

Is Semitransparency Useful for Navigating Virtual Environments?

Luca Chittaro

HCI Lab, Dept. of Math and Computer Science,
University of Udine,
via delle Scienze 206, 33100 Udine, Italy
++39 0432 558450

chittaro@dimi.uniud.it

Ivan Scagnetto

HCI Lab, Dept. of Math and Computer Science,
University of Udine,
via delle Scienze 206, 33100 Udine, Italy

scagnett@dimi.uniud.it

ABSTRACT

A relevant issue for any Virtual Environment (VE) is the navigational support provided to users who are exploring it. Semitransparency is sometimes exploited as a means to see through occluding surfaces with the aim of improving user navigation abilities and awareness of the VE structure. Designers who make this choice assume that it is useful, especially in the case of VEs with many levels of occluding surfaces, e.g. virtual buildings or cities. This paper is devoted to investigate this assumption with a proper experimental evaluation on users. First, we discuss possible ways for improving navigation, and focus on implementation choices for semitransparency as a navigation aid. Then, we present and discuss the experimental evaluation we carried out. We compared subjects' performance in three conditions: local exploitation of semitransparency inside the VE, a more global exploitation provided by a bird's-eye-view, and a control condition where neither of the two features was available.

Categories and Subject Descriptors

I.3.6 [Computer Graphics]: Methodology and Techniques – *Interaction techniques*. H.5.2 [Information Interfaces and Presentation]: User Interfaces – *Interaction styles, evaluation*. H.1.2 [Models and Principles]: User/Machine Systems – *Human factors*.

General Terms

Experimentation, Human Factors.

Keywords

navigation aids, evaluation, wayfinding.

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 republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

VRST '01, November 15-17, 2001, Banff, Alberta, Canada.
Copyright 2001 ACM 1-58113-385-5/01/0011...\$5.00.

1. INTRODUCTION

One of the most relevant usability issues for a Virtual Environment (VE) is the navigational support provided by its user interface. In current VEs, people often become disoriented and tend to get lost. Inadequate support to user navigation is also likely to result in users leaving the VE before reaching their targets of interest, or to leave users with the feeling of not having adequately explored the visited VE. These problems become even more critical in the case of the growing number of VEs on the Web, where users are very likely to leave prematurely the site if they encounter usability problems. It is also interesting to note that, besides the traditional VR applications that include 3D models of cities and buildings, VEs based on architectural metaphors are being increasingly used to visualize *abstract* information in domains as diverse as computer networks [17], databases [6], e-commerce [5], information systems [9], operating systems [12], and program code [13]. As a consequence, the issue of 3D navigation of VEs is going to attract researchers who do not currently belong to the VR community.

In some systems, semitransparency is exploited as a means to see through occluding surfaces, assuming that it will improve user navigation abilities and awareness of the VE structure. However, authors who employ transparency as a navigation aid typically acknowledge that this is just an assumption and mention a lack of user testing as a major limitation of their work (e.g., [17]). This paper is thus devoted to investigate the considered assumption with a proper experimental evaluation on users, which tests two different approaches to the exploitation of transparency, and includes a control condition where transparencies are not available.

2. IMPROVING NAVIGATION IN VEs

In general, navigation can be informally defined as the process whereby people determine where they are, where everything else is, and how to get to particular objects or places [14]. Human navigation abilities in the physical world have been studied both in psychology and architecture (a concise survey is provided by [18]). There are two distinct types of navigational knowledge of an environment (that can be generalized and apply to the case of VEs), each one supporting different behaviors. *Route knowledge* (also called *procedural knowledge*) is egocentric, describes paths

between locations, is usually gained through personal exploration, allows reaching a destination through a known route, but does not allow recognizing unfamiliar alternate routes (e.g., short-cuts). *Survey knowledge* is exocentric, describes the relationships among locations, can be gained also through study, provides the mental equivalent of a map (often referred to as a *cognitive map*), and allows to recognize alternate routes. Some psychological studies investigated the sources of spatial knowledge acquisition. For example, Thorndyke and Hayes-Roth [20] compared spatial judgment abilities of subjects who learned an environment only from personal exploration or only from a map, highlighting the difficulty of changing perspective (e.g., subjects who acquired knowledge only from the exocentric map perspective were most error prone in tasks that required to translate their knowledge into a response within the environment).

