

Physically Large Displays Improve Path Integration in 3D Virtual Navigation Tasks

Desney S. Tan, Darren Gergle, Peter G. Scupelli, Randy Pausch

School of Computer Science

Carnegie Mellon University

5000 Forbes Avenue, Pittsburgh, PA 15213, USA

{desney, dgergle, pgs}@cs.cmu.edu, pausch@cmu.edu

ABSTRACT

Previous results have shown that users perform better on spatial orientation tasks involving static 2D scenes when working on physically large displays as compared to small ones. This was found to be true even when the displays presented the same images at equivalent visual angles. Further investigation has suggested that large displays may provide a greater sense of presence, which biases users into adopting more efficient strategies to perform tasks. In this work, we extend those findings, demonstrating that users are more effective at performing 3D virtual navigation tasks on large displays. We also show that even though interacting with the environment affects performance, effects induced by interactivity are independent of those induced by physical display size. Together, these findings allow us to derive guidelines for the design and presentation of interactive 3D environments on physically large displays.

Categories and Subject Descriptors: H.5.1 [Multimedia Information Systems]: Artificial, augmented, and virtual realities; H.5.2 [Information Interfaces and Presentation]: User Interfaces - Screen design, User-centered design, Graphical user interfaces; J.4 [Social and Behavioral Sciences]: Psychology.

General Terms: Human Factors, Performance.

Keywords: Physically large display, field of view, visual angle, 3D virtual navigation, path integration, immersion, presence.

INTRODUCTION

With recent advances in technology, large wall-sized displays are becoming prevalent. Although many researchers have articulated qualitative benefits of group work on large displays, much less has been done to systematically quantify and exploit these benefits for individual users.

Recently, we have begun to isolate how physical display size affects the way individual users perform different tasks [26]. We found that users were more effective performing spatial orientation tasks when they worked on large displays as compared to smaller ones, even when identical images were

viewed at equivalent visual angles. However, because of the static 2D task we used, we could not generalize our observations to more real-world or interactive tasks. In this work, we report a study we conducted to explore how physical display size affects cognitive strategies and performance on an interactive 3D navigation task.

This work contributes a deeper understanding of the effects of display size by demonstrating that prior results do indeed generalize, and that users are more effective at performing 3D navigation tasks on physically large displays. Additionally, we show that interactivity significantly affects performance on this set of tasks, but that the effects caused by interactivity are independent of those caused by physical display size. Based on these results, we provide design guidelines for the presentation of interactive 3D environments on physically large displays.

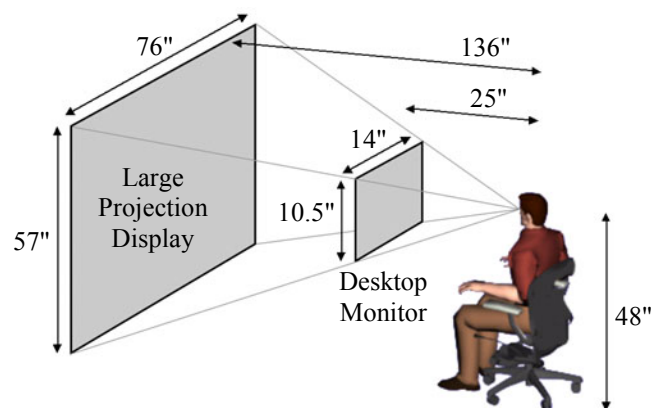


Figure 1. In our work, we maintained a constant visual angle for each of the two displays. Only size and distance to the user changed.

RELATED WORK

Physically Large Displays

Many large displays are created by combining multiple displays into a single display system. As such, a sizable proportion of large display research has focused on using the additional screen space offered by such systems to create more productive computing environments. For example, Baudisch et al. [2] provided a large low-resolution overview of the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2004, April 24–29, 2004, Vienna, Austria.

Copyright 2004 ACM 1-58113-702-8/04/0004...\$5.00.

working context around a smaller high-resolution focal screen. MacIntyre et al. [18] assisted users in managing multiple working contexts by presenting montages of images on large peripheral displays. Tan et al. [27] utilized peripheral projection displays to show pictures that served as memory cues. They hypothesized that the greater the sense of presence invoked by the display, the better the memory for learned information. Slater et al. [25] defined this sense of presence as “a state of consciousness, the (psychological) sense of being in the virtual environment.” When users are present in the virtual environment (VE), the location of their physical bodies is often construed as being contained within the space rather than looking at it from the outside. They claimed that being in this state makes users most effective in virtual environments. While it has been shown that the more inclusive, extensive, surrounding, and vivid the display, the higher the potential of presence [5], little work has carefully explored how display size affects the sense of presence.

