Journal of Visual Languages and Computing (1999) **10**, 3-17

In Search of the ‘Magic Carpet’:

Design and Experimentation of a Bimanual 3D Navigation Interface

Shumin Zhai Eser Kandogan* Barton A. Smith Ted Selker

IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120 USA

+1 408 927 1112, {zhai, eser, basmith, selker}@almaden.ibm.com

ABSTRACT

Hardware and software advances are making real time 3D graphics part of all mainstream computers. World-Wide Web sites encoded in Virtual Reality Modeling Language or other formats allow users across the Internet to share virtual 3D “worlds”. As the supporting software and hardware become increasingly powerful, the usability of the current 3D navigation interfaces becomes the limiting factor to the wide-spread application of 3D technologies. In this paper, we analyze the human factors issues in designing a usable navigation interface, such as interface metaphor, integration and separation of multiple degrees of freedom, mode switching, isotonic versus isometric control, seamless merger of the 3D navigation devices with the GUI pointing and scrolling devices and two-handed input. We propose a dual joystick navigation interface design based on a real world metaphor (bulldozer), and present an experimental evaluation. Results showed that the proposed bulldozer interface outperformed the status quo mouse-mapping interface in maze travelling and free flying tasks by 25% to 50%. Limitations of and possible future improvements to the bulldozer interface are also presented.

INTRODUCTION

We live in a three-dimensional (3D) world. Technological advances in computer graphics hardware, software and display systems will soon make real time 3D capabilities available to all mainstream computer systems. Furthermore, 3D

---

* Currently Eser Kandogan is a member of the technical staff at Silicon Graphics Inc. This work is done during his internship at IBM Almaden Research Center.

technology is also beginning to be integrated into the Internet technology. Languages such as Virtual Reality Modeling Language (VRML) have made it possible for designers and users on various platforms to escape from the flatland of 2D HTML pages to share their virtual 3D worlds on the desktops of users across the globe.

However, disorientation and difficulty of knowing how to navigate in current 3D interfaces are intolerably inhibiting the widespread applications of 3D technologies. Designing effective 3D navigation interfaces is both a challenge and an opportunity to the user interface design communities. First, it is a rare opportunity. As if we could move back in time to have the chance to reinvent interfaces for aircraft and motor vehicles, we are now at the beginning of a new era when nothing is not (yet) too late. Furthermore, the “virtual navigation craft” does not impose on us any mechanic or dynamic constraints as in aircraft and motor vehicles. The only fundamental constraint lies in accommodating human capabilities and limitations. In other words, designing 3D navigation interfaces is less like designing motor vehicles but more like searching for the “magic carpet” of the virtual world. A magic carpet takes us where we want to go without our having to worry about the mechanical details of getting there.

Much research has been done on input devices for 3D interfaces [see 5, 14 for recent reviews]. However, most of these studies were conducted with manipulation, not navigation (locomotion) tasks. There could be fundamental differences between manipulation and navigation tasks. For manipulation tasks, it is necessary to have all 6 degrees of freedom simultaneously available for the user to form coordinated manipulation [16]. In navigation tasks, it may be more efficient to control a few degrees of freedom at a time.

There are many practical requirements for a 3D navigation interface to rapidly gain wide acceptance. Particularly, we have to consider the following issues:

1. Integration with the existing GUI interfaces and tasks. Like any new interface technology, 3D virtual world navigation is in a bootstrap situation: not enough application and content may make it unjustified to have a special interaction device and lack of a usable interface and user population make developing 3D content less attractive. The solution lies in seamless integration of the 3D navigation devices with the existing GUI interfaces.
2. Low cost, for the same reason as above.
3. Both novice and expert “friendly”. An acceptable navigation interface should be easy to learn in a few minutes with a navigation metaphor familiar to most users while offering high performance to expert users.

THE STATUS-QUO 3D NAVIGATION TECHNIQUE

Although there are many VRML browsers available today, the user interfaces of these browsers are very similar to each other. The basic interface design was set in SGI's WebSpace^{TM} (later known as CosmoPlayer^{TM}), the first commercially available VRML browser [9], which in turn was based on the previous 3D user interface work [13]. With these existing browsers, 3D navigation is done by mapping 2 degrees of freedom (DOF) mouse cursor movements onto various translation and rotation degrees of freedom, according to the selected mode. For example, the WorldView^{TM} browser [12] has the following modes: walk (z-translation[2] and y-rotation in rate control), pan (x- and y-translation in rate control), turn (x- and y-rotation in rate control), roll (z-rotation in rate control), study (x- and y-rotation of the world, instead of self, in position control) and goto (Figure 1). A user can switch between these different mapping schemes by clicking on appropriate menu buttons. Some browsers also provide a restore mode to readjust the y coordinate of the view coordinates with that of the world coordinates. Much effort had been spent on fine-tuning the mapping transfer functions in various modes.