To improve the acquisition of navigational knowledge in VEs, some lessons can be learned from the design of real-world environments. For example, in the design of buildings, architects aim at reducing wayfinding problems for the people working or visiting the building [18] by increasing *visual access* (i.e. the number of parts of the environment which can be seen by a person from her position in space) or including navigational cues (e.g., room numbers, names of buildings, landmarks). *Landmarks* are distinctive environmental features (e.g., a statue, a river, a town square, ...) functioning as reference points [22].

2.1 Lines of Research

Two main lines of research can be identified among those projects which focus on improving user navigation in VEs. One category of projects is devoted to identify guidelines for designing more navigable environments. Some of these guidelines are being derived from other fields which have already faced the problem in the physical world. For example, extensive work exists on the design and placement of landmarks in such diverse areas as urban planning, geography, and psychology. An attempt to summarize the available knowledge on this topic in the form of guidelines is provided by Vinson [22], who aims at allowing users to apply their real world navigational experience. Some experiments have been carried out to determine interesting aspects of landmark design, e.g., it has been shown that landmarks should be *memorable* to adequately help users: in a study [15] contrasting the use of familiar 3D objects and abstract art to build landmarks, the former provided a significant improvement in navigability over the latter.

The second category of projects focuses on providing the user with electronic navigation aids to augment her capabilities to explore and learn. A well known example of a navigation aid is the electronic map of the environment to help users orient themselves.

The two above mentioned lines of research are obviously strongly related, e.g. different guidelines to design an easy-to-navigate VE can apply if the user is able to freely move to any position in 3D space or (s)he is instead constrained to predefined planes. In the following, we will concentrate on *navigation aids*.

2.2 Navigation Aids

The perspective of the VE provided by a navigation aid can be of two types: the *first-person* perspective aims at providing an egocentric viewpoint, as if the user were immersed in the environment (considering the current user's position, it shows the part of environment which should be in front of her own eyes); the *third-person* perspective shows instead an exocentric viewpoint where the user can see her current position explicitly marked in the environment. *Third-person* perspectives can require a considerable mapping effort to be correctly interpreted by the user (e.g., consider the typical real-world situation where someone is trying to find her way in a city by using a map and has to translate the exocentric view of the map into her egocentric view).

The first navigation aids to be proposed have been electronic analogues of the tools commonly used by people to navigate in unfamiliar real world environments. From this perspective, the most common choice has been to make an overview (in the form of an electronic map) of the environment available to the user. Besides this more traditional solution, novel navigation aids have been recently proposed by different authors, e.g. [7],[8],[11],[19].

Some approaches are based on augmenting the electronic map with features that are unavailable in a real-world map, such as the capability of self-orientation (e.g., the upward direction of the map can be arranged in such a way that it always shows what is in front of the viewer). The study presented by [8] analyzes the performance of users with a map of this kind and other three treatments: a control treatment, a grid treatment where a radial grid is superimposed on the world, and a map/grid treatment where both aids are present. The study reports that the treatments which included the map were those which supported the most effective searches (but no statistical comparisons between treatments are provided).

The *Worlds in Miniature (WIM)* metaphor [19] is an interaction technique that offers a miniature representation of the environment, standing between the user and the environment itself, and held in a virtual hand of the user. The user can directly manipulate both the WIM and the environment (changing something on one of the two changes the representation of the other and vice versa).

An attempt to provide miniature worlds which do not overlap the environment is given by the *Wordlets* approach [11]. Wordlets are 3D interactive thumbnails which are displayed outside the environment and can be explored and manipulated in the same way. They have been shown to be more effective than text and images [11] for building a guidebook to aid wayfinding in a virtual city. The guidebook represents landmarks leading to a destination.

Other navigation aids such as flying, spatial audio, and breadcrumb markers, are illustrated by [7], but the study of their effectiveness is only informal. Leading users through a preliminary tour of the environment (where the user can actively follow a pre-determined path or, alternatively, be passively moved through it) is a

technique used by [18] to allow users familiarizing with the environment before engaging in more specific tasks.

