Exploring Large Virtual Environments with an HMD when Physical Space is Limited

Betsy Williams,[*] Gayathri Narasimham,[*] Bjoern Rump,[*] Timothy P. McNamara,[*]
Thomas H. Carr,[*] John Rieser,[*] and Bobby Bodenheimer[*]
Vanderbilt University

Abstract

Virtual Environments presented through head-mounted displays (HMDs) are often explored on foot. Exploration on foot is useful since the afferent and efferent cues of physical locomotion aid spatial awareness. However, the size of the virtual environment that can be explored on foot is limited to the dimensions of the tracking space of the HMD unless other strategies are used. This paper presents a system for exploring a large virtual environment on foot when the size of the physical surroundings is small by leveraging people's natural ability to spatially update. This paper presents three methods of "resetting" users when they reach the physical limits of the HMD tracking system. Resetting involves manipulating the users' location in physical space to move them out of the path of the physical obstruction while maintaining their spatial awareness of the virtual space.

CR Categories: I.3.m [Computer Graphics]: Miscellaneous—Perception

Keywords: Virtual reality (VR), space perception

1 Introduction

Virtual reality provides people with opportunities to experience places and situations remote from their actual physical surroundings. They potentially allow people to learn about an environment which, for reasons of time, distance, expense, and safety, would not otherwise be available. Virtual reality systems could have a huge impact in education, entertainment, medicine, architecture, and training, but they are not widely used because of their expense and delicacy. However, head-mounted display (HMD) technology may become readily available to the public within the next several years. Other immersive virtual technologies, such as virtual caves, are less likely to achieve commodity status since they often involve greater expense in the form of large screens, projectors, and a locomotion device such as a treadmill or bicycle that allows a user to move about the environment. Since HMD systems hold the promise of being readily available to the public, constraints of the system need to be identified and addressed. A major drawback of HMD-based systems is the limited amount of space available for exploration. This work explores one way of manipulating virtual reality to extend the capabilities of an HMD system to explore and experience virtual spaces larger than the physical limits of the HMD tracking system. Specifically, this work looks at "resetting" subjects when they reach the limits of the tracking system. Resetting involves manipulating optical flow in such a way that allows users to move away from a physical obstruction such as a wall while experiencing a continuous sense of their location in virtual space.

HMD-based virtual environments are often explored on foot. Foot exploration is useful since the inertial cues of physical locomotion aid in spatial awareness. The size of the virtual environment that can be explored is limited to the dimensions of the tracking space of the HMD unless some other method of exploration is used. One obvious solution to exploring a large virtual environment is to use a joystick to translate freely in the virtual environment. This method has been shown to be inferior to physical locomotion [Chance et al. 1998; Ruddle and Lessels 2006; Lathrop and Kaiser 2002]. Other methods, e.g., [Templeman et al. 1999; Slater et al. 1995; Razzaque et al. 2001; Nitzsche et al. 2004; Usoh et al. 1999], have also been proposed, and we explore these in Section 2. Our previous work, [Williams et al. 2006], manipulated the translational gain of walking, so that one step forward in the physical environment corresponds to several steps forward in the virtual environment. Two experiments showed that increasing the translational gain of walking is a useful method of navigating large virtual spaces, and that it is superior to joystick exploration. We generally find that subjects' spatial orientation was similar in normal walking and walking with translational gain of ten, but to explore a battlefield or a city with this technique, additional strategies would be needed. Moreover, the physical limits of the tracking system may be reached no matter how high gain is scaled. Thus, this paper presents work on one such additional strategy that resets subjects when they walk and reach the end of the physical space. With this strategy, we are assessing the ability of people to rely on visual information for spatial updating during these resets. We evaluate three methods of resetting position while subjects walk in small physical tracking spaces in order to explore large virtual spaces. We call these three methods Freeze-Backup, Freeze-Turn, and 2:1-Turn.

After completing a reset, users travel along the same virtual path they had been traveling. In the Freeze-Backup method, the user obtains more space for virtual exploration by taking steps backwards while frozen in a fixed position in the virtual environment. In the other two methods, Freeze-Turn and 2:1-Turn, users overcome physical obstruction by physically turning 180° and maintaining their same position in virtual space before and after the turn. During a Freeze-Turn reset, the orientation of the user is frozen while the subject turns 180°. In the 2:1-Turn condition, the gain of the turn is doubled so that 180° turn in physical space corresponds to 360° turn in virtual space.

This work is important because our future goal is to extend and integrate the results of this paper with our prior work on scaling the translational gain of walking [Williams et al. 2006]. The resulting system should allow a person to seamlessly explore large virtual environments. The system envisioned here could be based in an office or small lab. In particular, if immersive virtual environments are to realize their potential as commodity-level components, a perceptually accurate interface that allows locomotion through them within the constraints of everyday space must be developed.