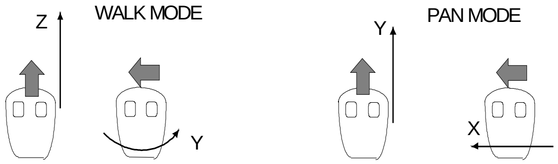

Figure 1: Mouse mapping navigation: Walk and Pan Mode

When navigating in VRML worlds with the current browsers, users typically find themselves off their targets, facing upside-down, or loose their locations. Multiple factors, such as low frame rate that delays the control feedback and sometimes badly designed "cueless" worlds themselves contribute to the "lost in virtual world" problem. The most critical factor, however, probably lies in the mouse mapping navigation technique itself. Among the many drawbacks of this technique, the following are the most noticeable:

---

[2] This paper uses the following 3D notation convention: axis is from the viewer's left side to the right side; y is the bottom up direction; z is the direction the viewer facing. Pitch is rotation about x (also labeled as x-rotation); yaw is rotation about y axis (y-rotation) and roll is rotation about z axis (z-rotation).

1. Mode switching. It is known that mode switching in user interfaces in general should be avoided. Mode switching causes inconsistent response to the same input. For the same mouse movement, the results are different depending on the current mode. It is well known that when consistent mapping exists, human information processing behavior tends to become an “automatic process” which requires little central capacity, attention or effort. In contrast, when consistent mapping is absent, human behavior tends to be a “controlled processes” which requires effort, attentive resource and central capacity [10, 11].

2. In many of these modes, cursor motions are mapped to movement “speed”. The farther one moves the cursor from the initial click position, the faster the movement is. In other words, cursor displacement is used for *rate control*, which is well suited for navigation where smooth and controllable speed is desirable. Experiments have shown that effective rate control requires a self-centering mechanism in devices such as isometric[3] or elastic joysticks. Isotonic devices such as the mouse are poor in rate control tasks [14]. Note that when one uses a rate controlled joystick such as the TrackPoint™ in IBM’s notebook computers to do virtual world navigation with today’s VRML interfaces, the self-centering effect in the joystick is not utilized, since the self-centering variable (force) is not directly mapped onto the speed.

Clearly we need to search for different devices, techniques and metaphors to replace the status-quo navigation interface. One apparent choice is the 6 DOF hand controllers such as the Spaceball™ (see [14] for a review of 6 DOF devices). Several reasons lessen the feasibility of this option. First, these devices have been relatively expensive due to the small market size and they are not integrated with the general GUI interface (e.g. pointing), thus preventing a critical mass of user population to overcome the bootstrap situation. More importantly, these devices are designed primarily as “manipulation”, not as “navigation” devices. Although both manipulation and navigation require multiple degrees of freedom controllability and they are mathematically equivalent problems, they may differ significantly in the human factors. When we manipulate objects, the multiple degrees of freedom tend to be integrated. People do not distinguish various degrees of freedom when tying shoe laces or swinging a golf club. On the other hand when we move (navigate) in real world, we rarely use all the 6 degrees of freedom simultaneously. We primarily stay on a 2D surface, move in a given direction (x- and z-translation) or turn around (y-rotation). We may move up and down (y translation) when proper means are available such as stairs and elevators. Except in gymnastics, large amount of pitch (x-rotation) and roll (z-rotation) rarely happen to our body.

---

[3] Isometric devices have constant location and infinite resistance to displacement. Isotonic devices have constant resistance (e.g. zero) to displacement. Elastic devices have resistance proportional to displacement.

Based on these observations and the criteria outlined in the introduction, we developed the bulldozer interface for VRML (or any other 3D interface) navigation.

THE BULLDOZER INTERFACE

The bulldozer interface is based on a dual-joystick configuration used by two hands. Although such a concept can be implemented with any pair of elastic or isometric joysticks, we chose TrackPoint^{TM,} the pointing stick used in IBM notebook computers for the current implementation. The circuit of both the joysticks connects to a single PS2 mouse port and communicates with various operating systems through the TrackPoint IV driver. There are currently two ergonomic designs to dual joystick input. One is to integrate the joysticks into the keyboard control surface, one below the C, V and one bellow M, N keys. The second design is a wrist-pad instrumented with two sticks (Figure 2). In both of the designs the two TrackPoints are used as general GUI input devices: one can be used as a pointing device and the other for scrolling. Research has shown such two-handed scrolling and pointing is more efficient than traditional graphical scrollbar solution [17].

Figure 2: The bulldozer keyboard wrist-pad

While navigating in 3D virtual worlds, the two joysticks are operated in an extended flyable bulldozer metaphor. A total of four degrees of freedom were made available to the user (x-, y-, z- translation and y-rotation). Pushing both joysticks forward, the user moves forward in the virtual world. Pulling both joysticks moves the user backwards. Pulling the left and pushing the right joystick turns the user to the left and vice versa. Pushing both of them horizontally in the same direction slides the user sideways. Pushing them in opposite directions moves the user up (outwards) and down (inwards) (Figure 3).

