

Exploring 3D Navigation: Combining Speed-coupled Flying with Orbiting

Desney S. Tan¹, George G. Robertson, Mary Czerwinski

Microsoft Research

One Microsoft Way

Redmond, WA 98052

+1 425 703-1527

desney@cs.cmu.edu, {ggr, marycz}@microsoft.com

ABSTRACT

We present a task-based taxonomy of navigation techniques for 3D virtual environments, used to categorize existing techniques, drive exploration of the design space, and inspire new techniques. We briefly discuss several new techniques, and describe in detail one new technique, Speed-coupled Flying with Orbiting. This technique couples control of movement speed to camera height and tilt, allowing users to seamlessly transition between local environment-views and global overviews. Users can also orbit specific objects for inspection. Results from two competitive user studies suggest users performed better with Speed-coupled Flying with Orbiting over alternatives, with performance also enhanced by a large display.

Keywords

Egocentric navigation, 3D virtual environments, user studies, interaction techniques

INTRODUCTION

Since most Virtual Environments encompass more space than can be viewed from a single vantage point, users have to be able to navigate efficiently within the environment in order to obtain different views of the scene. In fact, a 3D world is only as useful as the user's ability to get around and interact with the information within it.

Work done in 3D navigation has typically fallen into two efforts: research to understand cognitive principles behind navigation, and point design to create navigation techniques for specific tasks and applications. In our work, we build on previous research efforts to construct a taxonomy of navigation techniques. We use this taxonomy to inspire several new techniques and to further expand the design space. We present interesting results warranting further exploration of both this design space and the performance benefits from a large prototype display.

*LEAVE BLANK THE LAST 2.5 cm (1") OF THE LEFT
COLUMN ON THE FIRST PAGE FOR THE
COPYRIGHT NOTICE.*

RELATED WORK

There exists a vast body of work on general principles in 3D navigation. Thorndyke & Hayes-Roth [28], as well as many others [13,23,25,33], studied the differences in spatial knowledge acquired from maps and exploration. Darken and others have explored cognitive and design principles as they apply to large virtual worlds [6,7,23,29]. Furnas explored view traversability and navigability for effective navigation through large data structures [11]. In his evaluation of viewpoint motion control techniques, Bowman presented a high level taxonomy categorizing motor aspects of virtual navigation [3]. This is just one way of dividing the hierarchy of interrelated tasks. In contrast, we take a task-based approach in building our taxonomy.

Numerous techniques have been reported that try to maximize the utility of display space to view existing information. Giving the user the ability to zoom in and out and work at multiple scales is a popular method of virtually expanding the space. In the 2D realm, Masui provided a means to navigate through an information space by modifying the search area in each view [19]. Bederson in Pad++ built a graphical interface in which zooming is the principle means of navigation [2]. In these systems, users had to manually keep track of the different viewpoints from which they were working.

Igarashi integrated rate-based scrolling with automatic zooming, giving the user a global view of documents as the scroll rate increases [15]. We extend Igarashi's work and present solutions to two problems: allowing the user to transition precisely from overview to specific targets in the local view, and allowing users to inspect particular objects.

In 3D navigation systems, Fukatsu presented simultaneous overlaying of global bird's eye views as well as local views of the world [9]. Similarly, Stoakley used worlds in miniature to create such an effect [27]. In their implementation, Robinett explored scaling as a means to allow the user to quickly fly to different locations [23].

¹ Currently at Carnegie Mellon University School of Computer Science, 5000 Forbes Ave., Pittsburgh, PA 15213.

Ware described cyclopean scaling and depth modulated flying, scaling velocity based on the various depths of objects [4,30,31]. To aid the user in finding landmarks, Elvins created miniature virtual worlds called Worldlets that act as 3D thumbnails [8]. In these systems, users were either required to work simultaneously in two different unconstrained coordinate systems, or had to explicitly switch back and forth between the two.

Another means of expanding the space is to distort the views of the world. Information visualization techniques such as Fisheye Views [10], Perspective Wall [18], the Document Lens [22], and Desert Fog [16] are such techniques. These techniques are sometimes disorienting and place additional cognitive load on the user to mentally re-map the spatial relations.

We have incorporated lessons learned from these research efforts on 3D navigation in our current designs, discussed below.

TAXONOMY

Developing a taxonomy enables a more disciplined exploration of a design space. We began by categorizing existing navigation techniques and using the differences between the techniques to create a preliminary structure of navigation tasks. We then expanded and generalized this structure to include missing areas in the design space. This allowed for the identification and creation of new navigation techniques. These new techniques, in turn, allowed us to think of navigation in different ways and fueled further expansion of the space. Our current taxonomy, shown in Figure 1, is a work in progress.