There are two major factors to consider when thinking about display size: field of view, which is also referred to as visual angle, and physical display size. Grabe et al. [13] review several studies aimed at understanding the interaction of screen content with various display technologies. However, because little of this work differentiated between field of view and physical display size, we cannot decisively attribute the findings to one factor or the other. In fact, in most current work, researchers have placed large displays at distances that are not proportional to their increase in size over smaller displays. This provides different fields of view between the displays and leads to the confounding of field of view and physical display size.

Most of the work that has differentiated between these two factors focuses on field of view as the feature of interest. It is generally agreed that wider fields of view can increase the sense of presence in VEs [1, 22]. Czerwinski et al. [8] review prior literature suggesting that restricting the field of view leads to negative impacts on perceptual, visual, and motor performance in various tasks, possibly because users find it difficult to transfer real world experience into the VE. They also report evidence that wider fields of view offered by large displays lead to improved performance in 3D navigation tasks, especially for females.

It is only recently that researchers have begun to examine physical display size as a factor of interest. In exploring affordable alternatives to head-tracked head-mounted displays, Patrick et al. [19] found that users performed significantly better at remembering maps when using a large projection display as compared to a standard desktop monitor. They attributed part of this effect to a higher level of presence afforded by the large projection display, which may have provided better cues for map formation.

In order to better understand these results, we constructed an experimental setup in which a small and large display presented the same images at equivalent visual angles [26]. We found performance increases on spatial orientation tasks in-

volving static 2D scenes when users worked on the large display. We attributed these effects to large displays providing a greater sense of presence, which biased users into adopting egocentric strategies. Using these egocentric strategies, users imagined rotating their bodies within the virtual environment. This was much more efficient than using corresponding exocentric strategies, in which users imagined the environment rotating around their bodies. Just et al. [14] provide a comprehensive review suggesting that this choice of mental coordinate systems partially accounts for individual differences in spatial ability. Bell et al. [3] showed that people chose different coordinate systems for physically large spaces, as opposed to smaller ones. It is no surprise then, that the physical size at which information is presented biases the choice of mental representations, and hence affects task performance.

In this work, we extend the external validity of our previous work by exploring the benefits of physically large displays on 3D navigation tasks. 3D navigation is a logical extension to current tasks for several reasons: (1) Since VEs are only as effective as the user’s ability to move around and interact with information, improving performance on such tasks would greatly enhance overall productivity within these environments; (2) 3D navigation is a spatial task which seems well-suited to benefit from the greater sense of presence and egocentric strategies induced by large displays.

3D Navigation by Path Integration

When navigating, users continually update mental representations of their position and orientation within the environment. This is called spatial updating. Two strategies users employ to perform spatial updating are *piloting*, using external landmarks to get their absolute position within the environment, and *path integration*, sensing self-velocity and acceleration to derive their position relative to some starting point [12]. Path integration allows travelers to integrate isolated views of the environment into a cognitive map which may be used for subsequent piloting. Initial work [6, 16, 17] suggested that successful path integration requires proprioceptive and vestibular cues, or the physical awareness of our body’s position with respect to itself or to the environment. However, recent studies [eg, 15] have demonstrated otherwise, showing path integration to be effective using only visual cues. Interestingly, many of these studies have presented the environments on either physically large or wide field of view displays.

Riecke et al. [23] used a large half-cylindrical 180 degree wide projection screen and demonstrated that visual path integration, without associated proprioceptive or vestibular cues, is sufficient for elementary homing tasks. They claimed that additional peripheral cues provided by the display aided task performance. In other work, Péruch et al. [20] used a large video-projector screen and found that users navigated equally well in various field of view conditions, suggesting that task performance was independent of field of view. However, they did not explicitly discuss the influence that



Figure 2. The joystick used (left); User working on the small display (middle) and the large display (right).

the physically large display had in their studies. Our work contributes to this growing body of research, demonstrating that physical display size influences performance on these tasks.