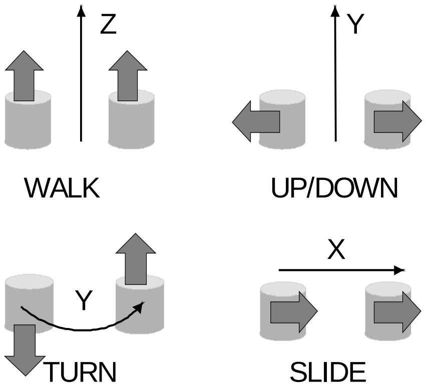

Figure 3: Extended flyable bulldozer metaphor

We gained deeper understanding of navigation in the iterative process of designing and testing various mapping schemes of the bulldozer interface. For example, we assumed that it would be advantageous to allow simultaneous turning and moving forward as in a real bulldozer. However, pilot testing showed this design to be inferior to making these two types of movement mutually exclusive. This is probably due to the need for separation of degrees of freedom in navigation tasks. When both moving forward and rotation are simultaneously available, it was difficult for the user to move straight ahead without wandering in different directions. Imagine how difficult it would be to move a shopping cart that has four, instead of two (front) rotating wheels. In fact there are poorly designed travel cases that do have four rotating rollers that offer too many simultaneous degrees of freedom. Note that in real bulldozers (or shopping carts), yaw (y-rotation) and z-translation are partially simultaneously available, but the constraints (the non-rotating back wheels in shopping carts and the linear tracks in bulldozers) make unintended rotations less likely to happen. Systematically exploring similar physical constraints in the virtual bulldozer interface should be explored in future work.

Further along the issue of separating degrees of freedom, we intentionally introduced a dead space in the transfer functions to reduce possible coupling between degrees of freedom. For example, when a user is intending to move forward and pushes both sticks forward, the two force vectors may not be perfectly parallel, and the slight angle between the two force vectors would produce an unintended up or down movement. The transfer functions introduced a buffer region in which a translation action is filtered out when it constitutes less than 30% of the total translation. This is also useful to remove any noise in the data. Also implemented in the transfer functions are the different sensitivity constants for each degree of freedom. Our observation has been that moving forward should be the most sensitive and yaw motion the least sensitive compared to other degrees of freedom.

The bulldozer software consists of a VRML browser, the TrackPoint IV driver, a Java applet and a Netscape plugin. The VRML browser, the Netscape TrackPoint plugin, the driver and the applet concurrently run on the system (Figure 4). TrackPoint data is buffered in the driver, until the Netscape plugin issues a poll to the driver, in which case the message array in the plugin is filled with the latest TrackPoint data. The Java applet polls the plugin in a loop (GET-PROCESS), and gets the collected TrackPoint data. Then, these data are processed in the applet to calculate the new viewpoint position and orientation. Finally, the applet sends out new events to the VRML Browser through its External Authoring Interface (EAI) to update the scene appropriately. This GET-PROCESS loop in the applet is executed repeatedly.

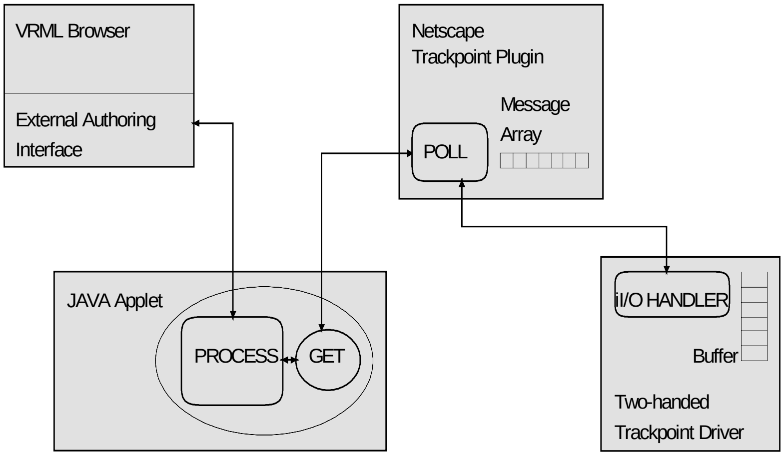

Figure 4: Execution Model

The bulldozer software runs on an IBM PC with a 100MHz Pentium processor and a Rendition Verite based graphics accelerator card. 10 frames per second (fps) update rate was achieved with the Worldview^{TM} VRML browser.