Task Selection

In our taxonomy, the design of navigation is treated as a task-based model. The designer begins by asking why the user is trying to navigate. We believe that navigation may be broken into three subtasks – exploration, to gain survey knowledge; search, to locate an object or route and to travel to it; and inspection, to establish and maintain a particular view of an object. Only in considering these subtasks can the designer, aware of the user’s goal, formulate the abstract solution satisfying the goal, and design the appropriate interface.

Travel Control

The abstract solution to the user’s goal can be defined in terms of travel control. We define travel as viewpoint motion (specified by position, orientation, and speed of the virtual camera) through a Virtual Environment. The designer must carefully construct functions that specify each of these factors. Each of these functions can be defined in terms of properties of the user, the environment, or a combination of the two. It is precisely this separation of position and orientation control, along with the consideration of different (and possibly non-linear) functions and frames of reference that allow us to create fundamentally new techniques.

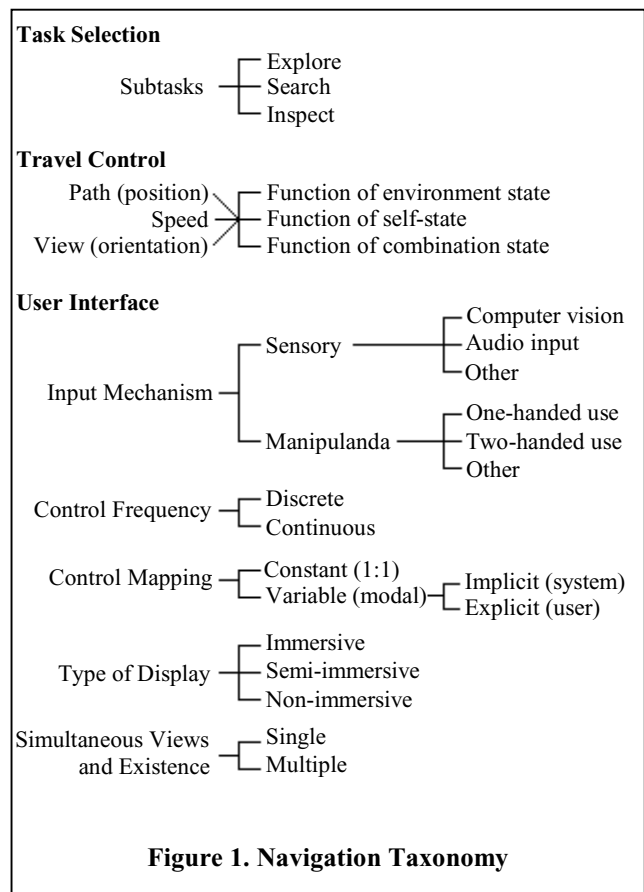


Figure 1. Navigation Taxonomy

User Interface

Besides deciding on the abstract solution to the user’s goal, the user interface and navigation metaphor must be considered. The input hardware determines the real-world affordances presented to the user as well as the degrees of freedom (DOF) over which the user has control. Within these constraints, the designer can decide on an intuitive mapping between input and camera control. Frequently there is insufficient input DOF for the task and user input must be moded. If the system can reliably determine the user’s goal, it can automate the moding; otherwise the user will have to be given explicit control of the different modes.

Output is a factor often overlooked in the design of navigation techniques. We assert that the means of output and the level of immersion that the user experiences greatly affect the navigation task and performance. The more immersed the user is, and the more easily the user can mentally integrate information acquired, the greater the chances of efficient navigation [21]. The goal of the designer, then, must be to present accurate information in a manner that is timely, intuitive, and informative to the user. Hence, we began experimenting with some very large displays to observe whether the output hardware would indeed contribute as significantly as predicted.

TECHNIQUES CONSIDERED

Inspection: Object Manipulation and Ghost Copy

Point-of-interest techniques allow users to specify objects in the environment toward which they would like to move [17]. In our taxonomy, this would be a function of environment state for position and orientation control since the user is moved directly to the target destination. Typically, a logarithmic function is implemented for speed control so that the speed is scaled with respect to the user's distance to the target in the environment. But, these techniques generally do not allow users to specify the desired position and orientation with respect to target objects. We extend this technique, including information about the user and the environment, to solve this problem.