2.3 Related Work on Semitransparency

Semitransparency is being used to improve perception and understanding of the user's working environment in novel styles of graphical user interfaces. A representative example is the *See-Through Interface* [1],[2] which provides semi-transparent interactive tools, called *Toolglass* widgets, that are used and combined together in an application work area. In this context, *magic lenses* are transparent windows that can be moved across the screen, changing how the output of objects falling under their scope is visualized. For example, a magnifying lens allows one to magnify parts of objects. In [21], applicability of magic lenses to 3D objects is explored more thoroughly, e.g. to provide Xray volumetric lenses (when the user passes the Xray lens over an object, the inside of the object is revealed). Viega *et al.* [21] suggest that previous ideas in navigation aids could be reformulated as 3D magic lenses: e.g., they propose to think about *Worlds In Miniature* (see Section 2.2) as a 3D volumetric lens (more precisely, as a volumetric reduction lens).

As an aside, it is interesting to note that semitransparency is exploited for navigation purposes in some videogames. A representative example of how it is generally used is provided by *Sanitarium* [10]. The game is based on a third-person perspective with a fixed camera capturing a portion of the environment; the user position in the environment is marked by an avatar which can be directly manipulated. Since the avatar can be led to areas hidden by doors, walls and obstacles, it can become hidden by objects (e.g., after entering a room, the walls of the room would hide the avatar). The provided solution is to make occluding objects physically disappear (and reappear when the avatar moves away) by means of a dissolve effect. In these cases, transparency is exploited to avoid occlusion effects, but not to allow the user to freely gain spatial information for her navigation purposes.

3. THE IMPLEMENTED NAVIGATION AIDS

An important choice for our experimental study concerned which specific navigation aids to employ. To make the evaluation more thorough, we decided to implement two different navigation aids, one based on a first-person perspective and the other on a third-person perspective. While the choice for the latter was relatively straightforward (a bird's-eye-view map is a representative solution adopted in many systems), designing the first-person perspective aid required a more careful consideration because no representative solution emerges from the literature. The choices made for the first-person perspective aid are discussed in detail in Section 3.1.

3.1 The STS Navigation Aid

The proposed navigation aid allows the user to visually inspect the parts of the environment which are adjacent to the one where

(s)he is positioned, by clicking on visually occluding surfaces to make them semitransparent. For conciseness, surfaces which support this functionality will be called STS (*See-Through Surfaces*) in the following. STS are meant to allow the user to gain survey knowledge about the relationships among her current location and adjacent locations. The functionality could be also considered as a magic lens (see Section 2.3) in a first-person perspective.

In the one-floor buildings we have designed for the user study, the STS functionality has been implemented by making walls sensitive to mouse clicks. When the user clicks on a wall, the wall becomes semitransparent. Semitransparency automatically deactivates after some seconds. If the user wishes to deactivate it earlier, (s)he can do so by clicking again on the wall. As an example, Figure 1 illustrates a corridor delimited by solid walls. Figure 2 shows the same corridor after a user has activated three STS: while in the former situation the user can only perceive to be in a corridor with a right turn at the end, in the latter much more information can be gained from the activated STS. Indeed, semitransparency on the right reveals a room with a couch and a table (establishing relations between current user position and a room she might have already visited or is going to visit). Moreover, the STS on the left and front allow the user to easily understand that she is in a corridor on the perimeter of the building (since she can see parts of the external garden), and semitransparency on the front provides also information about the user's relative position with respect to the entrance of the building (a column of the entrance is visible).

A design choice we had to make concerned how to delimit the single semitransparent surface which is activated by a mouse click. The criterion we adopted is to use intersections with other surfaces as delimiters. In other words, when the user clicks on a wall, that wall becomes semitransparent up to the point where (at its sides) it meets other walls. For example, in Figure 2, clicking on the wall at the right has made it semitransparent up to the point where it meets a perpendicular wall (which can be seen as opaque) at the end of the corridor. This approach was chosen because it makes it easy for the user to predict which surface will become semitransparent. For example, looking at Figure 1, three distinct walls can be perceived (each of the walls is perceived as a single object up to the point where its continuity is broken by an intersecting wall): a mouse click on one of the three will make semitransparent only that specific wall.

