Occlusion Management in Immersive and Desktop 3D Virtual Environments: Theory and Evaluation

Niklas Elmqvist and M. Eduard Tudoreanu

***Abstract*—We present an empirical usability experiment studying the relative strengths and weaknesses of three different occlusion management techniques for discovering and accessing objects in information-rich 3D virtual environments. More specifically, the study compares standard 3D navigation, generalized fisheye techniques using object scaling and transparency, and the BalloonProbe interactive 3D space distortion technique. Subjects are asked to complete a number of representative tasks, including counting, pattern recognition, and object relation, in different kinds of environments and on both immersive and desktop-based VR systems. The environments include a free-space abstract 3D environment and a virtual 3D walkthrough application for a simple building floor. Our results confirm the general guideline that each task calls for a specialized interaction—no single technique performed best across all tasks and worlds. The results also indicate a clear trade-off between speed and accuracy: simple navigation was the fastest but also most error-prone technique, whereas spherical BalloonProbe and transparency-based fisheye proved the most accurate but required longer completion time, making it suitable for applications where mistakes incur a high cost.**

***Index Terms*—Occlusion reduction, occlusion management, 3D space distortion, interaction techniques, evaluation.**

I. INTRODUCTION

Virtual worlds can in general be used for two different purposes: to either mimic reality in an effort to provide understanding about a real place, such as virtual walkthrough applications and photo-realistic rendering, or as a canvas for representing abstract information so that a viewer can make sense of it and reason about it, such as for information visualization. Regardless of purpose, most useful worlds are rich in objects due to the amount of information they have to convey. High object density inevitably leads to clutter and occlusion, causing the virtual world to be difficult to use effectively.

Fortunately, there are many ways of making sense of a crowded 3D world, such as distorting space, using 3D thumb-nails, making note of landmarks, utilizing navigational aids, and so on. However, it is often unclear for what kinds of tasks and types of worlds each technique is best suited, i.e., the context of the technique. Examples of such contexts include locating an object in an architectural walkthrough, sifting through volumetric 3D data, or identifying specific shapes in a large collection of objects. Very little work has been done to help designers understand in which situations various techniques work best, especially in an immersive setting.

In this paper, we try to remedy this problem by conducting a comparative study of some popular techniques for occlusion management in 3D environments. The aim of this study was primarily to identify situations where different techniques are most efficient in order to help designers make the best choice in terms of efficiency when building their virtual worlds. Subjects were asked to perform a number of representative tasks, including basic object counting, relating different types of objects, and recognizing world-sized patterns. We considered different types of worlds that fall into two main categories: abstract 3D spaces populated with 3D primitives, and architectural walkthrough-like environments. For each type of world, we varied the overall object density in order to detect possible points where techniques break down. In addition, we performed the study both in an immersive CAVE device as well as on a standard PC workstation.

Many different object de-cluttering techniques exist in the literature today, ranging from multiple views and space distortion techniques, to those employing transparency or direct manipulation. Out of necessity, this work deals only with a small sample of these. We focus on techniques that are interactive and directly controlled by a user exploring the virtual world and which do not require extra operations such as selection, sorting, and filtering. We also disregard automatic and query-based techniques for eliminating distracters—although such techniques are useful when the objects of interest are already known and selected or grouped together, usage scenarios where the user must explore and determine what they are looking for on the fly are more general and relevant for our purpose.

The techniques included in this study are generalized fisheye views [1], BalloonProbe space distortion [2, 3], and standard 3D camera navigation controls. More specifically, we study two different variants of generalized fisheye views based on object scale and transparency. For the BalloonProbe technique, we study both spherical and wedge-shaped probe geometry. We have built basic implementations of all techniques using a generic test platform, allowing us to conduct experiments of the different methods side by side with the exact same test parameters.

In the next section, we will go through the related work in this field. We then give a general model for the occlusion problem, followed by a discussion of using generalized fisheye

---

Manuscript Received on February 23, 2007.

Niklas Elmqvist is a postdoctoral researcher at the In-Situ Lab of INRIA Futurs located at Universit´e Paris-Sud XI in Paris, France. E-Mail: elm@lri.fr.