In separate VE navigation work, some researchers have found that the acquisition of spatial knowledge is facilitated by active navigation control [eg, 7, 21]. These researchers claim that proprioceptive cues provided by the input devices as well as cognitive benefits of decision-making immerse users more within the VEs and aid in encoding mental representations of the environments. Others however, have reported opposite results, showing that active control hurts performance in various navigation tasks [eg, 4]. Flach [10] argues that the different results could be due to the tradeoffs imposed by control of attention, kinds of information available, sensitivity to information, as well as activities involved. We decided to explore both how level of interactivity in the VE affects navigation by path integration, as well as how it interacts with effects caused by varying the physical size of displays.

PATH INTEGRATION STUDY

In our study, we examined physically large displays as an independent factor. We chose a 3D navigation task to address the questions of external validity and real world usefulness of prior results. Additionally, we explored the egocentric vs. exocentric hypothesis that has been proposed to explain effects. Interestingly, there have been two mental models suggested in connection with performing path integration, a traveler-centered model and an environment-centered model. These models relate directly to the proposed dichotomy of possible strategies, differentiating egocentric from exocentric representations (see Rieser [24] for a review). If large displays provide a greater sense of presence and bias users into adopting egocentric strategies, we would expect performance to increase on our 3D navigation task. Thus,

Hypothesis 1: Users will perform better in the path integration task when using a physically large display due to the increased likelihood that they adopt egocentric strategies.

We also examined the effectiveness of providing active navigation control. While prior literature provides evidence of

active control helping in some situations and hurting in others, we expected users to perform better when they had interactive control using the joystick due to the additional cues afforded by the physical manipulation. Therefore,

Hypothesis 2: Users will perform better in the path integration task when they are interactively moving themselves through the virtual environment.

Finally, we expected the benefits of the large display to be robust against other factors that could potentially provide a similar heightened sense of presence. Specifically,

Hypothesis 3: The effects induced by physical display size will be independent of those induced by interactivity.

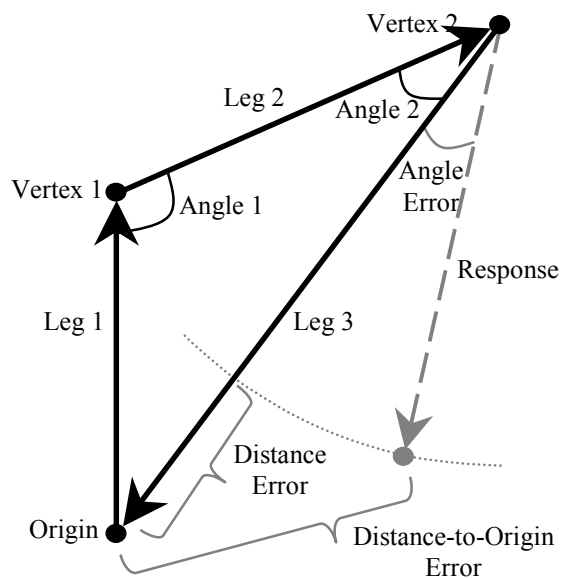
Participants

Sixteen (8 female) college students, aged 19 to 29 years old, participated in the study. Users were intermediate to experienced computer users who played an average of less than an hour of 3D video games per week. All users had normal or corrected-to-normal eyesight. The study took about an hour and users were paid for their participation.

Equipment

We closely replicated the hardware setup used in our previous experiments [26] (see Figures 1 and 2). We used two displays, an Eiki Powerhouse One LCD projector and an 18" NEC MultiSync 1810X LCD monitor. Both displays ran at a resolution of 1024×768 , and were calibrated to be of roughly equivalent brightness and contrast. We mounted the projector from the ceiling and projected onto a flat, white wall. The image projected on the wall was 76" wide by 57" tall. The image on the monitor was 14" wide by 10.5" tall. To keep the visual angles equivalent between the displays, we placed the monitor 25" and the projector 136" away from the user. The center points of both displays were set to eye-height, 48" above the ground. We moved the monitor in and out as necessary.

We had initially implemented a fairly complex mechanism to ensure that the user's eyes were constantly placed where the two displays would occupy equivalent visual angles. However, pilot-test video showed that users hardly moved their heads during the study. In fact, adjusting the seat before the



study and telling users not to further adjust it was sufficient to maintain similar visual angles between the displays. We used this latter method in our setup. In the rare case where users moved their chair during the study, we readjusted their position before proceeding.