In the Object Manipulation technique, the user drags on an object and manipulates a copy of the object, placing it in the position and orientation from which she would like to view it. When she releases the button, the copy fades away and the camera animates to the appropriate position and orientation in the world to attain the desired viewpoint of the object. In the Ghost Copy technique, the user drags and manipulates multiple copies of an object so as to get multiple simultaneous viewpoints of it. Since this is designed to be a relatively lightweight tool, the copies are destroyed when the user moves away or selects another object to view.

Search: Inverse Fog/Scaling and Ephemeral World Compression

In complex or dense worlds, searching for objects is difficult partially because many objects are occluded from any given viewing position. In the Inverse Fog and Scaling techniques, the user controls the radius of a sphere centered on her. Objects that fall within this sphere are either made transparent (in Inverse Fog) or scaled down (in Inverse Scaling) so as to reduce the amount of occlusion that objects within the sphere present. The closer to the center of the sphere (where the user is), the more transparent or smaller the objects become. This allows users to scan through worlds very quickly by simply pivoting or turning and removing occlusions so as to see objects that are different distances away. This technique can be moded with traditional techniques (such as point-of-interest or regular driving navigation) for efficient search.

In Ephemeral World Compression, the user is given the ability to scale the world in ways that aid in the task at hand. When trying to gain an understanding of the structure of the world or when searching for objects, the user can scale the entire world down so that she gets an overview of the world. Alternatively, the user can compress it radially around the camera so that all the objects fit within the viewing frustum. The user can navigate within this scaled down world and return the world to normal when she is where she would like to be.

Moded Tasks: Possession and Rubberneck Navigation

Existing techniques do not allow the user to see the world from the point of view of another actor or object. In the Possession metaphor, we define functions of environment state for both position and orientation to enable this capability. In this technique, the user clicks on an object to move to the exact position and orientation of that object, thereby 'possessing' it and seeing the world from the point of view of that object. Automatically determining point of view is a difficult problem and we would ideally like to use designer supplied hints. For example, it may be obvious that possessing a bunny object should cause the user to look through the eyes of the creature, but this is not as clear in the case of a Ferris Wheel. We currently place the user looking forward on the main pivot point of the object. In an intuitive moding, the user can also create new objects (usually cameras) on the fly by clicking anywhere that is not an object. She can position these objects before flying into and possessing them. In single user worlds, this allows users to quickly explore the world and search for landmarks, as well as to inspect objects from specific points of view. In multi-user worlds, users can position objects to suggest viewpoints that other users might be interested in.

The Possession navigation technique can be extended so that the user inherits the form as well as capabilities of her new "being." This presents interesting possibilities for further input moding and control, as the user possesses different canonical objects. Another extension to this technique might be to allow the user to use traditional navigation while possessing objects, thus allowing for egocentric object manipulation.

In existing techniques, the virtual body and head are controlled with separate mechanisms, making it difficult to walk about an environment while looking around. Rubberneck Navigation is a technique we developed earlier with colleagues from several labs (see Acknowledgements), which combines discrete movement control with continuous viewpoint control to allow users to use standard 2D input devices (mouse, trackball, touch-pad) to manipulate body as well as head movement. In this technique, the user moves the mouse to look around. She can at any time use either point-of-interest navigation to specify a movement path or Igarashi's Path Drawing technique [13], in which she holds down the button to lock the viewpoint and drags to draw a path along which to move. Once in motion, movement is decoupled from viewpoint and the user can continue to look around and explore the environment. This technique is also particularly useful in moving around an object while continually controlling the viewpoint to look at it. In an independent effort, Cohen et al. have implemented a technique similar to this in their Harold system [5].

Moded Tasks: Speed-coupled Flying with Orbit

How can a user easily and intuitively alternate between walking around an environment, getting an overview, and examining objects? In considering the control mapping, we



Figure 2. Local view of scene while standing still or moving slowly.

decided that a natural way to separate the intention to freely navigate from that of constrained object inspection is to determine where the user first starts dragging the mouse. If the user drags on free space (including the sky or ground plane), the system assumes that the user is trying to freely move around the environment. If, however, the user starts dragging the mouse on an object, the system assumes that the user is trying to examine it. We indicate the presence of this moding by changing the cursor as the user moves it around to match the function that would be called if the user were to start dragging at that point.

A particularly interesting technique for exploring 3D space that grew out of thinking about the various functions of travel control in relation to one another is Speed-coupled Flying. In this technique, we couple speed control to height (position) and tilt (viewpoint) control to give the user the ability to transition seamlessly between and navigate within local as well as global views of the world. This allows the user to acquire landmark, procedural, and survey knowledge and to effectively perform exploration and search tasks.