M. Eduard Tudoreanu is an Assistant Professor in the Department of Information Science at the University of Arkansas at Little Rock. His areas of research are information visualization and Virtual Reality. E-Mail: metudoreanu@ualr.edu

views and the BalloonProbe for alleviating the problem. The main part of the paper is our description of the two user studies we conducted and the results we gained from them. We close the paper with our conclusions.

II. RELATED WORK

Improving the scalability of virtual worlds with high object density has long been an important issue in the quest to increase the usefulness of 3D environments, and many techniques attacking this problem exist in the literature today. In the following text, we will try to describe the major ones.

2.1 Multiple Views

A popular approach to handle object congestion is to introduce additional views that present more information about the 3D environment. This is often done through a combination of overview and detail views. Baldonado et al. [4] present eight general guidelines for designing multiple view visualizations and give examples of existing applications. The Worlds-in-Miniature technique [5] uses an additional miniature 3D map of the environment, allowing the user to discover objects that would otherwise be occluded. The user can also directly interact with the WIM. Worldlets [6] are 3D thumbnails providing both global overview maps of a virtual world as well as local views optimized for detail. They are typically arranged into collections, serving as bookmarks into the3Dworld.Yet another multiple-view technique is bird's eye views [7], which combine overhead maps with the standard 3D view of the world.

2.2 Space Distortion

Space distortion can be used to manage object congestion in both 2D and 3D, and is typically done by providing one or several foci that serve as the center of attention, and a surrounding context, all integrated into the same view. Generalized fisheye views [1] pioneered and formalized this concept of focus+context views, and many variations on the theme exist. The Perspective Wall [8] uses perspective foreshortening to visualize linear information on a 3D surface, the Table Lens [9] allows for spreadsheet-like tabular visualization, and the Hyperbolic Tree Browser [10] represents large hierarchies using hyperbolic geometry.

A related approach distorts view space instead of object space; the view can be animated between perspective and parallel projection to facilitate object discovery [11], or multiple viewpoints can combined into one using non-linear projection [12, 13]. Singh and Balakrishnan [14] explore non-linear projection further by introducing fisheye, sticky and mosaic cameras that make use of previous exploration to distort the camera space in response to the user's interests.

2.3 Direct Manipulation

Another class of techniques for disambiguating between objects in high-density virtual worlds is invasive in nature, allowing the user to manipulate the objects in the environment directly in order to make sense of it. The EdgeLens [15] is a method intended for selective reduction of edge congestion in 2D graphs and operates by means of a probe-like lens that separates edges that would otherwise overlap each other or even hide graph nodes. The 3D explosion probe presented by Sonnet et al. [16] can be used to separate visual elements in a 3D scene temporarily to create interactive exploding diagrams. In contrast, the BalloonProbe [2] is closely related to both the EdgeLens and the explosion probe, and provides an inflatable force field controlled by the user that can be used in areas of locally high object congestion.

More complex methods for direct manipulation of objects in virtual worlds exist; representative of these is Selective Dynamic Manipulation [17]. SDM is a suite of 2D and 3D techniques that allow for complex manipulation, comparison and disambiguation. All techniques operate on a currently selected object set; object sets can be freely created, modified and destroyed by selecting individual objects. Visual object properties for a whole set can then be modified by using special object handles attached to each object. This allows a user to, for example, scale up a subset of the objects in a visualization to study their relative sizes without being distracted by objects outside the set.

2.4 Transparency

A recent trend in both 3D virtual environments as well as 3D games is to make use of transparency to expose hidden content. Chittaro and Scagnetto [18] investigate the merits of this practice and conclude that see-through surfaces seem to be more efficient than normal 3D navigation, although not as efficient as bird's eye (overhead) views. Diepstraten et al. introduce view-dependent transparency [19] where occluding surfaces are made semi-transparent to allow hidden objects to shine through. In another work, they instead cut holes in intervening geometry to expose the concealed objects [20]. Coffin and Höllerer [21] present a similar technique with active interaction where the user is controlling a CSG volume that is dynamically subtracted from the surrounding world geometry. A related approach is Viola and Gröller's work on importance driven volume rendering (IDVR) [22]; here, all 3D elements are assigned a value governing its relative importance, and the final image is a blending of all of the elements with corresponding transparency.

III. PRELIMINARIES

As can be seen from the previous section, there is a wealth of available techniques to use for comparison. Given the large scope of the tasks involved, it simply was not possible to select and implement more than a handful of these for our study. Therefore, we had to make a selection.