We ran the study on a 1.33 GHz computer with a GeForce4 MX graphics card. Our VEs updated at 60 frames per second. We used a switchbox to send the graphics output to only one of the displays at any given time. The user provided input with the control stick and trigger button on a Radioshack 26-444 Joystick (see Figure 2).

Task

We used a triangle completion task to test how each of our manipulations affected path integration. In this task, we led users along two legs of a triangle and then had them find their way back to their starting position unaided. We picked this task because it is simple, well defined, and ecologically inspired. It is also commonly regarded as the most direct way of measuring path integration ability [12]. We believe that our results extend to more complex navigation tasks.

To isolate effects, we created a virtual environment that provided optical flow and depth cues necessary for path integration, but that did not contain distinct landmarks used for piloting. The environment was a circular arena with two concentric circles of trees. The inner circle bounded the navigation area. It was 16 meters wide and contained ten 4 meter tall trees that were evenly spaced along the circle. The outer circle was 22 meters wide, and contained ten 5 meter tall trees that were darker in color than the trees in the inner circle. Users in pilot tests complained that the environment seemed static and unreal. To address this concern, we animated the trees to gently sway in the breeze. The ground had a uniformly speckled texture. The maximum speed a user

could move was 2 meters per second. The maximum turning speed was 30 degrees per second.

Each trial in the test consisted of two phases, the encoding phase and the return-to-origin phase. In the encoding phase, we led users along the first two legs of a triangle (see Figure 3). For each leg, they saw a pole at the next vertex of the triangle (see Figure 4). Their task was to turn and move to each successive pole. Users could only turn when they were standing at a vertex. Additionally, they could only move forward and backward in straight lines, and only while they were facing the next vertex. This prevented users from straying off the defined path. Upon getting to the last vertex, users began the return-to-origin phase. In this phase, the poles disappeared and users had to use the joystick to turn and move to the origin, using only the mental map they had constructed and the visual cues provided by the environment. Again, users could only turn when they were at the vertex. However, since they could move forward and backward, they could return to the vertex to adjust their response angle if they felt that it was not correct. They pressed the trigger on the joystick to indicate their answer when they were done navigating.

Procedure and Design

Prior to the start of the study, users performed the Map Planning (SS-3) subtest from the ETS Kit of Factor-Referenced Cognitive Tests [9]. This well-validated test is commonly used to evaluate spatial ability skills.

After reading through instructions, users performed a set of practice trials before beginning the actual test. In these practice trials, users saw an overview map of the triangle before performing the task. After each trial, they received feedback on the overview map showing where they ended up relative to the origin. Each of the six practice trials used a unique triangle that did not match any of the test triangles. To prevent users from becoming reliant on maps, they were warned that they would not have these maps during the test.

The study was a 2 (Display Size: small vs. large) \times 2 (Interactivity: passive viewing vs. active joystick) within-subjects




Figure 4. First person view of arena and pole that the user experienced while performing the tasks.

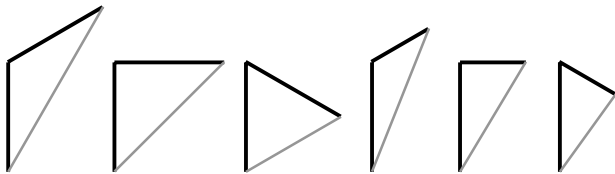


Figure 5. Triangles tested in the main study.

design. Users performed six trials in each condition, corresponding to six triangle configurations created by permuting three Angle 1 values (60, 90 and 120 degrees) and two Leg 2 lengths (3 and 5 meters). Leg 1 was always 5 meters long. These triangles can be seen in Figure 5. Each triangle was centered in the arena.

In the passive viewing condition, users had no control of their movement in the encoding phase. Instead, they passively viewed themselves moving along the first two legs of the triangle. We used a slow-in slow-out animation with linear acceleration to move the user at the maximum speeds. In the active joystick condition, users used the joystick to navigate the first two legs. In both conditions, users had joystick control in the return-to-origin phase. We balanced the order of the Display Size and Interactivity manipulations separately and fully randomized the order in which we presented different triangles in each condition.

Dependent measures included: (a) the overall *distance-to-origin error*, the absolute straight line distance between the point to which the user navigated and the actual origin; (b) the *angle-turned error*, the signed difference between the correct angle (Angle 2) and the angle the user turned; and (c) the *distance-moved error*, the signed difference between the correct distance (Leg 3) and the distance the user moved. These error measures can be seen in Figure 3.