The user is equipped with standard egocentric driving controls (i.e., dragging the mouse forward/backward moves the camera forward/backward, dragging the mouse left/right turns the camera left/right). The further the user drags in a particular direction, the faster the camera moves. Furthermore, in our technique, the user's forward speed is coupled to the camera's viewing height and angle. The faster the user moves forward, the higher the camera moves, getting a zoomed out overview of the environment. Contrast Figures 2 and 3 to get a feel for the difference between these two views. This coupling of speed to height and tilt keeps the visual flow across the screen constant, in an effort to allow the user to move and locate distant targets quickly. The basic technique as described is similar to Igarashi's speed-dependent scrolling [15]. However, we improve on that technique in two ways.



Figure 3. Overview of scene while moving fast.

As the user slows down or releases the button, the camera glides back down to the ground (or the level where it was before flying). Informal observations suggested that users used backwards movements mainly to navigate small distances after overshooting targets. In this case, they did not need an overview and did not expect to start flying.

Also, when releasing the button to land, users felt uncomfortable simply falling straight down to the ground. Igarashi [15] observed a similar problem in his system. Because of this problem, users tended to either overshoot or undershoot and be forced to take corrective action. Based on these observations from our first user study, we implemented an inertial based system in which the camera glides forward and eventually lands at a point in the world that was at the bottom of the view frustum when the button was released. Users had no need to take corrective action when the glide release was included in the technique.

In addition, we combine the speed-coupled flying technique with an object inspection technique we term *orbiting*. Orbiting allows the user to easily get desired viewpoints of particular target objects. When the user clicks and drags on an object, the object animates to the center of the screen and the user is switched to a mode of environmental state navigation. In this mode, dragging the mouse forward/backward moves the camera toward/away from the object; dragging left/right causes the user to move around a circle parallel to the ground plane and centered at the object, always keeping the viewpoint fixed on that object.

This technique is the combination of several points in the design space of 3D navigation. Our target was for planar tasks in 3D worlds, such as those found in chat rooms, architectural walkthroughs, or geographic data sets. This technique must be carefully augmented if used for higher dimensional non-planar tasks, such as finding a spacecraft in an asteroid belt.

EXPERIMENT 1

Given the novelty of these techniques, we ran two user studies to evaluate them. The studies show that the last of these techniques, Speed-couple Flying with Orbiting, allows

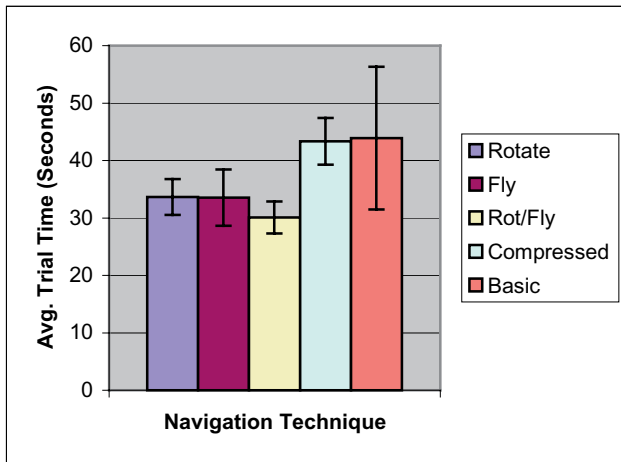


Figure 4. Results from Experiment 1 showing trial time for each technique.

users to more efficiently navigate in certain virtual environments, when compared to basic or other novel navigation techniques.

Method

Subjects

Seventeen experienced Windows users (7 female) participated in the study. All but three either did not play 3D games or played them very little (less than a half hour a week). On average, the 3 participants that played 3D games played them around 5 hours a week. The average age of the participants was 37.5, ranging from 11 to 60 years of age.

Materials and Procedure

Five navigation techniques were evaluated in this study: basic driving navigation as provided by Alice [1] (with and without Orbiting), Speed-coupled Flying (with and without Orbiting), and Ephemeral World Compression. Two 3D worlds were created with the Alice 3D authoring tool [1]. The first world, which we will refer to as the tutorial world, was 300 by 300 meters and contained 4 objects for navigation and manipulation purposes. The second world, which we refer to as the experimental world, was 500 by 500 meters and contained 23 objects, most of which consisted of carnival-themed structures, such as tents, roller coasters and rides (see Figures 2 and 3). Each world was designed so that there were “target” cubes and “target” drop-pads. The cubes and drop-pads were dual-coded to match each other via color and numeric coding. The cubes, numbered only on one face, were to be selected and carried to the matching numbered drop-pad, where it would be placed. In the tutorial world, there were only 2 pads and 1 cube for each trial. The tutorial consisted of the user finding the cube and placing it on its corresponding pad once for each of the five navigation conditions. In the experimental world, there were 4 cubes and 4 pads. A trial consisted of