VEs in our experiments have been built in VRML. In particular, to implement the STS functionality, we exploited the notion of event provided by VRML and the related routing mechanism. In detail, every surface involved in the STS functionality consists of a group including the surface itself, a touch sensor, a time sensor (to determine the time for automatic deactivation) and a route to/from a simple Javascript program handling the on/off switching of semitransparency. The script is used to set the proper transparency value and to overcome the lack of state information in VRML (a flag encodes the semitransparency status of the




Figure 1. A corridor in one of the buildings (This figure is reproduced in color on page 000).




Figure 2. Corridor of Figure 1 with 3 STS activated (This figure is reproduced in color on page 000).



Figure 3. BEV of a building (This figure is reproduced in color on page 000).

surface, allowing the script to distinguish whether a mouse click activates or deactivates semitransparency).

It is worth noting that the choice of the level of transparency is crucial: while an insufficient level of semitransparency would not allow the user to clearly distinguish the environment beyond a wall, an excessive level of semitransparency could mislead the user, giving the impression that passages are available where they are not. In VRML, every object has several associated attributes, one of which is the transparency of the material, that can be set to a real value ranging from 0 (total opacity) to 1 (complete transparency). In the implementation of the buildings, we found a value of 0.7 to be a good choice, allowing one to see clearly what is hidden by a wall, while retaining the perception that the wall is still in its place, i.e., no part of it has vanished. This conclusion was reached by testing various possibilities during the development of the environments, allowing some colleagues to visit the buildings and checking with them if there were any misperceptions. This pilot study was used to test the value of other possible parameters. In particular, another parameter that showed the need for a careful setting was the amount of time after which a semitransparent surface automatically deactivates. Setting no limit for this time (thus relying only on user input for the deactivation of semitransparency) does not only require more mouse clicks to the user, but also easily leads to visually confusing situations. Indeed, the user can leave surfaces in semitransparent state and proceed in the exploration of the environment. This can quickly lead to situations where several levels of semitransparent surfaces are seen one behind the other, with resulting difficulties in scene interpretation. Setting a too large time limit can still lead to the problem, while a too short one would not allow the user to collect sufficient information and could be irritating and time consuming (due to the repeated activations that the user would need). In our case, an amount of time of 8 seconds has shown to be a good compromise and has been adopted for the study.

3.2 The BEV Navigation Aid

As we did with the STS aid, we tried to choose a “best implementation” for the aid based on bird’s-eye-view (hereinafter, BEV). Among the possible implementations seen in the literature, we adopted one of the most informative: when the user activates BEV, the *full* screen space is used to show the whole building from above, making the top transparent (so that the view appears like a map) and highlighting current user’s position in the VE. Figure 3 is a screenshot of the BEV aid applied to the same building shown in previous figures (user position is indicated by a red ball, which in this figure is at the entrance). Clicking on the large arrow on the bottom right of the screen allows switching between the user’s egocentric perspective and the exocentric view provided by the BEV aid. The arrow orientation is upwards in first-person view, and downwards in BEV, to suggest the change in height of the viewpoint between the two perspectives.

4. THE EXPERIMENTAL STUDY

4.1 Task, Conditions and Hypotheses

The experiment concerned a wayfinding task where subjects had to find a path to a specific object (graphically represented by a well) starting from a predefined position (the entrance) in a VE of a building. Travel inside the environment was based on a first-person perspective “walk” mode, in which users controlled movement with the four arrow keys on the keyboard: by pressing the forward and backward keys, the user moved respectively forward or backward at a constant velocity, while the right and left keys were used to turn. Collision detection was used to prevent users to move through objects and surfaces. Each subject had to perform the task under three different conditions in a standard within-subjects design: a control (CTRL) condition where no navigation aids were available, a STS condition where STS was the available navigation aid, and a BEV condition where BEV was the available navigation aid. Figure 1 is a screenshot of a building under the CTRL condition; Figure 2 shows the same sub-section of the building with some semitransparencies activated under the STS condition; Figure 3 is a screenshot of a BEV for the same building.