Results

Overall Performance

In our primary analysis, we examined the distance-to-origin error as the variable of interest. We used a mixed model analysis of variance (ANOVA) in which Display Size (small vs. large) and Interactivity (passive viewing vs. active joystick) were repeated and Gender was treated as a between-subjects factor. We included all 2-way and 3-way interactions in the analysis. Because each participant performed each condition, observations within a pair were not independent and we modeled participants as a random effect nested within Gender. We originally included two covariates in the model: Time spent in the encoding phase, and the Spatial Abilities score. However, we removed these from the final analyses because they were not significant. The estimates and significance levels of the main factors of interest did not change in any significant fashion and the overall model fit was improved.

We found a significant main effect of Display Size ($F(1,339)=11.24$, $p<.001$), with the large display resulting in users having shorter error distances (2.88 vs. 3.48 meters). We also observed a significant main effect of Interactivity

($F(1,339)=12.38$, $p<.001$), with trials in the passive viewing condition demonstrating shorter error distances than trials in the active control condition (2.87 vs. 3.49 meters). See these results in Figure 6. We saw no interaction between Display Size and Interactivity, suggesting that the manipulations were independent of one another.

Examination of the effect of Gender on performance did not reveal a significant difference between males and females (3.12 vs. 3.24 meters, $F(1,14)=.07$, $p=.79$). Prior literature suggests differential effects of gender on performance with different fields of view. However, we controlled field of view to be constant across displays and saw no interaction between Gender and Display Size. This is consistent with our prior findings [26], in which we used a similar setup. All remaining interactions were not significant.

Systematic Component Errors

To test for systematic performance errors, we decomposed the aggregate distance-to-origin error and individually examined the distance-moved error and the angle-turned error. We used the same model as in the primary analysis, but replaced the dependent variable distance-to-origin error with the distance-moved and angle-turned errors.

We found a significant difference in Display Size for the distance-moved error ($F(1,339)=4.314$, $p=.03$). Users consistently underestimated the distance in both conditions (mean of 1.17 meter undershoot, overall). However, they underestimated significantly more in trials with the smaller display than the large (1.03 vs. 1.31 meter undershoots). The effect of Interactivity, while trending in the expected direction, was not significantly different for this measure (1.10 vs. 1.24 meter undershoot for passive viewing vs. active control, $p=.28$). While the mean result across all conditions demonstrated an underestimation of the angle (1.43 degree underturn, on average), we found no significant differences across the conditions for angle-turned error.

Effects of Triangle Shape

To examine effects of the different triangle configurations,

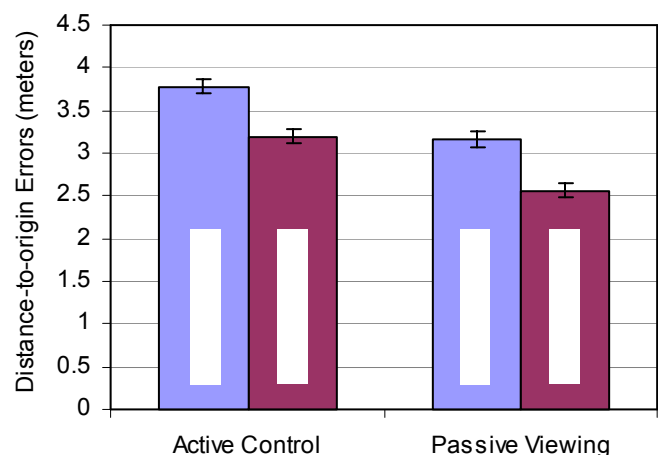


Figure 6. Main effects of Interactivity and Display Size. We saw no interactions between manipulations.

we performed an additional analysis to explore whether the correct distance and correct angles affected performance in any systematic way. We performed a similar analysis as in the previous sections with correct distance (Leg 3) and correct angle (Angle 2) added as independent variables. We also examined interactions to determine if the Display Size and Interactivity manipulations were more or less helpful depending on the difficulty of triangles.

We found, holding all other variables constant, that for every meter the correct distance increased, users accumulated an additional 0.635 meters in the distance-to-origin error ($F(1,354)=12.70$, $p<.001$). An examination of the interactions revealed that correct distance did not differentially affect performance across the various conditions. Correct angle had little effect on overall performance ($F(1,354)=1.47$, n.s.) and had no significant interactions.