[*]{betsy.williams, gayathri.narasimham, bjoern.rump, t.mcnamara, tom.carr, j.rieser, bobby.bodenheimer}@vanderbilt.edu

Copyright © 2007 by the Association for Computing Machinery, Inc.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions Dept, ACM Inc., fax +1 (212) 869-0481 or e-mail permissions@acm.org.
APGV 2007, Tübingen, Germany, July 26–27, 2007.
© 2007 ACM 978-1-59593-670-7/07/0007 $5.00

RIGHTSLINK

2 Background

Previous research has explored various techniques of navigating a large scale virtual environment. Haptic devices, such as a joystick or keyboard, allow users to virtually explore large environments [Ruddle et al. 1999; Bowman et al. 1999; Waller et al. 1998; Darken and Sibert 1996], but studies have shown that using physical bipedal locomotion rather than haptic devices produces significantly better spatial orientation [Chance et al. 1998; Ruddle and Lessels 2006; Lathrop and Kaiser 2002]. Templeman et al. [Templeman et al. 1999] and Slater et al. [1995] have participants "walk in place" to move through large virtual environments, but this technique lacks the same proprioceptive cues of walking. Another method of navigating a large virtual environment is manipulating rotation such that the locomotion of the subject fits within the limits of the tracking system [Razzaque et al. 2001; Nitzsche et al. 2004]. Razzaque et al. [Razzaque et al. 2001] examine subjects ability to locomote to a series of five targets they call "waypoints". In this study, the virtual room is slightly rotated while the subject walks to the waypoint, and then to a greater degree as the subject searches for these waypoints. This method requires a large tracking area for the rotational manipulation to be imperceivable, and is not a complete solution because a situation could easily occur in which the physical limits of the tracking system are reached. Virtual flying [Usoh et al. 1999] and teleporting are other ways of exploring large virtual environments, yet they lack locomotive feedback. Other systems involve large screen caves with a locomotion input such as a bicycle or treadmill. Cave systems are expensive, and most only contain three virtual walls. Treadmill systems are difficult and expensive to construct with enough degrees of freedom to allow for free exploration.

In this study, we scaled physical rotation. Kuhl [2004] and Pick et al. [1999] have shown that people can recalibrate rotations. Although we are not looking for recalibration, this work shows that people can maintain spatial orientation when the rotational gain of turning is not their own. Research shows that physical changes in direction are more important than physical translation in the development of spatial knowledge [Presson and Montello 1994; Rieser 1989; Rieser et al. 1995]. This finding is important because we manipulate translations in one resetting condition and manipulate rotations in the other two. The experiment presented in this paper uses a spatial orientation task where subjects turn to face a direction, similar to the pointing task of Rieser and others [Rieser 1989; Klatzky et al. 1998; Kearns et al. 2002].

3 Resetting Methods

Three resetting methods are evaluated. Resetting involves physical locomotion with optical manipulated flow in such a way that the user's sense of where they are relative to objects in their virtual environment is not changed. The three resetting methods are called Freeze-Backup, Freeze-Turn, and 2:1- Turn and explained as follows.

1. **Freeze - Backup.** In this method, the computer indicates to the user that they have reached the boundaries of the tracking system and needs to reset. The tracking system is no longer used to update the position of the subject in the virtual environment , so that the user's position in virtual space is no longer updated with movement in physical space. The user is then instructed to take steps backwards in physical space while user's position in virtual space remains fixed or frozen. When enough steps are taken, the computer indicates for the user to stop, the displays are unfrozen, and the user is allowed to continue along the same path that they were walking before the reset. During the backward walking, orientation tracking is active so that the user can look around.

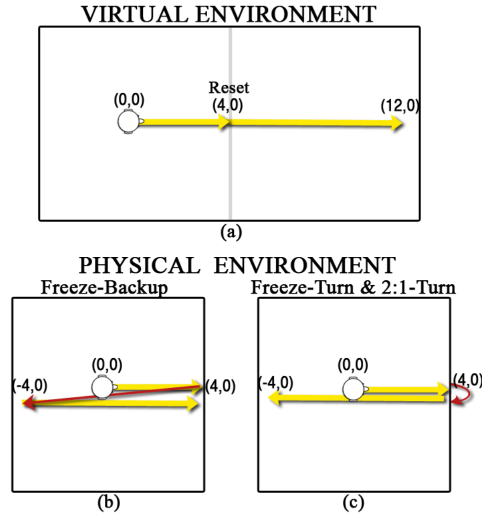

Figure 1: The path a subject perceives they have taken in the virtual environment is shown in (a), while the path in physical space that the subject takes under the different resetting methods is shown in (b) and (c). In this example, a person at position (0,0) in physical space views the virtual environment at position (0,0). In (a), the person walks forward in the virtual environment where they are alerted by a signal at (4,0) indicating they are near the tracking limits and need to reset their position in physical space. The person then continues walking to (12,0) in the virtual environment. The corresponding paths in the physical environment for the three resetting methods are shown in (b) and (c). Red arrows indicate physical movement during a reset.

Figures 1(a) and 1(b) show an example of the process. The rectangle shown in Figure 1b represents the physical limits of the tracking system while the larger rectangle shown in Figure 1a represents the virtual environment. In this example, the user starts at position (0,0) in both the real and virtual environments. The user then proceeds to (4,0) but cannot explore further because the limits of the physical space have been met. Therefore, the user undergoes a reset, and their position is frozen at (4,0) as they follow the red path and back up to (-4,0). During the backup phase, the user is instructed to simply walk backward and told when to stop, and is not guided backward. Thus, the user does not typically walk a straight path directly behind them as in this example. Once the user reaches (-4,0), the system instructs the user to stop, and the user continues along the the yellow path until they reach (4,0). The corresponding path in virtual environment from (0,0) to (12,0) is show on the right.