Since using the same building for the three conditions would have caused serious learning effects due to the acquisition of navigational knowledge, we designed three different buildings so that each subject visited each of the buildings only one time. The three buildings represented different and deliberately unfamiliar environments: one was a stone building (Figures 1, 2, and 3 are all taken from the stone building), the other two were respectively inspired to a fantasy-gothic look and to a science-fiction look. The only constant graphical element in the three environments was the 3D model of the object to be found (i.e., the well), while all other elements changed. All possible care was taken in order to ensure that the navigational complexity of the three environments was the same. To this purpose, the following parameters have been controlled and kept constant: size of the building, number of rooms, number of doors, number of landmarks, position and distance of the landmarks on the map, length of the path from the starting point to the destination, and number of choice points along that path. Anything that can count as a landmark has been considered, including both models of objects and types of textures. As an aside, it must be noted that a difference in the number and complexity of these landmarks can also cause a significant difference in rendering speed among the different VEs, which is an additional motivation to hold them constant.

Each room and landmark was different and unique. Landmarks were made memorable by adopting familiar 3D objects, such as common house furniture (e.g., tables, chairs, couches, lamps, ...), or other easy-to-recognize objects (e.g., we have used swords, plants, loudspeakers, crosses, coffins, ...).

Interaction with the environment took place through keyboard and mouse. The four arrow keys on the keyboard were used for traveling inside the building, while the mouse was devoted to the

activation of the available navigation aid. In the CTRL condition, the mouse did not allow one to activate any aid; in the STS condition, users could point to any wall with the mouse cursor and click on it to make it semitransparent; in the BEV condition, subjects could click the large arrow described in Section 3.2 to activate/deactivate BEV. The BEV perspective was presented in full-screen format, it was mutually exclusive with the first-person perspective travel mode, and the arrow keys had no effect while the user was in BEV perspective.

Our first hypothesis for the experiment was that both approaches

Table 1. Post-hoc comparisons.

	Mean Difference	Significance
CTRL vs. STS	126	p<0.05
STS vs. BEV	142	p<0.05
CTRL vs. BEV	268	p<0.001

Table 2. Subjective impressions.

	BEV	STS
Very Easy	15	11
Easy	3	6
Doable	0	1
Difficult	0	0
Very Difficult	0	0
MEDIAN	Very Easy	Very Easy

to the use of transparency improve user wayfinding performance. In making this hypothesis, we were motivated by two main considerations. First, from an architect's point of view, semitransparency can be seen as a way of increasing visual access, thus supporting an increase of user's navigational abilities (see also Section 2). Second, both functionalities can be easily and quickly understood even by the most casual users (this can be less likely with some novel navigation aids mentioned in Section 2). On one side, the capability of seeing through physical objects is a typical supernatural power for different cultures, and has also been popularized by fantasy and science-fiction literature (e.g. the X-ray vision super power in Superman's comic books) up to the point that it has even been used to sell bogus merchandise such as "X-Ray Specs". On the other side, perception studies have shown that the partial occlusion effect given by semitransparency is still a depth cue that can be readily perceived by the viewer. In particular, the thorough studies on human subjects by Zhai *et al.* [23] have shown that semitransparency effectively reveals spatial relationships among objects within VEs, particularly in the depth dimension, so that the user can perceive and locate objects with respect to each other effortlessly, easily comprehending the depth relation between a semitransparent surface and objects that are in front or behind it.

Our second hypothesis was that BEV would improve performance no less than STS. The second hypothesis was formulated by considering that, although BEV provides an exocentric perspective which could be difficult to map into the egocentric one used for moving, the "best implementation" chosen for BEV provides a global overview of the environment, containing much more information than the more local view provided by STS. In the context of the considered wayfinding task, the global view allows the user to see and identify the full path from start to destination in the building, and should thus greatly affect wayfinding performance, even if the mapping effort is difficult.

4.2 Experimental Design and Procedure

Subjects were recruited among new students in Computer Science on their first days of classes at our University. The main motivation for accepting to take part in the study was to visit our Department and have a look at the laboratories. The majority of students was 19 years old with a few exceptions. More specifically, the age of subjects ranged from 18 to 31, averaging at 20. We recruited a total of 22 subjects, all male. The experiment was successfully completed by 18 subjects. Data concerning the other 4 subjects had to be discarded, because three of them completely lost themselves in at least one of the environments (and asked for evaluator's help), while one suffered a mild form of motion sickness during the experiment.