the user successfully finding each cube (in any order) and dropping it on its respective drop-pad. Once all 4 cubes had been placed on all 4 drop-pads, the trial was completed and an interim screen was presented (the screen displayed the message, “Click to Continue”, at which time the next trial would begin). Each of the five navigation types was presented once in random order during the experimental session. In between each navigation condition (trial), the participant filled out a brief series of 3 questions about that navigation condition prior to proceeding to the next trial. A trial deadline of 5 minutes was imposed during the experimental tasks in order to ensure that participants progressed through the session in a timely manner. When all trials were completed, the participant filled out more survey questions, was debriefed and provided with a software gratuity for participating.

The study was run on 450 MHz Pentium II Dell computers with 15-inch Sony flat panel monitors. The Microsoft natural keyboard and Intellimouse input devices were used (although the wheel on the Intellimouse was not used in the study). Four subjects also experienced the navigation techniques using an experimental 39-inch display (~36 x 14 inches). This display was run at 2048 x 768 pixels, twice the resolution of the small display, and twice the virtual field of view (~75 degrees). This was exploratory and will be documented in depth in Experiment 2.

All dependent measures of interest were recorded on the participant’s computer. Because we were interested in improvements in navigation ease compared to the basic Alice driving navigation technique, we collected the following measures: overall task time, cube search time, cube carry time, cube carry distance (shortest possible path), actual cube carry distance, the number of cubes retrieved and placed before the deadline (percent correct) and user satisfaction questionnaire responses for each condition as well as overall preference.

Results and Discussion

Overall Trial Time and Percent Correct

The average trial time for each navigation technique is shown in Figure 4. A repeated measures Analysis of Variance (ANOVA) was performed on the overall trial time data. A main overall effect of navigation technique was observed, $F(4,60)=4.25$, $p<.01$. Paired comparisons using the Bonferroni technique showed multiple significant differences: Flying with Orbit was reliably faster than Basic navigation ($p<.05$), Basic navigation with Orbit ($p<.01$), and Ephemeral World Compression ($p<.001$). In addition, Flying was reliably faster than Basic with Orbit ($p<.01$) and Ephemeral World Compression ($p<.01$). A repeated measures ANOVA was performed on the percent correct data, revealing a lack of a significant main effect of navigation technique on accuracy, $F(4,64)=1.64$, $p>.05$.

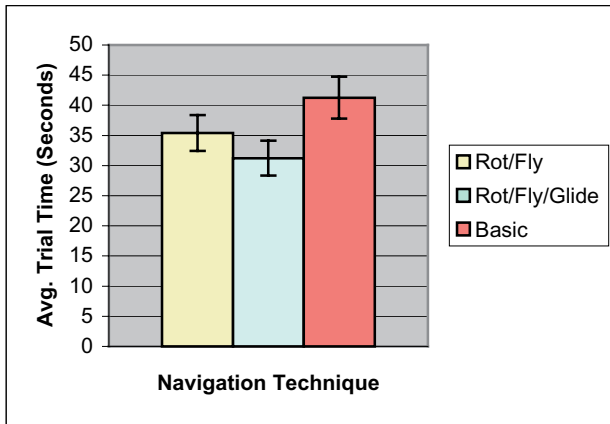


Figure 5. Results from Experiment 2 showing trial time for each technique.

Carry Distance (Theoretical minus Actual)

We subtracted participants’ actual distance over which they carried the cubes from the shortest possible path they could have taken (“theoretical distance”), and subjected this data to an ANOVA. A significant main effect of navigation technique was observed, $F(4,60)=2.5$, $p<.05$, although no paired comparisons reached significance using the Bonferroni correction for multiple tests.

User Satisfaction

There was a significant preference in the user satisfaction data for the Flying with Orbit navigation technique (preferred by 14/17 users, significant by binomial test). The 3 users that did not prefer the Flying with Orbit technique chose the Flying only technique as their most preferred. The user satisfaction data collected after participants used each navigation technique were submitted to an ANOVA. A significant main effect of navigation condition was observed, $F(4,56)=3.1$, $p<.05$. However, once again no single condition was rated reliably higher than the others once the Bonferroni technique was used in the paired comparisons test. Participants noted in their survey that they had particular problems knowing where they would “land” when flying. This usability issue was addressed in the prototype tested in Experiment 2.