In order to ensure fair comparison between the techniques, we chose our sample from interactive direct manipulation techniques controlled by a user exploring the virtual world. We only considered general-purpose visual techniques suitable for scenarios when the targets are not previously known automatic or query-based methods for filtering out distracters or identifying targets are designed for a specific task and do not lend themselves to comparison with a general-purpose technique. Similarly, we disregarded techniques requiring more than one interaction phase, i.e., selecting, filtering or grouping objects or object hierarchies prior to manipulating them. We also consider only single-view methods due to the difficulty of integrating and switching between multiple views in an immersive environment. The interaction techniques required to manage multiple views in immersive environments have not been thoroughly studied, and the effect of these techniques on our measurements cannot be predicted and distinguished from the effects of the occlusion management paradigm.

We chose generalized fisheye views and the BalloonProbe technique for the fact that they are both simple and low-level 3D interactions requiring no a priori selection or grouping of objects, and they are representative of the space distortion approach to object disambiguation in 3D. Including higher-level techniques such as SDM in the evaluation would certainly be interesting, but the comparison would not be ecologically valid since the scope and usage scenarios of the competing techniques would then be so very different.

In this section, we elaborate on the user tasks targeted and introduce the two occlusion management techniques chosen, including the two different variants we have implemented of each technique. We also classify each technique using the occlusion management taxonomy presented by Elmqvist and Tsigas [23].

3.1 Use Tasks

The main user tasks we are targeting with our evaluation are based on the ability to distinguish and identify objects in a given 3D environment. These low-level tasks are always performed in the context of a higher-level task specific to the current visualization. We select a representative set of such high-level tasks in our evaluation in order to give the test subjects a meaningful reference framework.

In this treatment, we refer to objects as being either targets or distracters, depending on whether they have any relevance to the current high-level task or not. We refer to the low-level tasks as target discovery, i.e., the process of finding the targets in a collection of objects, target access, i.e., the process of retrieving information in a target, and spatial relation, i. e., the commensuration of targets in the world with each other and their context.

3.2 Generalized Fisheye Views

Given a general data set to be displayed, a fisheye view consists of a representation of the data centered on a specific focal point in the data set with a degree-of-interest (DOI) function governing the level of detail of each data point depending on some notion of distance between the point and the focus. This is the basic concept of focus+context displays, where the focused detail area of the data set is integrated with the surrounding context in a single view.

The nature of the DOI function controls the level of detail for data points in the fisheye view and is entirely independent of the graphical representation of the data. The function usually depends on the distance between focal and query points and may be continuous, discrete, filtering, or use a semantic scale, etc. The level of detail, on the other hand, is a measure of the information shown in the visual representation.

1) Classification: Fisheye views may be employed as occlusion management techniques with different characteristics depending on the DOI function. In TABLE1, we show our classification of both the scale-based (Scale FE) versus the transparency-based (Trans FE) versions, respectively.

As can be seen from the table, fisheye occlusion management techniques are primarily designed for target access. For a transparency-based fisheye, we are able to detect even contained targets, whereas realistically only proximate targets can be reliably discovered with a scale-based version. On the other hand, depth cues are retained to a much higher degree for a scale-based implementation. Furthermore, scale-based fisheyes preserve both location and appearance of targets, whereas transparency-based ones will change both geometry as well as appearance, i.e., transparency-based fisheye will only preserve location.

2) Implementation: Our implementation of the fisheye view for a virtual 3D environment uses a continuous DOI function based on the Euclidean distance between the viewer's hand position (the focus point) and each object being rendered. We use a standard function with interest inversely proportional to distance, i.e., $i(d)=c/(d+c)$for a specific constant c.

We developed two different alternatives for the level of detail, one based on object scale and one on object transparency. For the former alternative, objects are scaled according to interest so that more interesting objects are larger than less interesting ones (see Fig. 1a and Fig. 4a). For the latter alternative, we analogously modify the overall transparency of an object as a function of the interest.

One interesting point to note is that a 3D environment already has a natural fisheye effect arising from perspective distortion. In other words, objects that are far away appear smaller than objects that are closer to the viewpoint. In some cases, it might even be beneficial to employ a DOI function directly proportional to the distance, causing distant objects to become larger in order to avoid this effect, or nearby objects to become more transparent. This is beyond the scope of this paper, however.