The physical position of a user in x, y, z space using a right-handed coordinate system is obtained from the tracking system. The position in the center of the room on the floor is (0, 0, 0). The x, y, and z directions while standing in the center of the room facing to the front correspond to front to back movement, user height, and right to left movement, respectively. Movement is limited in the x and z directions due to the finite range of the tracking system. Since the y-direction indicates movement perpendicular to the ground pane, this value typically represents the user's eyeheight, and does not limit the exploration of the virtual environment. Orientation is obtained from the rotational sensor located on HMD which updates rotation about the x-axis (pitch), y-axis (yaw), and z-axis (roll). The details of the algorithm are in Appendix A.1.

RIGHTSLINK

2. **Freeze - Turn.** In this method, when the tracking device finds that the subject has reached a boundary, the computer indicates to the participant that they need to reset by turning around. The display of the HMD is frozen, freezing the participant's position and yaw angle in virtual space, and the participant turns 180 degrees. The display is unfrozen, tracking is updated, and the subject is able to continue traveling along his route.

   Figure 1(a) and 1(c) show an example of the area walked in the physical space and the corresponding area walked in the virtual space. In this example, the user starts in position (0,0) in both the physical and virtual environments viewing the virtual world through the HMD. The user walks to position (4,0) where there is no more physical space and desires to continue along this same path. Thus, the user turns around with a display frozen in that y-direction at position (4,0) to reset. When the screen is unfrozen, the user has turned 180° and continues walking to (-4,0) in the physical space. The yellow path in Figure 1a shows the locomotion that the user perceived in the virtual environment. Appendix A.1 explains the algorithm by which this reset is accomplished.

3. **2:1 - Turn.** In this method, when the subject reaches the boundaries of the tracker, the computer indicates that they should turn and keep turning until completing a visually full turn in the virtual environment. The rotational gain of the yaw angle during this turn is scaled by two, such that the user rotates 180° in the physical environment, but rotates 360° in the virtual environment. See Appendix A.3 for details regarding the algorithm.

   Figure 1(a) shows an example of a path taken in a virtual environment where one reset is undergone. The corresponding physical path is seen in Figure 1(b). Note that the path taken by the Freeze-Turn and 2:1 conditions are similar since they both involve the user turning around to reset.

4 Experimental Evaluation

Since all three of these methods are reasonable forms of resetting, we conducted an experiment to evaluate which one worked best. Additionally, we wanted to determine what the cognitive cost of a reset was in each method. We would like to assess if users become increasingly disoriented during long uses of the system. A priori, we can make several observations about the performance of the various methods. First, the backup method requires walking backward in an HMD, an action that is less stable than walking forwards. The 2:1-Turn condition switches users between a "normal" (1:1) rotational gain and a 2:1 rotational gain, which may prove disorienting. The Freeze-Turn system disassociates proprioceptive cues from optical flow, which may also be disorienting.

4.1 Materials

Twelve naive subjects participated in this study. The virtual world was viewed through a full color stereo NVIS nVisor SX Head Mounted Display with 1280 x 1024 resolution per eye, and a field of view of 60° diagonally. The size of the physical room in which the experiments were performed was approximately 5m by 6m, and within the room the limits of the four camera video position tracking system was approximately 5m by 5m. The virtual room was 50m by 50m shown in Figure 2, ten times the size of the physical limits of the tracking system. Objects were placed in the room in different orientations to give the subject a sense of the size and scale of the environment. The environment contained 7 different color tables scattered throughout the environment, 14 posters on the wall, a refrigerator, a fish tank, three sofa areas, two bookshelves, a group of six chairs, a computer desk, a computer, doors, a group of slot machines, and a pool table.

Figure 2: This figure shows the virtual environment used in the experiment.

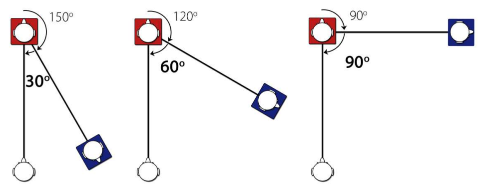

Figure 3: This figure shows a top down view of the three different angles (30°, 60°, and 90°) of the two segment path followed by each participant. The subject walks to a red chair then turns 150°, 120°, or 90°, respectively, and walks to the blue chair.

4.2 Procedure

The goal of this experiment was to assess how well subjects maintain spatial awareness of an environment after undergoing resets. We tested subject's spatial knowledge in each of the three resetting conditions. Since there were six orders of three reset conditions, two subjects completed the each of the six different orders in a counter-balanced fashion. During each testing condition, the participant completed a total of eighteen trials. A trial consisted of walking a path and then turning to face a remembered target object. Before each trial participants were placed in a starting position, and then asked to remember the location of one object or a set of three objects. Trials involving three objects were included so that subjects needed to keep in mind all three objects during the walk to the test position. In this condition, three objects were named at the start of the trial and subjects were told that they would be asked to turn and face any one of the tree after they walked to the test position. Participants were given about sixty seconds to remember the objects locations and freely rotate around from the starting position to view them before traveling the path. Objects were selected so that they did not appear along the participants' path. The correct angle of response from the facing position at the end of the path to the object that the subject was asked to face varied from 30° to 180°.