First, subjects filled out a brief questionnaire on their prior experience with computers and 3D environments. All subjects were computer literate, spending at least 4 hours per week using computers (mean number of weekly hours was 10), and were regular users of 3D computer games (every subject played for at least 1 hour per week, and the mean number of weekly hours devoted to computer games was 4). Subjects were also asked if they were left-handed, because the keyboard and mouse positions were arranged assuming a right-handed user. Only one subject was left-handed and he was invited to rearrange the mouse and keyboard position in case it was not comfortable for him.

Next, subjects were allowed to spend unlimited time in a very simple virtual building (unrelated to those used for the experimental task) until they felt familiar with the controls and the navigation aids (both navigation aids were available in this initial training environment). When the subjects felt ready, the experiment began and subjects were introduced to the experimental task. For a quicker comprehension, the task was presented to subjects as a 3-levels computer game where, for each level, they had to enter a different building and find as quickly as possible a source of water (graphically represented by a well) inside. Subjects had a 1-minute break between the completion of a level and the start of the following one.

A within-subjects randomized design has been used. We considered the availability of navigation aids as the independent variable for the experiment, while the dependent variable was the time required to complete the task. The order of visit of the three




Figure 4. Mean time to complete the task.

buildings and the order of the three navigational conditions changed independently for each subject in such a way that: (i) every navigational condition was presented an approximately equal number of times as a first, second, and third condition, (ii) every building was visited an approximately equal number of times as a first, second, and third environment, and (iii) there was no fixed association between a specific building and a specific navigational condition (to counterbalance the effects of a possibly higher navigational difficulty of an environment over the others, in case the effort to keep complexity of the three buildings constant might have left something unaccounted for).

Finally, subjects filled out a second questionnaire where they were asked to qualitatively rate their subjective impressions on the two navigation aids, and could add free comments.

The hardware used for the experiment was a standard 19 inch Trinitron monitor and a Pentium III PC equipped with an Open GL hardware accelerator (Nvidia TNT2 Ultra). The full screen was devoted to present the selected view of the current environment.

4.3 Analysis and Results

A one-way analysis of variance (ANOVA) has been performed. The within-subjects variable was the availability of navigation aids with three levels: no aids (CTRL), STS, and BEV. The dependent variable was the time required to complete the task. The result of ANOVA ($F(2,34)=16.05$, $p<0.001$) indicated that the effect was significant. We thus employed the Scheffé test for post-hoc comparisons among means. The values of means (376 sec. for the CTRL condition, 250 sec. for the STS condition, and 108 sec. for the BEV condition) are graphically illustrated by Figure 4, while results of the post-hoc analysis are given in Table 1. It turns out that user performance in the control condition is significantly lower than performance both in the STS condition ($p<0.05$) and in the BEV condition ($p<0.001$), and performance in the STS condition is significantly lower than performance in the BEV condition ($p<0.05$).

Results of subjective impressions collected with the second questionnaire are summarized in Table 2. Subjects were asked to rate how difficult they found using the two navigation aids. The table shows clearly that none of the subjects found them difficult

to use, and both aids were very favorably rated by users (with BEV getting slightly better scores).

4.4 Qualitative Observations

We observed that the strategies adopted by subjects in using STS were quite different. More precisely, users differed in: (i) frequency of use (ranging from a few to several tens of activations during the visit of a single environment); (ii) number of simultaneously activated STS (from a single one to as many as reachable on the screen); and (iii) simultaneity of movement (some subjects stopped to activate STS, while others continued to move). In general, subjects who benefited the most from the use of the functionality seemed to be those who tended to use it more frequently, on more than one STS at a time, and while continuing to move.

Many users tended to lose a noticeable amount of time to operate the mouse in order to click the different STS inside their viewpoint. This suggests that an alternative implementation of the STS functionality could be tried, by reformulating it as a X-ray Vision empowerment: if the user turns X-ray Vision on, STS could automatically activate as soon as they enter the user's visual field, without the need of pointing and clicking them.

5. CONCLUSIONS

The experimental study presented in this paper has confirmed the hypothesized positive effects on user navigation performance of VEs. In the following, we propose some considerations on the study, and identify further lines of experimentation.