3.3 BalloonProbe

The BalloonProbe technique, introduced by Elmqvist [2], provides an inflatable force field probe connected to the user's 3D input device. The probe can be applied to areas of high object congestion in order to disambiguate between objects and reduce the local occlusion. See Fig. 2 for a schematically overview of the balloon distortion. The balloon itself is centered on the 3D cursor controlled by the user, who can not only toggle the inflation/deflation of the balloon to its full size, but also modify this size depending on the situation. This gives the user an intuitive and useful way of quickly separating objects in areas of locally high inter-object occlusion.

There are two main ways of using the probe:

- **Distracter removal.** Give the users the means to remove distracters from the environment or the view.
- **Target separation.** Give the users the means to separate and isolate targets in the environment.

The task governs the force field geometry that is best suited to solving it. For removing distracters, the displaced objects are not interesting and we just want to get them out of the way, so we use a wedge-shaped force field of two half planes hinged around the probe focus point. The behavior of this probe will be akin to "parting branches" (or distracters) in order to see the

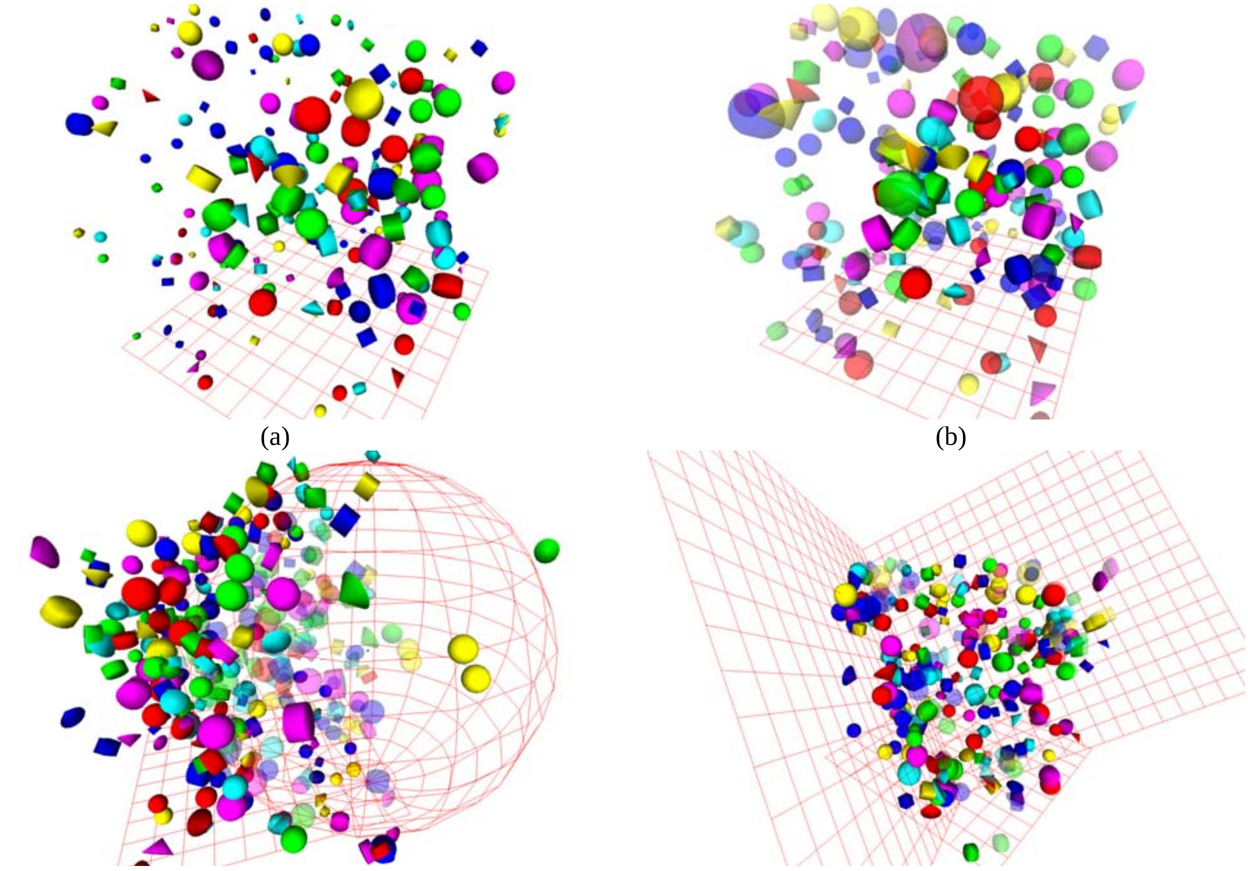

(a) (b)