The travel path consisted of a two segment route, where subjects walked to a red chair and then to a blue chair. The red and blue chairs are meant to serve only as signs, showing the way they should walk to reach the test position. The angle between the starting point,

RIGHTSLINK

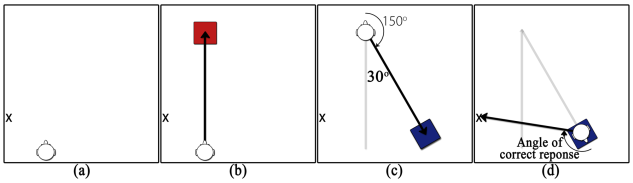

Figure 4: A trial consists of walking a two segment path and turning to face a remembered target object. In this example, the subject is asked to remember the location of an object denoted ‘x’ as shown in (a). Once the subject indicates to the experimenter that they have memorized the position, the red chair appears and the subject is instructed to walk to it as indicated by (b). Once the subject has reached the red chair, it disappears, and the blue chair appears (c). In this example the experimenter instructs the participant to find the blue chair on their right, requiring them to turn 150°. Once the subject reaches the blue chair (d), they are asked to turn to face object ‘x’ with eyes closed. The correct angle of response is shown in (d).

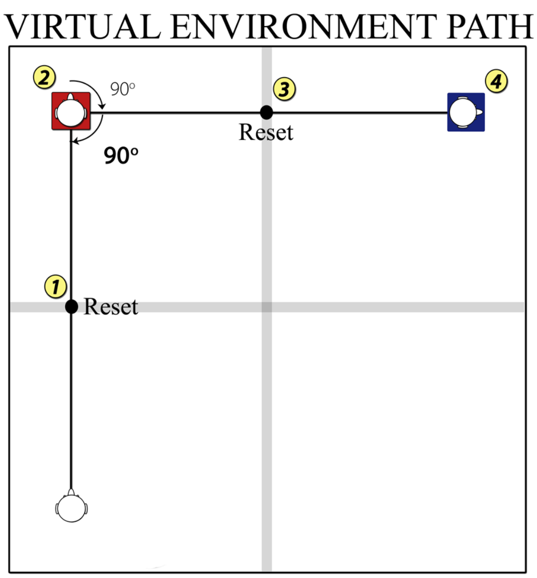

Figure 5: This figure shows a two segment path traveled by a participant in the virtual environment during a two reset trial. In this particular example, the angle of the path is 90°. The subject starts the trial positioned in the bottom left corner. They walk to the red chair, and are reset once at position 1 along their path to the red chair. Once the subject reaches the red chair at position 2, the subject turns 90° to the right to find the blue chair and walks towards it. Along the way to the blue chair the subject is reset at position 3 and then continues to reach the blue chair at position 4. The resets do not change the position and orientation of the user in the virtual environment.

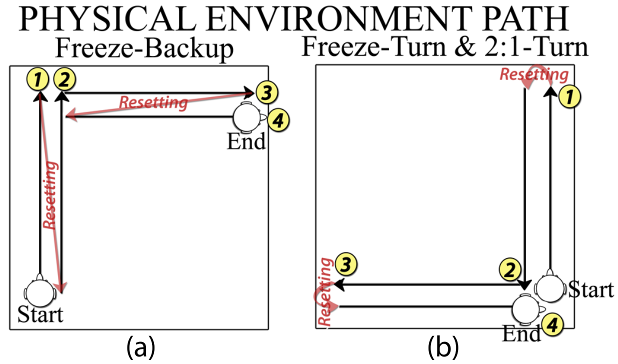

Figure 6: This figure shows the physical path taken by the subjects when traveling the virtual path of Figure 5. In the Freeze-Backup condition shown in (a), participants walk toward the red chair and are reset at position 1. To complete a reset, they take steps backwards as shown by the red arrow. During the reset, their position is frozen in virtual space. After they have taken enough steps backwards, the screen unfreezes and they continue along their path to the red chair at position 2. Next the participant walks to the blue chair, is reset at 3 then continues along until reaching position 4. The physical path followed during this trial in 2:1-Turn condition and the Freeze-Turn condition is the same as shown in (b). The subject is reset at position 1 and turns 180° to continue to the red chair positioned at 2. The subject turns 90° to the right and continues toward the blue chair and is reset at position 3.

red chair, and blue chair was either 30°, 60°, or 90° as shown in Figure 3. Figure 4 shows an example of a trial. After the subject memorizes the location of the object or objects (Figure 4(a)), a red chair appears and the participant is instructed to walk to the red chair (4(b)). Once the subject had arrives at the red chair, the red chair disappeared and a blue chair appeared(4(c)). The experimenter instructed the subject on which direction to turn (right or left) to find the blue chair. The subject was not allowed to look around at the target object or objects while walking the two segment path. At the end of the path, the experimenter instructed the subject to close his eyes and turn to face a remembered target object (4(d)). Time was recorded using a stopwatch and the rotational position was recorded by the computer. The average distance from the final location subject at the blue chair to the target object was approximately 20m and ranged from 3m to 40m. The starting position was varied randomly within 10m of the center of the room and the orientation varied randomly by 90°.

While locomoting along the path, the subject was reset zero times, two times, or four times depending upon the length of the path. In the zero reset condition, the subject completed two 4m segments. Note that the zero reset condition is the same under all three resetting methods; it was included in the experimental design to provide a baseline across trials. In the two and four reset conditions the subject traveled two 8m paths and two 12m paths, respectively. The position of the reset on the path was engineered so that they were spaced an equal distance apart. For example resets in the 8m path of the two reset condition occurred at 4m. Likewise, in the 12m segment length of the four reset condition, resets occurred at the 4m and 8m. Figures 5 and 6 show an example of a trial where the subject is reset twice. Figure 5 represents the path traveled in the virtual environment and Figure 6 shows the paths traveled under the three different resetting condition. Since there were two different numbers of objects to remember, three path angles, and three different numbers of resets, there were eighteen trials per condition representing each possible combination. Each condition took ap-

RIGHTSLINK

proximately 45 minutes to complete and thus were completed on consecutive days.