In a similar fashion to the breakdown we performed with the systematic component errors, we also looked at the effect of correct distance on the distance-moved error as well as the effect of correct angle on the angle-turned error. We found that for each meter the correct distance increased, users underestimated the distance by an additional 0.465 meters ($F(1,358)=137.90$, $p<.001$). Similarly, we found that for each degree the correct angle increased, users further undershot the actual angle by an additional 0.501 degrees. We found no differential effects across conditions. These results are consistent with previous research showing that triangle shape significantly affects error rates [15].

Summary

In this study, we found evidence suggesting that users did indeed perform better on the path integration task when working on the large display, supporting hypothesis 1. Contrary to our initial hypothesis 2, we found that actively controlling navigation with a joystick did not help, but instead hurt, performance on the task. Finally, we found that effects induced by interactivity were independent of those induced by display size, supporting hypothesis 3.

FOLLOW-UP INVESTIGATION

Our implementation of the triangle completion task in the first study contained two fairly distinct subtasks: wayfinding, which included sensing the outbound path, forming a mental representation of the environment, and then computing the return path; and locomotion, or actually executing that path with motor movements to control the joystick. Since the errors observed could have been a result of either of these sub-processes, we ran a follow-up investigation to test the contribution that locomotion had on the error. Specifically, we wanted to know how well a user could use the joystick to turn a specified angle and move a specified distance. We hypothesized that users were proficient with the joystick and that the errors observed in the main study could be attributed mainly to wayfinding errors.

Participants

Eight (4 female) college students, who had not participated in the main study, participated in the follow-up. We selected users to be of approximately the same demographic as before. The follow-up took about half an hour and users were paid for their participation.

Materials and Procedure

We used the same physical setup as before. We simplified the triangle completion task to reduce the wayfinding component and test only how accurate users were in using the joystick to turn and move specified angles and distances.

Before each trial, we provided users with the angle and distance that they would have to turn and move. We told users the angle that they would have to turn (eg, 60 degrees to the right). Unfortunately, since the virtual world contained no absolute unit of distance, telling a user to move 3 meters, for example, would not have been very useful. Hence, we specified the distance a user had to move by having them first travel a path of identical distance.

To reinforce these specifications, we showed users an overview map containing two legs of equal length connected at a single vertex. The user's task when placed in the virtual environment was to move straight ahead along the first leg, learning the distance they would have to travel. Following this, they had to turn the specified angle and move a distance equal to that of the first leg in order to reach the end-point of the second leg. They hit the trigger on the joystick to indicate when they were done navigating.

Before the study, users read the instructions, tried six practice trials in which they received feedback, and then performed the test. We tested angles and distances that represented the range performed in the return-to-origin phase of the main study. Using angles of 60, 90, 120, or 150 degrees, and distances of 3, 4, 5, 6, 7, or 8 meters, we created twenty-four test trials. Users performed these trials only in the small display \times active joystick condition. We expected that the largest locomotion errors would occur in this condition since it was the one in which users made the largest overall errors in the main study. This would serve as a good estimate of how much the locomotion errors were contributing to the overall error.

Results

We calculated 95% confidence intervals for our dependent measures. The distance-to-origin error had an interval from 0.31 to 0.39 meters, while the magnitudes of the distance-moved error had an interval from 0.18 to 0.22 meters, and the angle-turned error from 2.31 to 2.75 degrees. When compared to the mean magnitude of errors from this condition before, 3.78 meters, 1.71 meters, and 31.52 degrees, respectively, we see that locomotion errors account for a very small portion of the overall errors. This confirmed our hypothesis that wayfinding errors accounted for most of the errors seen in the main study.

GENERAL DISCUSSION AND DESIGN GUIDELINES

Our findings provide strong evidence that users perform 3D navigation tasks involving path integration more effectively on physically large displays than on smaller displays, even when the same environments were viewed at equivalent visual angles. In fact, in our simple triangle completion task, users performed about 17% better on the large display. Since more complex navigation tasks involving path integration can be decomposed into a series of triangle completions [12], we could imagine the improvements cascading and leading to much greater overall benefits of using large displays. While there could be other ways to increase navigation performance, few of the alternatives provide as simple an extension to current tasks and methods as increasing the physical size of displays.