Comparing against the criteria we outlined in the introduction, the bulldozer interface offers the following promise:

1. Integration with existing GUI devices, interfaces and applications. The dual joystick offers multiple functions. For example, while one of the joysticks offers the usual pointing function, the other on the non-dominant hand can be used for scrolling, eliminating the need of using the scroll bar. Studies have shown such a dual stream pointing and scrolling interface offers significant advantages in document (e.g. web) browsing tasks [17]. Incorporated into the keyboard and the wrist-pad, the dual joystick design can also be used for two-handed menu operations [1, 6] or two-handed graphical

manipulation [7]. They are ergonomically designed to avoid accidental operation when users rest their hands on them (Figure 1). We have also developed a generic software solution that requires no change to the VRML browser.

2. Low cost. Due to the large notebook user base, miniature joysticks are mature technology mass-produced at relatively low cost.

3. Both novice and expert user "friendly". The bulldozer interface offers higher navigation performance than the status quo for both novice and expert users.

The last point was a hypothesis based on the following analyses:

1. The bulldozer interface is modeless.
2. It works in rate control scheme that is well suited for navigation.
3. It offers more degrees of freedom, and the integration and separation of these degrees of freedom were designed according to the nature of human navigation.
4. It is based on an intuitive physical metaphor, not an artificial mapping of control motions. For all of the degrees of freedom, there is an isomorphic mapping between control and effect, except moving up and down.

However, the usability of an interface technology can not be reliably predicted based on analysis alone, due to the complexity of human performance. We therefore conducted the following experiment to test our hypothesis to test our prediction.

EXPERIMENTS

Purpose

The bulldozer interface was experimentally tested against the status-quo mouse-mapping navigation interface. As analyzed earlier, the bulldozer navigation interface differs from the status-quo interface in many ways, including the number of degrees of freedom, mode switching and metaphor. The purpose of this experiment was not to quantify how each of these factors influences the user’s navigation performance, but rather to estimate user performance and preference to our overall design of the bulldozer interface. Due to its fast rendering performance on PC machines, WorldView^{TM} browser was used as the representative of the status-quo, mouse mapping interfaces (hereafter referred as “mouse interface”) in the experiment.

Tasks

The experimental tasks were chosen to represent typical maneuvers in virtual 3D world navigation. Two VRML worlds were designed for the experiment. The first VRML world was a maze with a single tunnel path that consisted of three sections: A. Wide tunnel with occasional turns; B. Narrow tunnel with frequent turns; C. Wide roads with occasional barricades the subject had to jump over or move underneath. Figure 5 shows the overview and the subject's view of the maze. In the experiment subjects were required to move along the tunnel as quickly as possible, starting from section A, all the way through the maze without "accidents" (hitting on wall or barricade). When accidents did happen, they were required to backup and then continue. The backup time was a penalty added to the total completion time. Timing for each section was recorded separately. The sections from A to C were designed to increase in difficulty of navigation.

During pilot experiments, we observed that the mouse interface performed well when the turns were not too sharp. We therefore designed both wide and narrow sections in the maze task. We expected the bulldozer interface to outperform the mouse interface in section B (narrow) and section C (barricade). In section C, in addition to moving forward and turning, the subjects also had to move up and down to avoid the barricades.

Figure 5: Overview (top) and the subject's view (bottom) of the maze travelling task

The second VRML world used in the experiment consisted of three 3D targets with varying heights and orientations. Each target object had a red rectangle on its surface. The subjects were asked to adjust their orientation in front of each red rectangle and then pass through it on all three targets in a given order. Although the subjects were not required to follow a particular path and were allowed to fly freely between the objects, they had to adjust their 3D position and orientation to successfully pass through each rectangle. Figure 6 shows one of the three targets in the second VRML world.

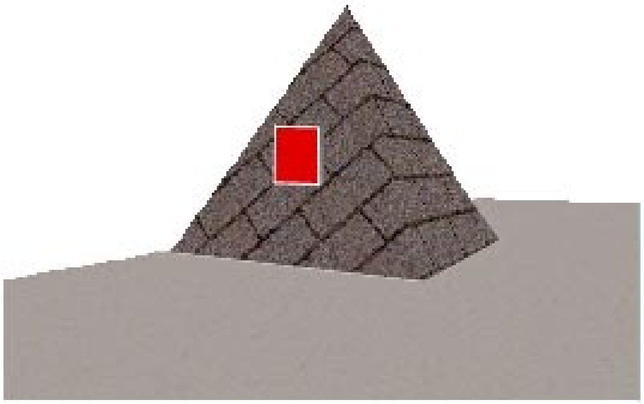

Figure 6: A target in the free flying task

Subjects

A total of 12 subjects participated in the experiment, 3 females and 9 males. Most of them were between 20 to 30 years of age. All twelve subjects were regular users of the mouse. Eight of them used isometric joystick regularly, three subjects had little experience and one had no previous experience with isometric joystick. Daily computer usage among subjects varied from less than an hour (2 subjects), to above 10 hours (1 subject), with most of the subjects using between 4 to 10 hours. Eight of the subjects had played action computer games more than 20 times in the past, four of them had not.