We varied the path angle by 30°, 60°, and 90° so that subjects did not use strategies based upon the the path angle. The participant had to memorize the location of one object or three objects. We used three objects as well as one object to see how difficult the task becomes when the cognitive burden is higher. In the case of memorizing one object participants could use strategies, but in the case of three objects subjects generally are forced to spatially update along their path.

5 Results

We analyzed the results of our experiment in terms of the errors and latencies in turning to face the targets. Turning error is defined as the difference in angle when turning to face a given target relative to one's actual position in the virtual room. The angle of correct response to the target object from the end of the path varied as repeated within-subject trials from 30° to 180°. Latency was measured from the time the subject was given the object to face until the subject came to rest at a final position. The independent variables in this experiment were reset condition (Freeze-Backup, Freeze-Turn, and 2:1-Turn), number of resets (2 or 4), number of target objects (1 or 3) and angle of turn (30°, 60° or 90°). All independent variables were within-subjects. As noted above, the zero reset condition was identical across all reset conditions, to provide a baseline under an ideal condition. Thus it was not included in the statistical analysis.

Graphs of the mean turning errors and mean latencies collapsed across various factors are shown in Figure 7 through 12. Figures 7 and 8 show the mean turning error and latency as a function of reset condition with the zero reset condition included as a baseline, respectively. Figures 9 and 10 break this information out by number of resets, and Figures 11 and 12 break this information out further by number of objects.

A multivariate repeated measures analysis of variance on the turning error found a main effect of reset condition $F(2,22) = 5.4$, $p < .05$. Participants made fewer errors with Freeze-Backup than with other reset conditions. There were no other main effects or interactions. A repeated measures analysis on latency show a main effect of number of objects to remember, $F(1,11) = 29.9$, $p < .01$. Participants were faster when they had to remember fewer objects. There was a significant interaction of reset condition $\times$ number of objects, $F(2,22) = 9.8$, $p < .05$. Subjects were fastest when they had to remember fewer objects but were generally faster in the 2:1-Turn condition (see Figure 12).

Finally, at the end of completing all three trials, subjects were asked to pick which method they preferred best. Seven subjects preferred the Freeze-Backup method, four preferred the 2:1-Turn method, and one subject preferred the Freeze-Turn method.

6 Discussion

This paper examines methods for exploring large HMD-based virtual environments when the physical space housing the HMD is limited. We studied three methods for resetting a user's location in physical space while hoping they could maintain their spatial orientation in the virtual space. In the Freeze-Backup method, the user obtains more space for virtual exploration by taking steps backwards while frozen in a fixed position in the virtual environment. In the other two methods, Freeze-Turn and 2:1-Turn, users overcome physical obstruction by physically turning 180° and maintaining

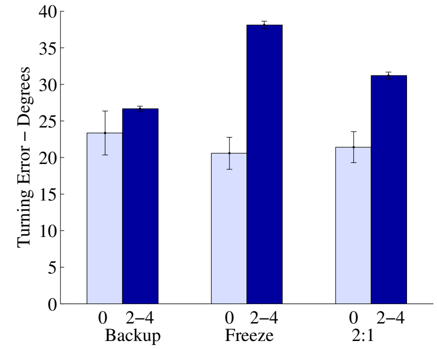

Figure 7: Mean turning error as a function of condition and number of resets. Under each resetting condition the mean of the zero resets is compared to the mean of the two and four resets combined.

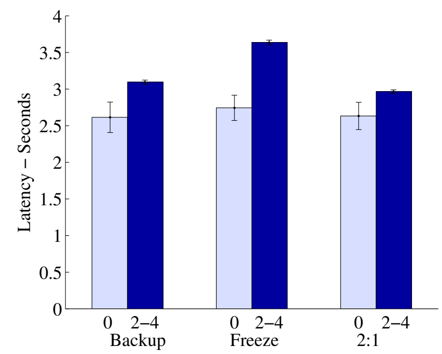

Figure 8: Latency as a function of condition and number of resets. Under each resetting condition the mean of the zero resets is compared to the mean of the two and four resets combined.

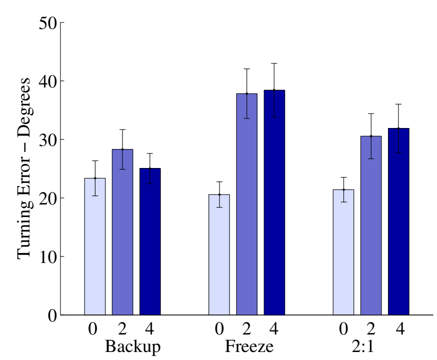

Figure 9: Mean turning error as a function of condition and number of resets. Under each resetting condition the mean is categorized by number of resets: zero, two, or four.

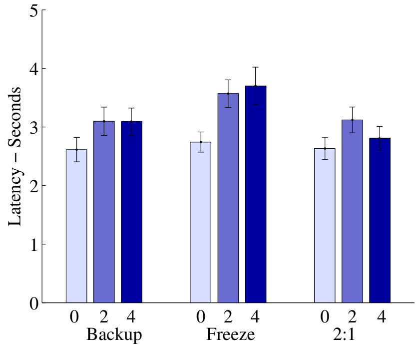

Figure 10: Latency as a function of condition and number of resets. Under each resetting condition the mean is categorized by number of resets: zero, two, or four.