Big Display v. Small Display

As was mentioned, four users experienced the experimental world while viewing a large, experimental display. Although the overall pattern of their data did not differ significantly with regard to the navigation conditions studied, there were some interesting trends in their data that encouraged us to run another study using display size as a variable of interest. While it has been shown that large projection screens may be effective substitutes for immersive displays such as head-mounted displays [20], we know of little work done to quantify the differences between these large semi-immersive displays and regular non-immersive displays. On average, overall trial times were 26.7 seconds using the big display, versus 38.6 using

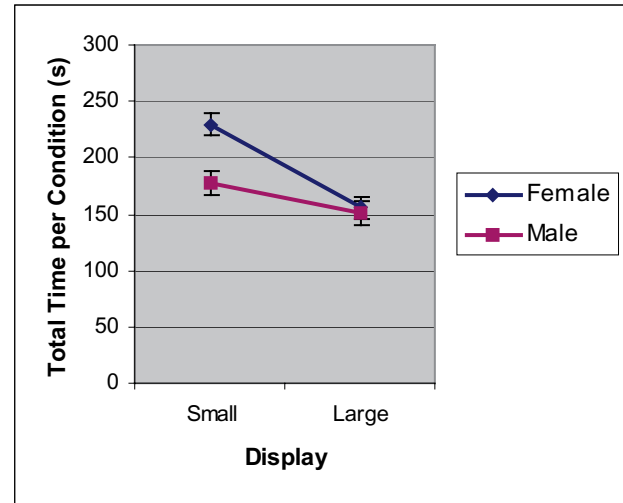


Figure 6. Results comparing small versus large displays for male and female participants.

the smaller display. Since there were too few participants to make generalizations, we decided to run Experiment 2 using display size as an explicit factor.

EXPERIMENT 2

Method

A goal in Experiment 2 was to replicate the finding that the Flying with Orbit navigation technique was a significant improvement beyond Basic navigation in 3D worlds. In addition, we refined the Flying with Orbit condition by adding a “gliding” behavior to the Flying with Orbit condition’s landing behavior. We thought that the smoother landing might alleviate users’ problems knowing where they were going to land, as observed in Experiment 1. Finally, we included the large display as a variable in this study, in order to more rigorously assess the preliminary findings reported for Experiment 1 with regard to the benefits of the larger display. Therefore, Experiment 2 was a 3 (Basic, Flying with Orbit, and Flying with Orbit and Glide navigation techniques) x 2 (big v. smaller display) within subjects design. The navigation technique order to presentation was randomized, and the display order was counterbalanced across participants. All other aspects of the experiment were identical to Experiment 1, including the dependent measures collected during the experimental sessions. Thirteen participants (6 female), with an average age of 31.6 (range 19-53 years of age), participated in the study for approximately 1.5 hours.

Results and Discussion

Overall Trial Time

One subject was discarded from analysis because she failed to complete any of the six conditions. The remaining twelve subjects (7 men, 5 women) combined only failed to complete a condition a total of six times. The average trial time for each navigation technique is shown in Figure 5. A repeated measures analysis of total time to complete each

navigation technique (i.e., for putting four blocks on their respective landing pads) revealed a significant main effect, $F(1,11) = 221.1$, $p < 0.001$. Basic Navigation and Flying with Orbit were significantly different, as measured through pair-wise comparisons using the Bonferroni correction for multiple tests. In addition, Flying with Orbit was borderline reliably different from Basic Navigation ($p=.058$).

Big v. Smaller Display

The total time to complete a task was significantly affected by display type $F(1,11) = 21.8$, $p < 0.001$. Participants were reliably faster with the larger display (152.8 seconds versus 199.0 seconds with the smaller display). We believe this was largely due to the increased field of view of the larger display (~75 degrees of visual angle from edge to edge of screen for the large display, ~33 degrees for the small display). The peripheral vision afforded by the larger field of view (both physical and virtual) was important for the search and navigation tasks because users were able to see more of the environment at any given time. With the larger field of view, they seemed better able to sample and identify objects and did not miss as many objects. Also, although not explicitly tested, we hypothesize that users in this condition had a better sense of their movement and location in the environment. This would corroborate findings reported of egocentric navigation with large fields of view in virtual environments [32].

Although there was no statistical difference in time to task completion based on gender there was a significant interaction between gender and display type, $F(1,10) = 10.8$, $p < 0.01$. Time to complete was the same for men and women using the large display (151 s versus 155 s), but men were faster on the small display (177 s versus 230 s). These results are summarized in Figure 6. It may be that larger displays can somehow compensate for any gender differences observed during 3D navigation, although the factor(s) behind this effect will need to be isolated and explored more thoroughly in future research.