A within-subjects, order balanced design was used in this experiment. Each subject was tested on both the bulldozer and the mouse interface. Six subjects performed with the bulldozer interface first, followed by the mouse; the rest of the subjects had the reverse order.

With each interface, the subjects were first given instructions on the usage of the input device for navigation, and then they were allowed to practice, which lasted as long as the subjects needed to explore the different degrees of freedom available (typically less than 10 minutes). Following the practice session, the subjects proceeded to the navigation tasks in the two VRML worlds. Navigation with each interface in each VRML world was repeated three times to examine learning effects. A subjective evaluation questionnaire was given to the subjects immediately upon completion of all tasks with each interface. The total experiment for each subject lasted less than an hour. Participation in the experiments was voluntary. A motivating cash award was set for the top performer whose total navigation time (with both interfaces) was the shortest.

Results

Figure 7 shows the mean completion times with 95% confidence bars shown for the three sections of the maze. A repeated measure variance analysis showed that subjects' performance with the bulldozer interface was statistically faster than the mouse in all three sections (A, B, C). While performance in the wide section of the maze was already significant ($F_{1,11}$=8.89, $p < 0.05$), greater differences were measured in the narrow and barricade sections ($F_{1,11}$=14,96, $p < 0.005$ and $F_{1,11}$=66,05, $p < 0.0001$ respectively). Taking the mouse condition as the reference, on the average subjects performed 25% faster with the bulldozer interface in the wide section of the maze. In the narrow section, the difference was 32% and in the barricade section it rose up to 50%. A statistically significant learning effect was also observed in the wide and narrow sections of the maze based on trial comparisons. No significant interface ordering effects were observed.