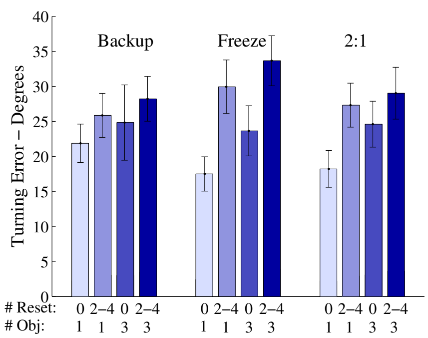

Figure 11: Mean turning error as a function of condition, number of resets, and number of objects memorized. Under each resetting condition the mean is grouped into six different categories representing each of the possible combinations of number of resets and number of objects.

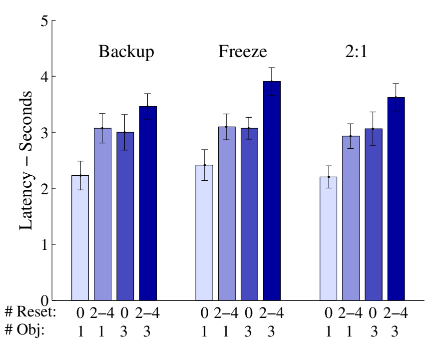

Figure 12: Latency as a function of condition, number of resets, and number of objects memorized. Under each resetting condition the mean is grouped into six different categories representing each of the possible combinations of number of resets and number of objects.

their same position in virtual space before and after the turn. During a Freeze-Turn reset, the orientation of the user is frozen while the subject turns 180°. In the 2:1-Turn condition, the gain of the turn is doubled so that 180° turn in physical space corresponds to 360° turn in virtual space.

Our results indicated that the lowest errors occur in the Freeze-backup condition, while latencies were lowest for the 2:1 condition. There are several interesting observations about these results. Updating one's position in the Freeze-Backup condition involves ignoring proprioceptive cues that would result in the change of perspective being a geometric translation. The lower turning errors are generally consistent with prior literature indicating that it is easier to judge changes in perspective when the geometry of the change is a translation rather than a rotation [Rieser 1989; Presson and Montello 1994; Philbeck et al. 2001]. This result has generally extended to both errors and latencies, so the finding that the 2:1-Turn condition is at worst as good as the Freeze-Backup condition is surprising. The errors and the latencies were the worst for the Freeze-Turn condition. The other two conditions were both better than the Freeze-Turn condition.

The design of our experiments was focused on trial latencies, and we did not keep track of the total time to complete the testing phase of experiment in each of the resetting conditions except for two subjects. The mean testing time of the two subjects for the Freeze-Backup, Freeze-Turn, and 2:1-Turn was 32.2 minutes, 23.2 minutes, and 22.1 minutes, respectively. These results are consistent with our memories of the testing times for the other subjects.

In verbal reports seven of twelve subjects indicated they preferred the Freeze-Backup condition with four of twelve indicating they preferred the 2:1-Turn method. However, even though the Freeze-Backup condition has the lowest turning errors and is preferred by a majority of users, a satisfying design for a commodity-level HMD system would likely consist of either the Freeze-Turn or 2:1-Turn methods. The disadvantages of the Freeze-Backup method are the potential danger of backing into a wall or tripping over the tether, and the longer length of time and walking involved in resetting. However, since the Freeze-Backup condition is the best resetting condition, we envision its use in training applications where spatial orientation is important, where time is not an issue, and trainees are able to have a guide to make sure that they do not back into a wall or get stuck in a corner. For example, this method could be used to test emergency exits for proposed architectural designs in case of a fire, and other emergency type applications. We generally prefer the 2:1-Turn method since resets are relatively fast, and gain could be further exploited so that the amount a user turns during a reset could easily be manipulated to provide maximum physical space for forward exploration.

Users had to remember the locations of one or three objects this experiment, and remembering the higher number of objects resulted in longer trial latencies. There was also a significant interaction between reset condition and number of objects. The 2:1-Turn condition was generally the fastest. This finding indicates that there is little cognitive cost to the rotation that the 2:1-Turn condition introduces when compared to the purely translational actions of the Freeze-Backup condition. This result is interesting when compared to much prior work, e.g. [Rieser 1989], that shows maintaining spatial orientation under general rotations is computationally more expensive than under translations.

The results of these experiments also show that there is a cost to resetting in terms of a user's orientation to remembered objects. To mollify this cost, a complete system would likely involve a method or provision for the users to reorient themselves periodically in the virtual environment. Given the current state of virtual environ-

RIGHTSLINK

ments, the need for reorientation is not a severe drawback, although it is one we would like to eliminate. There is ample evidence that people have difficulty maintaining orientation in virtual environments [Ruddle 2001; Allen and Singer 1997; Péruch et al. 2000]. Typically, these difficulties are attributed to poor idiothetic cues, such as the absence of proprioception and other sources of information provided by self locomotion (in the case of desktop virtual environments) and the limited field of view of HMDs. This topic is one we are actively pursuing.

Future work also involves integrating these resetting methods into a system that incorporates scaling of translational gain as a method for exploring virtual environments [Williams et al. 2006]. We would like to deploy immersive virtual environments widely, for learning and training, and it seems likely that physical space is a constraint that must be overcome for their widespread adoption.