First, from the point of view of generalizing the obtained better results for BEV vs. STS to other BEV implementations in the literature, it must be noted that the BEV used in this paper was able to provide an overview of the entire environment in a single screen. Other implementations of BEV might be more difficult to use and/or less informative. For example, if the BEV of the entire environment were too large to fit the screen, gaining a full overview of the VE would require scrolling operations which could considerably lower overall performance. Some implementations choose to reduce the space devoted to the map in order to show both the first-person and third-person view of the environment in the same screen, e.g. a small "radar" window is often shown in a corner of a large first-person view. These solutions preserve the continuity of the first-person navigation experience (switching to a full screen BEV breaks instead that continuity). From a wayfinding perspective, a smaller viewable map area makes it more difficult to determine the desired path, but seeing both the first-person and third-person views simultaneously allows the user to establish relations between them much more easily than switching between different screens. It would be thus interesting to contrast these different approaches to BEV in more detail.

It should also be noted that the paths in our VEs were two-dimensional (i.e., horizontal with turns) as we are used with corridors in everyday life. Research [3],[4] that studied navigation with (less familiar) three-dimensional paths (i.e., those moving

along all three principal axes, including the vertical one) showed that adding a third dimension to the paths significantly decreases user's performance in gathering information and maintaining spatial orientation in the VE. It would thus be interesting to compare STS and BEV in more complex VEs with three-dimensional corridors: while the STS aid would still allow the user to easily get the relative positions of close items wrt to her position by seeing through surfaces around her, examining the BEV to derive the same information could become much more difficult.

Another factor to consider is that subjects in our experiment had some experience in 3D navigation from the use of videogames. This could suggest additional experiments on users who are not familiar with videogames, although we would not expect changes in the final ranking of the three conditions.

In deriving design guidance from the study results, BEV and STS should not be considered as mutually exclusive alternatives. The types of information they provide are different and complementary in perspective (egocentric vs. exocentric), scope (local vs. global), and level of detail (fine-grained vs. coarse-grained), suggesting a combined exploitation. Evidence that combined use of aids providing local and global information has a synergic effect is emerging in recent psychology studies on the use of maps in very-large-scale environments [16], for which it would not be possible to fit a detailed BEV into a single screen. First-person and third-person navigation aids differ also in the cognitive maps the user develops of the environment. In particular, navigating an environment directly is more likely to result in survey knowledge which is orientation-independent, while survey knowledge acquired from an external map tends to be orientation-specific [8]. This suggests investigating the acquisition of survey knowledge of the VE (e.g., in terms of estimation abilities for relative distances and orientation) as an additional line of research.

Finally, an interesting direction of research in integrating STS with other empowerments concerns the possibility of *moving* through walls. While in our experiment users were constrained to stay within the corridors as it is typical of some VE applications (e.g., training, games,...), other researchers [3],[4] have investigated navigation scenarios where collision detection was disabled, allowing users to freely move through walls to gain spatial knowledge of the VE. The addition of STS to those scenarios could prove worthwhile: indeed, when users cannot see what lies behind a wall, they are forced to adopt "blind" strategies (e.g., traveling along parallel lines) to exploit the capability of moving through walls; with STS, they could see in advance if the location hidden by a wall is worth visiting, without the need of actually traveling to it.

6. ACKNOWLEDGMENTS

This work has been partially supported by the Italian Ministry of University and Research (MURST), under the COFIN 2000 program.

We would like to thank one of the anonymous reviewers for the valuable comments which helped us improve the paper.