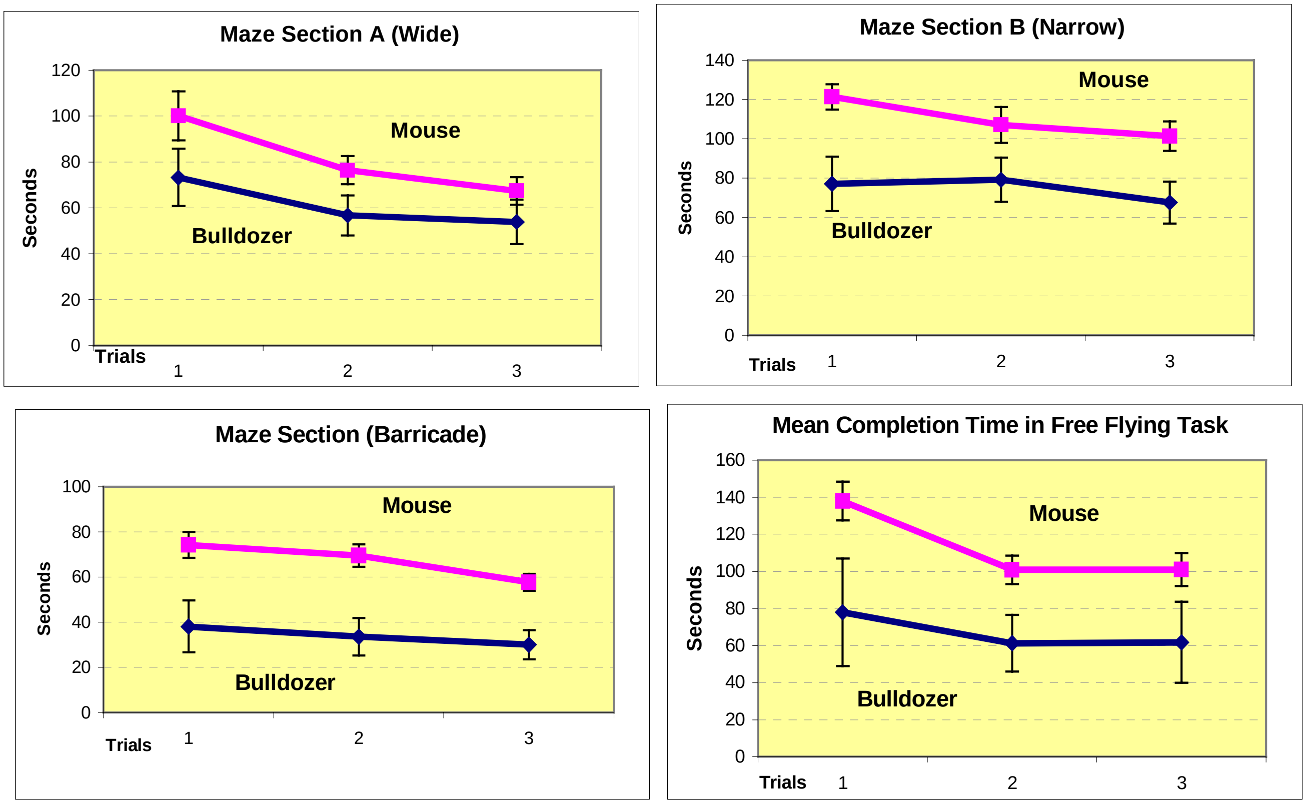

Figure 7: Mean completion times in the maze travelling and free flying tasks

The mean completion times for the free flying task were also significantly shorter for the bulldozer interface compared to the mouse interface at the level of 0.01 ($F_{1,11}=10.72$, $p < 0.005$). Subjects on the average performed 41% faster with the bulldozer (Figure 7). While no ordering effect was observed, a significant learning was observed. The shortest completion times were all achieved with the bulldozer interface in the third trials but by different subjects ($t_A=29s$, $t_B=37s$, $t_C=11s$, $t_F=27s$).

Subjective Evaluation

In the first set of questions subjects were asked to rate general characteristics of both of the methods. In all of the rating dimensions the bulldozer interface received a higher score than the mouse interface (Figure 8). In particular, the mouse was found to be frustrating, while the bulldozer was more stimulating.

In the second set of questions, subjects evaluated specific characteristics of the methods such as learning, usage, predictability and speed (Figure 8). While subjects rated the ease of learning and usage of operations equally well, the bulldozer interface was rated to be more predictable. The speed in both conditions was found to be moderate, which was about 10 frames per second on the average.

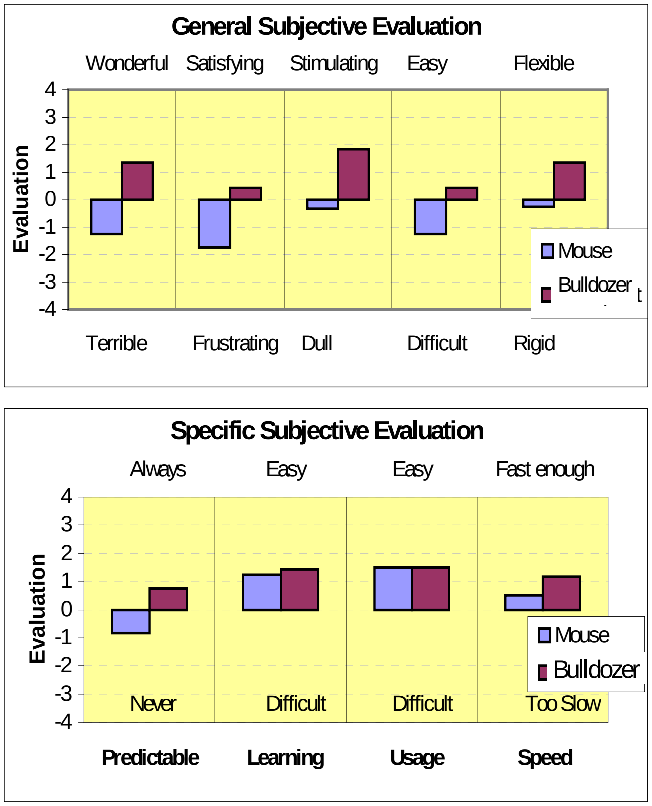

Figure 8: Subjective evaluation results

CONLUSIONS AND DISCUSSIONS

The experiment results showed that the bulldozer interface provided a clear and substantial amount of advantage over the status quo mouse-mapping interface, both in terms of navigation performance and in terms of user's preference. Although all of the subjects had more experience with the mouse, their mean navigation performance was faster with the bulldozer interface in all conditions. Even the subject with no previous experience with the isometric joystick performed 26% faster in the wide section of the maze, 46% and 48% in the narrow and barricade sections, respectively. Furthermore, the advantage of the bulldozer interface became greater in more difficult tasks, such as narrower tunnels with more frequent turns.

We therefore conclude that the bulldozer interface offers a promising alternative to navigation in VRML or other 3D applications. All of the following factors might have contributed to the relatively superior performance of the bulldozer interface.

Mode Switching

In the status-quo mouse-mapping interface, users had to explicitly switch between different modes in order to achieve higher degrees of freedom than the mouse has (2 DOF). Not only the switching action takes time and motion, it may also impose cognitive load of coping with different control responses to the same mouse input: the same mouse cursor movement resulted in different responses (rotation, translation etc) depending on the mode. Subjects repeatedly commented that using mode switching was frustrating. Indeed, the advantage of the bulldozer interface over the mouse interface was greater in the barricade section of the maze (50%) and in the free flying task (41%) than in the maze sections (25% and 32% in wide and narrow sections respectively). Both the barricade section of the maze and the free flying task required mode switching between walk (z-translation and y-rotation) and pan (x-, y-translation) modes.