Acknowledgments. The authors thank Bernhard Riecke and Jon Kelly for helpful discussions, and the anonymous reviewers for insightful comments. This material is based upon work supported by the National Science Foundation under Grant IIS-0237621, by a Vanderbilt Discovery Grant, and by the National Institute of Mental Health under Grant 2-R01-MH57868. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

References

ALLEN, R. C., AND SINGER, M. J. 1997. Landmark direction and distance estimation in large scale. In *Proceedings of the Human Factors and Ergonomics Society 41st Annual Meeting*, 1213–1217.

BOWMAN, D. A., DAVIS, E. T., HODGES, L. F., AND BADRE, A. N. 1999. Maintaining spatial orientation during travel in an immersive virtual environment. *Presence 8*, 6, 618–631.

CHANCE, S. S., GAUNET, F., BEALL, A. C., AND LOOMIS, J. M. 1998. Locomotion mode affects updating of objects encountered during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence 7*, 2, 168–178.

DARKEN, R., AND SIBERT, J. 1996. Navigating in large virtual worlds. *The Int. J. of Human-Computer Interaction 8*, 1, 49–72.

KEARNS, M., WARREN, W., DUCHON, A., AND TARR, M. 2002. Path integration from optic flow and body senses in a homing task. *Perception 31*, 349–374.

KLATZKY, R. L., LOOMIS, J. M., BEALL, A. C., CHANCE, S. S., AND GOLLEDGE, R. G. 1998. Spatial updating of self-position and orientation during real, imagined, and virtual locomotion. *Psychological Science 9*, 4 (July), 293–298.

KUHL, S. A. 2004. Recalibration of rotational locomotion in immersive virtual environments. In *APGV 2004*, 23–26.

LATHROP, W. B., AND KAISER, M. K. 2002. Perceived orientation in physical and virtual environments: Changes in perceived orientation as a function of idiothetic information available. *Presence 11*, 1, 19–32.

NITZSCHE, N., HANEBECK, U., AND SCHMIDT, G. 2004. Motion compression for telepresent walking in large target environments. *Presence 13*, 1, 44–60.

PÉRUCH, P., BELINGARD, L., AND THINUS-BLANC, C. 2000. Transfer of spatial knowledge from virtual to real environments. In *Spatial Cognition II*, C. Freska, Ed. Springer, Berlin.

PHILBECK, J. W., KLATZKY, R. K., BEHRMANN, M., LOOMIS, J. M., AND GOODRIDGE, J. 2001. Active control of locomotion facilitates nonvisual navigation. *Journal of Experimental Psychology: Human Perception and Performance 27*, 141–153.

PICK, H. L., RIESER, J. J., WAGNER, D., AND GARING, A. E. 1999. The recalibration of rotational locomotion. *J. Exp. Psych: Hum. Perc. Perf. 25*, 5, 1179–1188.

PRESSON, C. C., AND MONTELLO, D. R. 1994. Updating after rotational and translational body movements: coordinate structure of perspective space. *Perception 23*, 1447–1455.

RAZZAQUE, S., KOHN, Z., AND WHITTON, M. C. 2001. Redirected walking. *Eurographics Short Presentation*.

RIESER, J. J., PICK, H. L., ASHMEAD, D. A., AND GARING, A. E. 1995. The calibration of human locomotion and models of perceptual-motor organization. *J. Exp. Psych: Hum. Perc. Perf. 21*, 480–497.

RIESER, J. J. 1989. Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology 15*, 6, 1157–1165.

RUDDLE, R. A., AND LESSELS, S. 2006. For efficient navigation search, humans require full physical movement but not a rich visual scene. *Psych. Science*. To appear.

RUDDLE, R. A., PAYNE, S. J., AND JONES, D. M. 1999. Navigating large-scale virtual environments: What differences occur between helmet-mounted and desk-top displays? *Presence 8*, 2, 157–168.

RUDDLE, R. A. 2001. Navigation: Am i really lost or virtually there? *Engineering Psychology and Cognitive Ergonomics 6*, 135–142.

SLATER, M., USOH, M., AND STEED, A. 1995. Taking steps: The influence of a walking technique on presence in virtual reality. *ACM Trans. on Human Interaction* 2, 3, 201–219.

TEMPLEMAN, J. N., DENBROOK, P. S., AND SIBERT, L. E. 1999. Virtual locomotion: Walking in place through virtual environments. *Presence 8*, 6, 598–617.

USOH, M., ARTHUR, K., WHITTON, M. C., BASTOS, R., STEED, A., SLATER, M., AND BROOKS, F. P. 1999. Walking > walking-in-place > flying, in virtual environments. In *SIGGRAPH 99*, 359–364.

WALLER, D., HUNT, E., AND KNAPP, D. 1998. The transfer of spatial knowledge in virtual environment training. *Presence 7*, 2, 129–143.

WILLIAMS, B., NARASIMHAM, G., MCNAMARA, T. P., CARR, T. H., RIESER, J. J., AND BODENHEIMER, B. 2006. Updating orientation in large virtual environments using scaled translational gain. APGV.

A Algorithm Details

The specifics of each resetting algorithm is discussed in this appendix. The physical position of a user in x, y, z space using a right-handed coordinate system is obtained a camera tracking system. The position in the center of the room on the floor is (0, 0, 0). The x, y, and z directions while standing in the center of the room facing to the front correspond to front-to-back movement, user height, and right-to-left movement, respectively. Movement is limited in the x and z directions due to the finite range of the tracking system. Since the y-direction indicates movement perpendicular to the

RIGHTSLINK

ground pane, this value typically represents the user's eyeheight, and does not limit the exploration of the virtual environment. Orientation is obtained from the rotational sensor located on HMD which updates rotation about the x-axis (pitch), y-axis (yaw), and z-axis (roll).

