virtual Manipulation. *ACM Transactions on Computer-Human Interaction*, 5(3): 260-302, 1998.

[9] K. Hinckley, R. Pausch, D. Proffitt, J. Patten and N. Kassell. Cooperative Bimanual Interaction. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI)*, pp. 27-34, 1997.

[10] B. Insko. Passive Haptics Significantly Enhances Virtual Environments, Ph.D. dissertation, University of North Carolina at Chapel Hill, 2001.

[11] R.W. Lindeman, J.L. Sibert and J.K. Hahn. Towards Usable VR: An Empirical Study of User Interfaces for Immersive Virtual Environments. In *Proceedings of ACM Conference on Human Factors in Computing Systems (CHI)*, pp. 64-71, 1999.

[12] R.W. Lindeman, J.L. Sibert and J.N. Templeman. The Effect of 3D Widget Representation and Simulated Surface Constraints on Interaction in Virtual Environments. In *Proceedings of IEEE Virtual Reality (VR)*, pp. 141-148, 2001.

[13] M.R. Mine, J. Frederick P. Brooks and C.H. Sequin. Moving objects in space: Exploiting proprioception in virtual environment interaction. In *Proceedings of ACM SIGGRAPH*, pp., 1997.

[14] I. Poupyrev, N. Tomokazu and S. Weghorst. Virtual Notepad: Handwriting in Immersive VR. In *Proceedings of IEEE Virtual Reality Annual International Symposium (VRAIS)*, pp. 126-132, 1998.

[15] D. Schmalstieg, M. Encarnação and Z. Szalavári. Using Transparent Props For Interaction With The Virtual Table. In *Proceedings of ACM Symposium on Interactive 3D Graphics (I3D)*, pp. 147-154, 1999.

[16] M.A. Tapia and G. Kurtenbach. Some Design Refinements and Principles on the Appearance and Behavior of Marking Menus. In *Proceedings of ACM Symposium on User Interface Software and Technology (UIST)*, pp. 189-195, 1995.