7. REFERENCES

- [1] Bier, A., Stone, C., Pier, K., Fishkin, K., Baudel, T., et al. Toolglass and Magic Lenses: The See-Through Interface. In Proceedings of SIGGRAPH '93 (Anaheim CA, 1993), ACM Press, 73-80.
- [2] Bier, E. A., Stone, M. C., Pier, K., Fishkin, K., and Baudel, T. A taxonomy of see-through tools. In Proceedings of CHI '94 (Boston MA, 1994), ACM Press, 517-523.
- [3] Bowman, D. A., Koller, D., and Hodges, L. F. A Methodology for the Evaluation of Travel Techniques for Immersive Virtual Environments. *Virtual Reality: Research, Development, and Applications*, 3(2), 1998, pp. 120-131.
- [4] Bowman, D., Davis, E., Badre, A., and Hodges, L. Maintaining Spatial Orientation during Travel in an Immersive Virtual Environment. *Presence: Teleoperators and Virtual Environments*, 8(6), 1999, 618-631.
- [5] Chittaro, L., and Coppola, P. Animated Products as a Navigation Aid for E-commerce. In Proceedings of CHI 2000 (The Hague, The Netherlands, 2000), Extended Abstracts Volume, ACM Press, 107-108.
- [6] Costabile, M. F., Malerba, D., Hemmje, M., and Paradiso, A. Building Metaphors for Supporting User Interaction with Multimedia Databases. In Proceedings of 4th IFIP Conference on Visual DataBase Systems (VDB-4), Chapman & Hall, 1998, 47-66.
- [7] Darken, R.P., and Sibert, J.L. A Toolset for Navigation in Virtual Environments. In Proceedings of UIST '93 (Atlanta GA, 1993), ACM Press, 157-165.
- [8] Darken, R.P., and Sibert, J.L. Wayfinding Strategies and Behaviors in Large Virtual Worlds. In Proceedings of CHI '96 (Vancouver, Canada, 1996), ACM Press, 142-149.
- [9] Dieberger, A., and Frank, A. U. City Metaphor for Supporting Navigation in Complex Information Spaces. *Journal of Visual Languages and Computing* 9, 1998, 597-622.
- [10] Dream Forge Entertainment/ASC Games, Sanitarium, www.ascgames.com, 1998.
- [11] Elvins, T.T., Nadeau, D.R., Schul, R., and Kirsh, D., Worldlets: 3D Thumbnails for 3D Browsing. In Proceedings of CHI '98 (Los Angeles CA, 1998), ACM Press, 163-170.
- [12] Hopkins, J. F., and Fishwick, P.A. A Three-Dimensional Synthetic Human Agent Metaphor for Modeling and Simulation. *Proceedings of the IEEE* 89(2), 2001, 131-147.
- [13] Knight, C., and Nunro, M. Virtual but Visible Software. In Proceedings of IV 2000: International Conference on Information Visualization, IEEE Press, 2000.

- [14] Jul, S., and Furnas, G.W. Navigation in Electronic Worlds. SIGCHI Bulletin 29(4), 1997, 44-49.
- [15] Ruddle, R.A., Payne, S.J., and Jones, D.M. Navigating Buildings in "desk-top" virtual environments: Experimental Investigations Using Extended Navigational Experience. Journal of Experimental Psychology: Applied 3(2), 1997, 143-159.
- [16] Ruddle, R.A., Payne, S.J., and Jones, D.M. The Effects of Maps on Navigation and Search Strategies in Very-Large-Scale Virtual Environments. Journal of Experimental Psychology: Applied 5(1), 1999, 54-75.
- [17] Russo Dos Santos, C., Gros, P., Abel, P., Loisel, D., Trichaud, N., and Paris, J.P. Metaphor-aware 3D Navigation. In Proceedings of InfoVis 2000: IEEE Symposium on Information Visualization, IEEE Press, 2000.
- [18] Satalich, G.A. Navigation and Wayfinding in Virtual Reality: Finding the Proper Tools and Cues to Enhance Navigational Awareness. MS Thesis, www.hitl.washington.edu/publications/satalich, 1995.
- [19] Stoackley, R., Conway, M.J., and Pausch, R. Virtual Reality on a WIM: Interactive Worlds in Miniature. In Proceedings of CHI '95 (Denver CO, 1995), ACM Press, 265-272.
- [20] Thorndyke P.W., Hayes -Roth B. Differences in Spatial Knowledge Acquired from Maps and Navigation. Cognitive Psychology 14, 1982, 560-589.
- [21] Viega, J., Conway, M.J., Williams, G., and Pausch, R. 3D Magic Lenses. In Proceedings of UIST '96 (Seattle, Washington, 1996), ACM Press, 51-58.
- [22] Vinson, N.G. Design Guidelines for Landmarks to Support Navigation in Virtual Environments. In Proceedings of CHI '99 (Pittsburgh PA, 1999), ACM Press, 278-284.
- [23] Zhai, S., Buxton, W., and Milgram, P. The Partial-Occlusion Effect: Utilizing Semitransparency in 3D Human-Computer Interaction. ACM Transactions on Computer-Human Interaction 3(3), 1996, 254-284.

This paper will appear in the Proceedings of the 8th ACM Symposium on Virtual Reality Software & Technology (VRST 2001), ACM Press, New York, November 2001.