A.1 Freeze-Backup

The algorithm first initializes the reset offset *resetOffset* so that

$$resetOffset_x = 0 \quad (1)$$
$$resetOffset_z = 0 \quad (2)$$

After a reset, the position of the user in the virtual space must be calculated by offsetting the physical position of the user by some amount. Therefore, the virtual position (*vePos*) and orientation (*veOri*) at any point in time while the user is not undergoing a reset can be calculated as

$$vePos_x = currPos_x + resetOffset_x \quad (3)$$
$$vePos_y = currPos_y \quad (4)$$
$$vePos_z = currPos_z + resetOffset_z \quad (5)$$
$$veOri(x,y,z) = currOri \quad (6)$$

where vectors *currPos* and *currOri* indicate the current position and orientation of the user. Before the first reset, the users position in x, y, z space and orientation (pitch, yaw, and roll) is equal to that of the physical space. Once the user has reached a boundary, a message automatically appears requesting that the user stop walking. Once the user stops moving, their position in the virtual position is fixed and their current location in the virtual environment must be recorded and the reset offset updated accordingly.

$$resetOffset_x = currPos_x + resetOffset_x \quad (7)$$
$$resetOffset_z = currPos_z + resetOffset_z \quad (8)$$

During the reset, the user takes steps backwards, yet the virtual position is not updated. However, virtual orientation is updated, enabling the user to look around from a fixed position while backing up. Therefore while the user is undergoing a reset, the user's position in the virtual world must be calculated:

$$vePos_x = resetOffset_x \quad (9)$$
$$vePos_y = currPos_y \quad (10)$$
$$vePos_z = resetOffset_z \quad (11)$$
$$veOri(x,y,z) = currOri \quad (12)$$

The user stops backing up when the experimenter indicates that they have backed up enough. Then to complete the reset and enable to continue along the path he or she was traveling prior to the reset the following calculation is made:

$$resetOffset_x = resetOffset_x - currPos_x \quad (13)$$
$$resetOffset_z = resetOffset_z - currPos_z \quad (14)$$

A.2 Freeze-Turn

In this resetting condition, the user turns around with the virtual screen frozen until he or she feels that they have turn approximately 180°, and then the screen is unfrozen and the user continues along their path. Thus, head movement about the y-axis must be manipulated. This manipulation is controlled by θ_y, and is initialized with the reset offset in the x ($resetOffset_x$) and z ($resetOffset_z$) directions. The two variables $rotAxis_x$ and $rotAxis_z$ specify the origin of the transformation. Thus, the variables in this resetting conditions are initialized by equations 1, 2, and the following:

$$\theta_y = 0 \quad (15)$$
$$rotAxis_x = 0 \quad (16)$$
$$rotAxis_z = 0 \quad (17)$$

To calculate the users' position in the virtual environment while they are not resetting, the current physical location of the user in the x and z directions must be translated by $resetOffset_x$ and $resetOffset_z$ and rotated about the y-axis. Thus, the rotation matrix is defined as

$$R = \begin{bmatrix} cos(\theta_y) & sin(\theta_y) \\ -sin(\theta_y) & cos(\theta_y) \end{bmatrix} \quad (18)$$

Current virtual position and orientation is calculated as follows:

$$\begin{bmatrix} vePos_x \\ vePos_z \end{bmatrix} = \begin{bmatrix} currPos_x - rotAxis_x \\ currPos_z - rotAxis_z \end{bmatrix} R + \begin{bmatrix} resetOffset_x \\ resetOffset_z \end{bmatrix} \quad (19)$$
$$vePos_y = currPos_y \quad (20)$$
$$veOri_x = currOri_x \quad (21)$$
$$veOri_y = currOri_y + \theta_y \quad (22)$$
$$veOri_z = currOri_z \quad (23)$$

When the tracker senses the user out of bounds, the computer alerts the user by message on the HMD display instructing them to stop locomoting. To reset, the user turns around while frozen in their current position. Therefore, to start the reset the following calculations are made:

$$startAngle_y = currOri_y \quad (24)$$
$$\begin{bmatrix} resetOffset_x \\ resetOffset_z \end{bmatrix} = \begin{bmatrix} currPos_x - rotAxis_x \\ currPos_z - rotAxis_z \end{bmatrix} R + \begin{bmatrix} resetOffset_x \\ resetOffset_z \end{bmatrix} \quad (25)$$
$$rotAxis_x = currPos_x \quad (26)$$
$$rotAxis_z = currPos_z \quad (27)$$

The variable $startAngle_y$ stores the y-direction orientation of the user at reset. Thus, during a reset, virtual position is calculated using equations 9, 10, and 11, and orientation is calculated as:

$$veOri_x = currOri_x \quad (28)$$
$$veOri_y = startAngle_y + \theta_y \quad (29)$$
$$veOri_z = currOri_z \quad (30)$$

Position is frozen in the x and z directions and orientation is frozen about the y axis. To end the reset θ_y must be updated:

$$\theta_y = \theta_y - (currOri_y - startAngle_y) \quad (31)$$

A.3 2:1-Turn

The algorithm for this resetting condition is exactly the same as the Freeze-Turn condition, with the exception of equation 29 which calculates virtual orientation around the y-axis during the resetting phase:

$$veOri_y = (currOri_y - startAngle_y) * 2 + \theta_y \quad (32)$$

RIGHTSLINK