Carry Distance (Theoretical minus Actual)

There was also a significant main effect of distance traveled (compared to theoretical distance) in meters, $F(1,6) = 63.3$, $p < 0.001$. A significant main effect for display, $F(1,6) = 6.6$, $p < 0.05$ was obtained. With the large display, subjects traveled fewer unnecessary meters (336 m) than when using the small display (405 m) on average, per cube. There was also a significant main effect for navigation type, $F(1.09,6.54) = 8.62$, $p < 0.05$. Participants traveled quite a bit more distance in the Basic Navigation task, on average 630 extra meters. When traveling in the Flying with Orbit and Flying with Orbit and Glide conditions, on average, participants traveled 230 m and 252 m extra per cube.

User Satisfaction

There were no significant main effects in terms of user satisfaction ratings across the various conditions tested, $F(2, 72)=.6$, $p=.7$. However, users significantly preferred the Flying with Orbit and Glide navigation technique (preferred

by 8/13 participants, 4 chose Flying with Orbit and 1 chose Basic navigation). Also 9/13 participants chose the large display as their favorite (borderline significant via binomial test, $p=.087$). The large display was significantly more often chosen as presenting the most information (10/13 participants, $p<.03$, binomial test). Only 7 participants chose the large display as the most efficient for the tasks in the study, 5 chose the smaller display.

FUTURE DIRECTIONS

We believe that there is much work that can be done to refine the taxonomy presented, as well as to further explore uncharted areas in the design space. The large prototype display we have constructed has been shown to enhance performance with navigation tasks, especially with female users. One of the areas we have grown particularly interested in as a result of these usability results is the fabrication and testing of new display devices, both as they apply to 3D navigation and also more generally to conventional desktop tasks. We would like to isolate and better understand the effects of various characteristics of different display devices.

CONCLUSION

In this work, we have presented a taxonomy of navigation techniques, as work in progress. We have described several new techniques that have grown out of this taxonomy, including Object Manipulation, Ghost Copy, Inverse Fog/Scaling, Ephemeral World Compression, Possession, Rubberneck Navigation, and Speed-coupled flying combined with Orbiting. User studies have shown that the last of these techniques, Speed-coupled Flying with Orbiting allows users to more efficiently navigate in certain environments. These studies have also revealed that our prototype large screen displays enhance performance with navigation tasks, warranting further investigation.

ACKNOWLEDGEMENTS

We would like to thank Gary Starkweather for building the large prototype display that was used in testing; Brent Field for assisting in user testing and data analysis; and David Thiel for invaluable assistance in producing the video figure. We would also like to acknowledge work done with us by Mathew Conway, Ken Hinckley, Takeo Igarashi, Dennis Proffitt, and Randy Pausch to develop the Rubberneck Navigation technique.

REFERENCES

1. Alice 3D Authoring Tool, see <http://www.alice.org/>.
2. Bederson, B., Hollan, J., Perlin, K., Meyer, J., Bacon, D., and Furnas, G. Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics, *Journal of Visual Languages and Computing*, 7 (1996), 3-31.
3. Bowman, D., Koller, D., and Hodges, L. Travel in Immersive Virtual Environments: An Evaluation of Viewpoint Motion Control Techniques. *Proceedings of*