The effects we observed might be explained by the hypothesis that large displays provide users with a greater sense of presence within the virtual environment, biasing them into using more efficient egocentric strategies. One concern with this explanation is that other mechanisms, such as interactivity, may affect task performance by evoking similar strategies. These mechanisms might then negate the effects provided by the large display. Our results show that this is not the case, and that effects induced by large displays are independent of those induced by differing levels of interactivity. This means that designers can safely use different control mechanisms and continue to experience the benefits of their large display systems.

However, contrary to our initial hypotheses, our findings suggest that active joystick control is detrimental in the set of tasks we tested. We believe that this negative effect can be explained by the attention-cue tradeoff imposed by the new interaction mechanism and environment. The unfamiliar task of using a joystick to navigate the 3D virtual environment required a great deal of attention for our users, who indicated that they did not normally play 3D video games. This additional attention requirement probably impaired the creation of mental representations during the encoding phase in the main study. Because of the disparate reports of the effects of interactivity in various navigation tasks, we advise that researchers examine this manipulation carefully for their specific tasks and demographic before designing any interface and display system.

We conducted the follow-up investigation to examine whether the errors seen in the main study were mainly cognitive wayfinding errors or mechanical control errors. We found that mechanical control errors accounted for only a very small portion of the total error, indicating that most of the error could be attributed to cognitive processes. This is consistent with assumptions in prior path integration literature, which attribute all errors to the encoding process [11]. This indicates that to increase performance with these measures, designers should spend time optimizing cognitive cues, rather than control mechanisms.

FUTURE WORK

In order to further explore the benefits of large displays in navigation tasks, we would like to test other components of effective navigation. We plan on using other measures, such as map memory, in the experimental framework that we have built. We would also like to quantify the utility of training tasks on such displays by examining how well spatial skills learned within virtual environments transfer to real-world navigation.

Given our results and working hypotheses, we would suggest that large displays benefit any task that is either intrinsically egocentric or that can be dichotomized into egocentric or exocentric representations and strategies. We believe that most first person 3D navigation tasks fall into this category. There is no doubt that other design and control tasks would also benefit from these effects. However, it is unclear how much intrinsically exocentric tasks, such as 3D modeling or object manipulation, would be affected. Exploring this remains future work.

To determine the optimal size of displays, we would like to understand how small variations in physical display size affect task performance and whether or not there exist any singularities where performance drastically changes. Finally, we would like to explore how other factors such as field of view interact with physical display size and further affect task performance.

CONCLUSION

We examined the effects that physical display size and interactivity had on path integration in 3D virtual navigation tasks. Our study extends prior findings demonstrating that large displays immerse users more within virtual environments and bias them into using more efficient spatial strategies. In fact, we have shown that users perform 3D navigation tasks requiring path integration more efficiently on large displays than on smaller ones, even when identical scenes were viewed at equivalent visual angles. We have further shown that the distraction imposed by active navigation control using a joystick may outweigh any additional cues it might have provided, at least for the set of tasks we tested. However, effects induced by interactivity seem to be independent of those induced by display size. Our follow up investigation showed that locomotion errors were small and that our results could mainly be attributed to wayfinding errors. Following these results, we have provided in-depth discussion and recommendations for the design and presentation of 3D environments on physically large displays.

ACKNOWLEDGEMENTS

We thank Mary Czerwinski, Dennis Proffitt and his Perceptual Psychology Lab, Jessica Hodgins, Scott Hudson, Robert Kraut, Susan Fussell, and the Stage 3 Research Lab for their discussions on large display systems. We are also grateful to Microsoft Research and IBM Research for their support of this work.