(c) (d)

Fig. 1. Example screenshots of the techniques in the abstract environment. (a) Scale-based fisheye view. (b) Transparency-based fisheye view. (c) Spherical BalloonProbe. (d) Wedge-shaped BalloonProbe. See Color Plate 8.

targets.

For separating targets, on the other hand, we just want to scatter clustered targets without losing track of them, and instead use a spherical force field like in the original BalloonProbe system. The behavior is then more akin to an actual balloon inflating between targets and pushing them apart to present them for inspection.

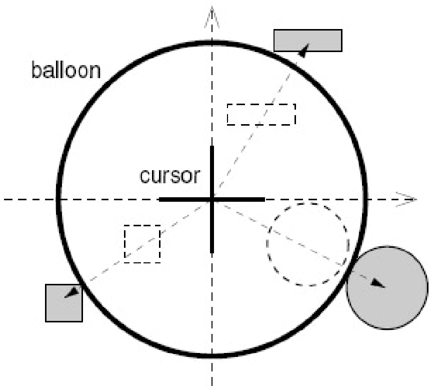

Fig. 2. 2D overview of the BalloonProbe technique with a spherical probe.

1) Classification: The BalloonProbe is an instance of the Interactive Exploder [23] design pattern in that it manages occlusion in the object space by distorting space in a direct manipulation way. TABLE 1 shows our classification of the technique using the occlusion management taxonomy.

More specifically, the BalloonProbe technique is primarily designed for accessing information in occluded objects by distorting space so that the objects are brought into view, either by separating conflicting targets or removing distracters. Due to displacing the location of targets, the technique enjoys high disambiguation strength, although fully contained targets may remain contained. By the same token, target location is obviously not preserved. The interaction is active and uses a single view.

2) Implementation: We implemented both versions of the BalloonProbe, i.e., using the wedge-shaped as well as the spherical probe geometry. The user can inflate and deflate the probe to and from full size using an input toggle button. Another input controls the size of the probe, i.e., the radius of the sphere or the angle between the half-planes. Alternatively, this could be controlled using only two buttons for directly inflating and deflating the probe to the desired size.

3.4 Software Platform

We implemented a common test platform for both Virtual Reality as well as standard desktop computers to allow for comparing the various techniques side by side under equal conditions and on potentially different hardware. The platform software is written in C++ using standard OpenGL for 3D rendering. The desktop version uses GLUT, whereas the CAVE version uses the CAVELib SDK from VRCO. Both versions provide a unified framework for implementing tasks, techniques, and scenarios independent of each other using a generic scene graph and extension mechanism.

The software platform supports a simple 3D flying navigation system using the available input devices for each hardware setup—for the CAVE, the view is controlled by the wand and the 3D-tracked shutter glasses of the user, whereas for the desktop, the view is controlled by the mouse to pan and move the distance of the cursor and by the keyboard for navigation.

The actual scenario for each trial differs depending on the condition, but all objects (both targets and distracters) are simple 3D primitives, such as spheres, cones, cylinders, and boxes. Each object is colored in a single color and uses standard smooth shading.

IV. USER STUDY (CAVE)

We designed the user study with the purpose of identifying the relative strengths and weaknesses of different occlusion management techniques. Our hypothesis was that each technique has a specific context where they perform best. The BalloonProbe technique provides local space distortion and should be effective for tasks with a local scope, whereas fisheye views provide more contexts and should accordingly be better for more global tasks. Moreover, all techniques should perform better than the base case, the standard 3D flying navigation metaphor with no specific occlusion management method. The measures of effectiveness we considered were not only the traditional time and accuracy to perform a task, but also the virtual distance traveled and the number of degrees of rotation required to complete it. The distance and rotation constitute navigational characteristics that are important for designers of immersive worlds because paradigms that require more extensive movement have the potential to cause more fatigue and dizziness to a viewer.