- the 1997 Virtual Reality Annual International Symposium (VRAIS)*, 45-52.
4. Chapman, D., Colin W. Manipulating the future: predictor based feedback for velocity control in virtual environment navigation. *Proceedings of the 1992 Symposium on Interactive 3D Graphics*, pp 63 – 66.
 5. Cohen, J., Hughes, J., Zeleznik, R. Harold: A World Made of Drawings. *Proceedings of the First International Symposium on Non-Photorealistic Animation and Rendering*, 83-90, 2000.
 6. Darken, R., Sibert, J. A Toolset for Navigation in Virtual Environments. *Proceedings of the UIST '93*, ACM Press, 157-165.
 7. Darken, R., Sibert, J. Navigating in Large Virtual Worlds. *The International Journal of Human-Computer Interaction*, 8(1), 49-72, 1996.
 8. Elvins, T., Nadeau, D., Schul, R., Kirsh, D. Worldlets: 3D Thumbnails for 3D Browsing. *Proceedings on CHI '98*, ACM Press, 163-170.
 9. Fukatsu, S., Kitamura, Y., Masaki, T., Kishino, F. Intuitive Control of Bird's Eye Overview Images for Navigation in an Enormous Virtual Environment. *Proceedings of the 1998 ACM Symposium on Virtual Reality Software and Technology*, 67-76.
 10. Furnas, G. Generalized Fisheye Views. *Proceedings of CHI '86*, ACM Press, 16-23.
 11. Furnas, G. Effective View Navigation. *Proceedings of CHI '97*, ACM Press, 367-374.
 12. Gleicher, M., Witkin, A. Through-the-lens Camera Control. *Proceedings of SIGGRAPH '92*, ACM Press, 331-340.
 13. Hunt, E., Waller, D. Orientation and wayfinding: A Review. ONR Technical Report N00014-96-0380. 1999.
 14. Igarashi, T., Kadobayashi, R., Mase, K., and Tanaka, H. Path Drawing for 3D Walkthrough. *Proceedings of UIST '98*, ACM Press, 173-174.
 15. Igarashi, T., Hinckley, K. Speed-dependent Automatic Zooming for Browsing Large Documents, *Proceedings of UIST 2000*, ACM Press, 139-148.
 16. Jul, S., Furnas, G. Critical Zones in Desert Fog: Aids to Multiscale Navigation. *Proceedings of UIST '98*, ACM Press, 97-106.
 17. Mackinlay, J., Card, S., Robertson, G. Rapid Controlled Movement through a Virtual 3D Workspace. *Proceedings of SIGGRAPH '90*, ACM Press, 171-176.
 18. Mackinlay, J., Robertson, G., Card, C. The Perspective Wall: Detail and Context Smoothly Integrated. *Proceedings of the CHI '91*, ACM Press, 173-179.
 19. Masui, T., Minakuchi, M., Borden, G., Kashiwagi, K. Multiple-view Approach for Smooth Information Retrieval. *Proceedings of UIST '95*, ACM Press, 199-206.
 20. Patrick, E., Cosgrove, D., Slavkovic, A., Rode, J., Verratti, T., Chiselko, G. Using a Large Projection Screen as an Alternative to Head-mounted Displays for Virtual Environments. *Proceedings of CHI 2000*, ACM Press, 478-485.
 21. Robertson, G., Czerwinski, M., van Dantzich, M. Immersion in Desktop Virtual Reality. *Proceedings of UIST '97*, ACM Press, 11-19.
 22. Robertson, G., MacKinlay, J. The Document Lens. *Proceedings of UIST '93*, ACM Press, 101-108.
 23. Robinett, W., Holloway, R. Implementation of Flying, Scaling and Grabbing in Virtual Worlds. *Proceedings of the 1992 Symposium on Interactive 3D Graphics*, 189-192.
 24. Ruddle, R., Payne, S., Jones, D. The Effects of Maps on Navigation and Search Strategies in very-Large-scale Virtual Environments. *Journal of Experimental Psychology: Applied*, 5, 1999, 54-75.
 25. Siegel, A., White, S. The development of spatial representations of large-scale environments. In H. Reese (Ed.), *Advances in Child Development and Behavior*, 10, 10-55. New York: Academic Press.
 26. Song, D., Norman, M. Nonlinear Interactive Motion Control Techniques for Virtual Space Navigation. *Proceedings of 1993 IEEE Virtual Reality Symposium*, IEEE Computer Society, 111-117.
 27. Stoakley, R., Conway, M., Pausch, P. Virtual Reality on a WIM: Interactive Worlds in Miniature. *Proceedings of CHI '95*, ACM Press.
 28. Thorndyke, P., Hayes-Roth, B. Differences in Spatial Knowledge Acquired from Maps and Navigation. *Cognitive Psychology*, 14, 1982, 560-589.
 29. Vinson, N. Design Guidelines for Landmarks to Support Navigation in Virtual Environments. *Proceeding of CHI '99*, ACM Press, 278-285.
 30. Ware, C., Osborne, S. Exploration and Virtual Camera Control in Virtual Three Dimensional Environments. *Proceedings of the 1990 Symposium on Interactive 3D Graphics*, 175-183.
 31. Ware, C., Fleet, D. Context Sensitive Flying Interface. *Proceedings of the 1997 Symposium on Interactive 3D graphics*, 127.
 32. Wickens, C.D. and Hollands, J.G. *Engineering Psychology and Human Performance*, 3rd Ed. Prentice-Hall, N.J., 1999, pp. 158-171.
 33. Witmer, B., Bailey, J., Knerr, B., & Parsons, K. Virtual spaces and real world places: Transfer of route knowledge. *International Journal of Human-Computer Studies*, 45, 413-428.