REFERENCES

1. Arthur, K.W. (2000). Effects of field of view on performance with head-mounted displays. *Dissertation Abstracts International*, 61(5B), 2614.
2. Baudisch, P., Good, N., Belloti, V., & Schraedley, P. (2002). Keeping things in context: A comparative evaluation of focus plus context screens, overviews, and zooming. *Proceedings of CHI 2002*, 259-266.
3. Bell, S. (2002). Spatial cognition and scale: A child's perspective. *Journal of Environmental Psychology*, 22, 9-27.
4. Booth, K., Fisher, B., Page, S., Ware, C., & Widen, S. (2000). Wayfinding in a virtual environment. *Graphics Interface*.
5. Bystrom, K.E., Barfield, W., & Hendrix, C. (1999). A conceptual model of the sense of presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, 8(2), 241-244.
6. Chance, S.S., Gaunet, F., Beall, A.C., & Loomis, J.M. (1998). Locomotion mode affects the updating of objects encounter during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence: Teleoperators and Virtual Environments*, 7(2), 168-178.
7. Cutmore, T.R.H., Hine, T.J., Maberly, K.J., Langford, N.M., & Hawgood, G. (2000). Cognitive and gender factors influencing navigation in a virtual environment. *International Journal of Human-Computer Studies*, 53, 223-249.
8. Czerwinski, M., Tan, D.S., & Robertson, G.G. (2002). Women take a wider view. *Proceedings of CHI 2002*, 195-202.
9. Ekstrom, R.B., French, J.W., Harman, H.H., & Derman, D. (1976). Kit of factor-referenced cognitive tests. Educational Testing Service: Princeton, NJ.
10. Flach, J. (1990). Control with an eye for perception: Precursors to an active psychophysics. *Ecological Psychology*, 2, 83-111.
11. Fujita, N., Klatzky, R.L., Loomis, J.M., & Golledge, R.G. (1993). The encoding-error model of pathway completion without vision. *Geographical Analysis*, 25, 295-314.
12. Golledge, R.G. (Ed.). (1999) Wayfinding behavior: Cognitive mapping and other spatial processes. Johns Hopkins University Press: Baltimore, MD.
13. Grabe, M.E., Lombard, M., Reich, R.D., Bracken, C.C., & Ditton, T.B. (1999). The role of screen size in viewer experiences of media content. *Visual Communication Quarterly*, 6, 4-9.
14. Just, M.A., & Carpenter, P.A. (1985). Cognitive coordinate systems: Accounts of mental rotation and individual differences in spatial ability. *Psychological Review*, 92(2), 137-172.
15. Kearns, M.J., Warren, W.H., Duchon, A.P., & Tarr, M.J. (2002). Path integration from optic flow and body senses in a homing task. *Perception*, 31, 349-374.
16. Klatzky, R.L., Loomis, J.M., Beall, A.C., Chance, S.S., & Golledge, R.G. (1998). Spatial updating of self-position and orientation during real, imagined, and virtual locomotion. *Psychological Science*, 9(4), 293-298.
17. Loomis, J.M., Klatzky, R.L., Golledge, R.G., Cicinelli, J.G., Pellegrino, J.W., & Fry, P.A. (1993). Nonvisual navigation by blind and sighted: Assessment of path integration ability. *Journal of Experimental Psychology: General*, 122(1), 73-91.
18. MacIntyre, B., Mynatt, E.D., Volda, S., Hansen, K.M., Tullio, J., & Corso, G.M. (2001). Support for multitasking and background awareness using interactive peripheral displays. *Proceedings of UIST 2001*, 41-50.
19. Patrick, E., Cosgrove, D., Slavkovic, A., Rode, J.A., Verratti, T., & Chiselko, G. (2000). Using a large projection screen as an alternative to head-mounted displays for virtual environments. *Proceedings of CHI 2000*, 478-485.
20. Péruch, P., May, M., & Wartenberg, F. (1997). Homing in virtual environments: Effects of field of view and path layout. *Perception*, 26, 301-311.
21. Philbeck, J.W., Klatzky, R.L., Behrmann, M., Loomis, J.M., & Goodridge, J. (2001). Active control of locomotion facilitates nonvisual navigation. *Journal of Experimental Psychology: Human Perception and Performance*, 27, 141-153.
22. Prothero, J.D., & Hoffman, H.D. (1995). Widening the field of view increases the sense of presence within immersive virtual environments. Technical Report, University of Washington, Seattle, WA, R-95-4.
23. Riecke, B.E., van Veen, H.A.H.C., & Bühlhoff, H.H. (2000). Visual homing is possible without landmarks: A path integration study in virtual reality. Max-Planck-Institut für biologische Kybernetik, Germany.
24. Rieser, J. J. (1989) Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 1157-1165.
25. Slater, M., & Usoh, M. (1993). Presence in Immersive Virtual Environments. *Proceedings of IEEE VRAIS*, 90-96.
26. Tan, D.S., Gergle, D., Scupelli, P.G., Pausch, R. (2003). With similar visual angles, larger displays improve spatial performance. *Proceedings of CHI 2003*, 217-224.
27. Tan, D.S., Stefanucci, J.K., Proffitt, D.R., & Pausch, R. (2001). The Infocockpit: Providing location and pace to aid human memory. *Workshop on Perceptive User Interfaces 2001*.