To formalize the above discussion, our hypotheses are:

H1: Technique is a significant effect for time, accuracy, distance, and rotation.

H2: There is interaction between technique and task with respect to our measurements.

H3: Any occlusion management technique outperforms simple navigation with no technique on all four metrics.

4.1 Subjects

We recruited 16 volunteer participants, 4 females and 12 males, all drawn from a pool of undergraduate and graduate students in computer science and engineering at the University of Arkansas at Little Rock. We estimate that participant ages ranged from 20 to 35 years of age. All participants had normal or corrected-to-normal vision, were not color blind, and were able to use the CAVE system and its input devices freely for the duration of the test sessions. They were paid $60 for their effort and competed for an additional prize of $60 for the most accurate result. Three prizes were awarded, one for each of the three tasks in the study.

4.2 Equipment

The study was conducted on a three-sided CAVE environment consisting of front, right and floor display walls. The front and right walls is Fake space reflex with PixelPipe technology that results in virtually no seam. Each wall is 10 × 8 powered by a Christie DLP projector with a resolution of 1280 × 1024 pixels. The CAVE is run by three dual-processor Intel Xeon 3.06 GHz personal computers with 3Dlabs Wildcat4 7210 graphics card, one computer for each display wall. Each display wall provides an active stereoscopic image using CrystalEyes shutterglasses, which are connected to a six-degree of freedom Flock of Birds tracker. Input is performed using a wand with at least three active buttons also tracked by the Flock of Birds. In the early stages of the study, we replaced a defective NeoWand with a Wanda.

4.3 Design

The experiment was designed with both between-subject and within-subject variables. Only one independent variable, TASK, was between-subjects to avoid making the trials excessively long. Henceforth, the tasks are termed “count”, “pattern”, and “relate”. More detailed description of the task, technique, and world are presented in the following subsections.

All other independent variables were within-subject: TECHNIQUE, DENSITY, WORLD, and TRIAL. TECHNIQUE is one of “sphere” for a spherical BalloonProbe, “wedge” for a wedge-shaped BalloonProbe, “scale” for a fisheye based on scale, “transparency” for a fisheye based on transparency, and “no technique” for no occlusion management technique ENSITY, with the levels “low” and “high”, referred to the total number of objects populating the scenarios (100 and 200 objects, respectively). WORLD has two possible values, floating 3D objects in space, termed “abstract”, and a single-level office environment with walls and doors in addition to the target and distracter objects, termed “architectural”. The last variable merely captures the fact that four randomly generated trials were performed for each condition.

The dependent variables were completion TIME, ERROR, DISTANCE traveled, and total angular ROTATION performed during each trial. The trials were randomized using a fixed pseudorandom seed, ensuring that each subject did the exact same tasks in the same order. Subjects received the TECHNIQUE variable in randomized order to counterbalance systematic effects of practice; for the other variables, the ordering was as specified above.

1) Tasks: The TASK independent variable represented the task the subject is asked to perform for a specific condition. We created three different types of tasks designed to capture many different aspects of object discovery and access, but all of them involved the lower-level task of recognizing a yellow cone as a target object. The three high-level tasks are presented below:

- **Count.** Count the number of yellow cones that appear in the virtual world.
- **Pattern**. Identify the global pattern formed by the yellow cones in the world.
- **Relate.** Find the third object “spying” on the yellow cone

and green box target objects. The targets always appear in triplets in different areas of the world, at a small but variable distance from each other and in different spatial configurations. By observing the objects in the vicinity of the two known target objects, the user is required to find the third target object. The distracters in one neighborhood are different than the distracters in another neighborhood.

The Count task was relatively simple and amounted to the subject merely counting the number of instances of a specific target in a given environment. This task was designed primarily to test discovery in the 3D environment and essentially required no global scope other than forcing the subject to remember which targets had been previously visited.

Relate entailed finding the instances of the two specified targets in the world and then isolating which third object was always present within a certain radius from the two targets (i.e., in a sense "spying" on the two targets). Here, we required a little more correlation between several different sites with the two targets in order to filter out the distracters.

Finally, the Pattern task charged subjects with finding the large-scale shape the target objects together formed in the world; the shapes were the letters C, K, R, X and Y defined as 5 × 6 grids of objects laid out on the horizontal plane (see Fig. 3). This task was designed to test global cognition of the 3D environment, and the shapes were chosen so that the individual shapes were easily confused with each other.