Isometric vs. Isotonic Rate Control

The time saved by eliminating mode switching alone does not count for all the performance gains in the bulldozer interface. In the section A (wide) and B (narrow) of the maze task which needed only one mode (walk) of the mouse interface, the bulldozer interface still yielded significantly faster performance. Among many others, isometric rate control in the bulldozer interface could be a contributing factor.

The problem of isotonic device (mouse) for rate control [14] was very apparent in the status-quo interface, especially for navigating in a narrow tunnel with many turns, where the user had to release the mouse, bring the mouse back into the window and readjust their initial position. When using isometric joysticks in rate control, the moving speed was directly proportional to the applied force and the self-centering mechanism in the isometric devices made it easy to slow down or stop at places the user intended to.

The Integration and Separation of Degrees of Freedom

The bulldozer interface made four degrees of freedom available to the user without requiring explicit mode switching. However, as discussed earlier, some degrees of freedom should be integrated (simultaneously activated) and others should separated (de-coupled) by means of mutual prohibition or thresholding. For instance, pilot tests showed that it was preferable to have forward movements (z-translation) and turns (y-rotation) separated so that it is easier to move straight ahead. On the other hand, we allowed simultaneous side translation and rotation. We observed that subjects quickly made use of such a capability and performed circular movement facing the center of the circle by simultaneously moving sideways (x-translation) and turning (y-rotation). This style of navigation proved to be very useful in the free flying task in which the users could move themselves while keeping the target object in view.

Control Response Compatibility

It is known that control and response compatibility is a critical factor to human control performance [3]. With the bulldozer interface, the control directions of three of the four degrees of freedom, the x- and z-translation and the y-rotation are isomorphic to those of the response. Such a compatible design made the use of the bulldozer interface intuitive. We observed that even those subjects who were not familiar with the bulldozer metaphor rapidly learned the usage of the two-handed Bulldozer interface.

Unfortunately, in one of the four degrees of freedom, the control and response direction in y-translation (moving up and down) had to be different: moving up or down was done by laterally pushing the two joysticks apart or towards each other. The subjects occasionally confused the up and down control or had to consciously think about the correct control action. As subjects gained more practice, this problem tended to subside.

The bulldozer interface outperformed the mouse-mapping interface by a large margin, even when vertical movement was involved. Despite the fact that the y-translation is used much less frequently than the other three degrees of freedom, it is desirable to improve the current interface to have control-response compatibility in all degrees of freedom. This can be done if the two joysticks are made long enough to be grasped and pulled upwards and pushed downwards for up and down translation. However such a design makes the physical interface overspecialized for 3D navigation and compromise our goal of having a generic interface (such as the wrist-pad design) that is also well-suited for other mainstream GUI applications. The bulldozer interface can also be implemented using standard game-type (elastic, self-centering) joysticks. Again, this would compromise our goal of using the same hardware for standard pointing and scrolling tasks as well.

Two-handed input

The bulldozer interface uses two-handed input. The fact that we use both hands in dealing with the natural world have motivated researchers to study such a capability for using computers [1, 2, 6, 7]. Researchers first empirically demonstrated that users were capable of simultaneously providing continuous input to computers from both hands without significant cognitive overhead. The theory that the two hands have *asymmetrical* functional division of labor [4] provided insight into designing two-handed input systems [1]. Recently, it was demonstrated that in addition to time-motion efficiency, two-handed input also brings cognitive advantage [7]. Our experiments showed that *symmetrical* actions from both hands also work well in human-computer interaction, as in many daily tasks.

FUTURE WORK

We are only at the beginning of the search for the magic carpet for virtual 3D world navigation. Much future work has been pointed out in the foregoing discussion. Additional future work includes:

Comparison with other 6 DOF devices. For this study we have limited ourselves to comparing the Bulldozer interface with mouse mapping techniques, the most commonly used VRML navigation method. Although specialized 6 DOF input devices such as a Spaceball is not well integrated with general 2D GUI functions such as pointing and scrolling, quantitative experimental comparison between the Bulldozer interface and these integrated 6 DOF devices should be conducted in the future.

Sounds. Perhaps due to real world experience, people tend to expect sounds while navigating. Simple sound that is proportional to speed or more complex auditory signals can improve to navigation performance and awareness. It may help to compensate for slow visual feedback in the graphical system.

Transformation Refinement. As pointed out earlier, we believe that there is much room for improvement in the transfer functions of the bulldozer interface. Better transformation schemes that handle the issue of integration and separation between various degrees of freedom may exist. For example, we may achieve better performance by simulating a real physical bulldozer (or shopping cart) mechanics and dynamics (e.g. chain length) in the transfer function to allow steady forward movement as well as quick turning.

Different metaphors. It is also possible that fast performance can be achieved with mappings that are different from the bulldozer metaphor. For example, one of the two sticks may be used for x- and z-translation and the other may be used for y-rotation and y-translation.