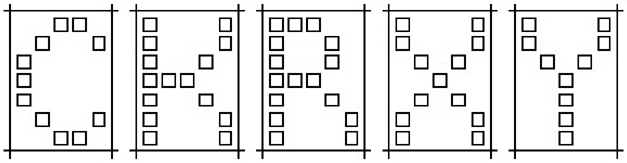

Fig. 3. The five patterns used in the pattern task.

2) Techniques: The TECHNIQUE variable represented the occlusion management technique currently in use for a specific trial. It had five different levels: "no technique" for no occlusion management (only the 3D camera controls described in Section III-D), "sphere" for a spherical BalloonProbe force field, "wedge" for a wedge-shaped BalloonProbe force field, "scale" for a scale-manipulating fisheye view, and "transparency" for a transparency-based fisheye view. Even if the probe interaction technique typically supports inflating and deflating the force field using a button, this functionality was disabled for purposes of the test and the force field was always active.

3) Worlds: The WORLD variable represented the specific type of virtual world for a specific condition. It had two levels: "abstract" for a free-space abstract3Denvironment, and "architectural" for a virtual 3D walkthrough application. These two types were chosen to represent the two basic classes of virtual worlds that are commonly used in virtual environments; the abstract 3D environment is similar to an information visualization application where abstract data lacking a natural visual mapping are represented by more or less arbitrary 3D geometry, whereas the virtual architectural walkthrough represents the class of virtual worlds that try to mimic reality in some sense. The abstract world only has a basic 3D grid at the bottom of the environment to aid in navigation (see Fig. 5a for a screenshot). The architectural scenario, on the other hand, provides a basic floor plan of a 3D building with floor, walls, ceiling and doors, all randomized (see Fig. 5b for a screenshot).

4.4 Procedure

Before starting the first session, the participants read and signed a consent form that included a brief description of the experiment. The reading took between five and ten minutes and was kept intentionally short because we considered the interaction techniques intuitive enough to require little explanation.

The participants were assigned into one of the three tasks (count, pattern, or relate) in a round-robin fashion. A random order in which they were to perform the five techniques was assigned by shuffling of cards; note that to perform the techniques in all possible orders would have required 120participants. Each technique was to be performed in a separate session.

Each participant received five answer sheets, one for each session. The sheets were kept in the lab, and participants only had access to the answer sheet while performing the corresponding session.

The first trial was performed under the supervision of a test administrator in order to make sure that the participant understood the task, navigation technique, and trial control mechanism. Participants were reminded that it was their choice whether to sit on a chair provided in the CAVE or stand up and move about the CAVE.

The five sessions, corresponding to the five TECHNIQUES, consisted of 16 trials, and the participants were encouraged to schedule them on separate days or at least with a few hours in between the sessions, because of possible fatigue and dizziness. A few people were able to perform two trials immediately following each other, but never more than two in one day. We conservatively estimated the time required to finish the test at four hours.

A trial began with a blank screen containing brief instructions for the task (for a given participant, the task did not change from trial to trial), the buttons required to initiate the trial, and the trial number. The participants started a trial by pressing a combination of buttons on the wand. The same combination also paused an ongoing trial or resumed a paused one. A paused trial produced a blank screen with a textual reminder on how to resume the trial, but the text was displayed with a different color than the color used between the trials to allow users to recognize that they had not ended the current trial. A different combination of buttons was used to end the trial and move to the beginning of the next one, a blank screen with instructions. Participants were instructed to pause the visualization whenever they needed to ask a question of the test administrator, and to end the task before writing down the answer on the session sheet.

One button was used to increase and another button to decrease the intensity of each technique. The effects of the intensity setting varied with the technique: enlarging the spherical BalloonProbe, changing the angle of the wedge probe, or making the objects in the world more transparent or larger for the fisheye techniques. There was no intensity associated