Side view display. We also observed that subjects initially had difficulty in deciding where to start turning with either interface due to the lack of side views. It is like driving a car with only a 90 degree front view. At times the subjects turned too early and found themselves facing the wall. However, most of the subjects quickly learned to estimate the correct point for turning. Providing side views may overcome such type of problems.

Manipulation. Although VRML is primarily used for *navigation* through 3D “worlds”, occasionally one may need to *manipulate* an object in the 3D sense. In such a case, an additional mode may have to be made available in which different mapping functions of the two sticks are required. Because manipulation and navigation are cognitively different and separate tasks, mode switching between such cases may be necessary and beneficial, provided different display, such as semitransparent hand cursors [15] are used in the manipulation mode.

In short, in searching of the magic carpet for virtual worlds, there is still a long journey ahead and there is much the user interface community can contribute towards such a goal.

ACKNOWLEDGMENTS

We thank our colleagues at the IBM Almaden Research Center, especially Steve Idhe and Kim May for their input to this research.

REFERENCES

1. Bier, E. A., Stone, M. C., Fishkin, K., Buxton, W., Baudel, T. A taxonomy of see-through tools. In *Proc. of CHI '94.* ACM/SIGCHI, New York, 1994, pp. 358-364.

2. Buxton, W. and Myers, B. A Study in Two-Handed Input. In *Proc. of CHI '86.* ACM/SIGCHI, New York, 1996, pp. 321-326.

3. Fitts, P. M., Seeger, C. S-R compatibility: spatial characteristics of stimulus and response codes. *Journal of Experimental Psychology* 46, (1953), 199-210.

4. Guiard, Y. Asymmetric Division of Labor in Human Skilled Bimanual Action: The kinematic chain as a model. *Journal of Motor Behavior* 19, 4, (1987), 486-517.

5. Hinckley, K, Tulio, J, Proffitt , D., Kassell, N., Pausch, R., Usability Analysis of 3D Rotation Techniques, *Proc. of UIST’97*, Banff, Alberta, Canada.

6. Kabbash, P., Buxton, W., Sellen, A. Two-handed input in a compound task. In *Proc. of CHI '94.* ACM/SIGCHI, New York, 1994, pp. 417-423.

7. Leganchuk, A., Zhai, S., Buxton, W. Manual and Cognitive Benefits of Two-Handed Input: An Experimental Study, *accepted for publication in ACM Transaction on Computer-Human Interaction.*

8. Mackinlay, J. D., Card, S. K., and Robertson, G. G. Rapid controlled movement through a virtual 3D workspace, In *Proc. of ACM SIGGRAPH '90* 24, F. Baskett, Eds., (August 1990), 171-176.

9. Mohageg, M., Myers, R. Marrin, C., Kent, Jim, Mott, D., Issacs, P. A user interface for accessing 3D content on the World Wide Web. In *Proc. of CHI'96.* ACM/SIGCHI, New York, 196, pp. 466-472.

10. Schneider, W. Shiffrin, R.M. Controlled and Automatic Human Information Processing: I Detection, search and attention, *Psychological Review*, 1977, 84, 1-66.

11. Shiffrin, R.M. and Schneider, W. Controlled and Automatic Human Information Processing: II Perceptual learning, automatic attending and a general theory, *Psychological Review*, 1977, 84, 127-190.

12. Worldview™ VRML 2.0 Browser, Intervista Inc., http://www.intervista.com/

13. Zeleznik, R. C., Herndon, K. P., Robbins, D. C., Huang, N., Meyer, T., Parker, N., and Hughes J. F., An interactive 3D toolkit for constructing 3D widgets. In *Proc. ACM SIGGRAPH '93*. ACM/SIGGRAPH, New York, 1993, pp. 81-84.

14. Zhai, S., Human Performance in Six Degree-of-Freedom Input Control, *Ph.D. Thesis, University of Toronto*,http://vered.rose.toronto.edu/people/shumin_dir/publications.html. See also Zhai, S., User Performance in Relation to 3D Input Device Design, *Computer Graphics* 32(4), November 1998. pp 50-54 for a cursory review of 3D input.

15. Zhai, S., Buxton, W. and Milgram, P. The partial occlusion effect: utilizing semitransparency in 3D human computer interaction. *ACM Trans. On Computer-Human Interaction* 3, 3, (Sept 1996), 254-284.

16. Zhai, S., Milgram, P. Quantifying Coordination and Its Application to Evaluating 6 DOF Input Devices, *Proc. of CHI'9: ACM Conference on Human Factors in Computing Systems.*

17. Zhai, S., Smith, B. A. and Selker, T. Improving Browsing Performance: A study of four input devices for scrolling and pointing tasks. In *Proc. of INTERACT97: The Sixth IFIP Conf. on Human-Computer Interaction*. July 1997, pp. 286-292.

Journal of Visual Languages and Computing (1999) **10**, 3-17