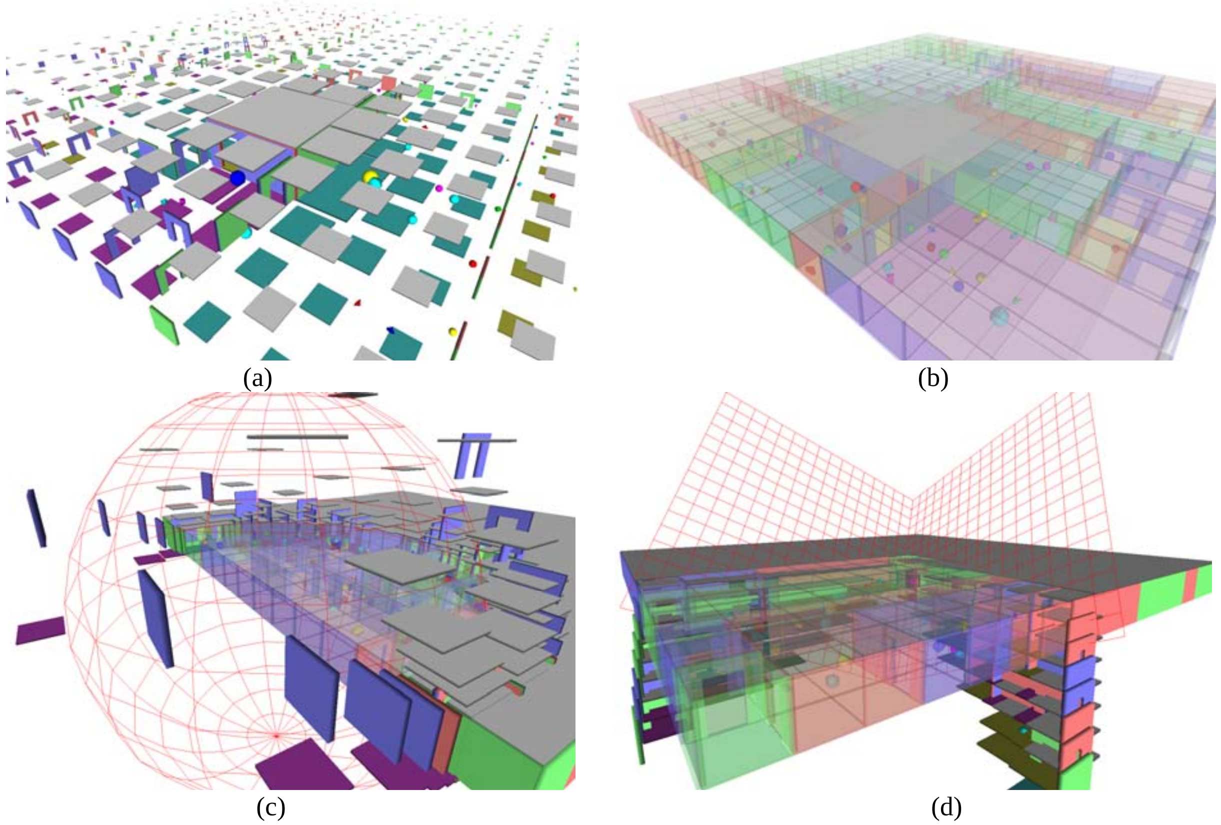
(a) (b)

(c) (d)

Fig. 4. Example screenshots of the techniques in the architectural environment. (a) Scale-based fisheye view. (b) Transparency-based fisheye view. (c) Spherical BalloonProbe. (d) Wedge-shaped BalloonProbe. See Color Plate 9.

with the “no technique” condition.

The software silently recorded, for each trial, completion times, the correct answer based on the randomly generated virtual world, the virtual distance navigated, and the total angular rotation performed by the participant. Trial timing started when the user advanced from the instructions screen and stopped when the subject ended the trial, pausing whenever the trial was paused.

The accuracy of the user answers was determined after the study when the answers written on the session sheets were checked against the correct answer recorded by the computer. The dependent variable ERROR captured how far the participant’s answer was from the correct result in each trial. For the counting task, the error was the absolute difference between the number of target objects present in the scene and the participant’s answer. For the pattern task, confusing K and R was considered half as erroneous as any other mistake because the two patterns are quite similar. For the relate task, the answer had two components—the shape of the unknown object and the color of the object—and the error was the sum of mistakes for each component.

The random generator was designed to produce the same scenes in each of the five sessions. Therefore, each TECHNIQUE went through the same 16 different worlds, which allows for a more accurate comparison of the effectiveness of the five techniques. Participants were not told about this feature of the experiment, and they thought all 80 worlds were different from each other. In fact, there were only 16 worlds on which the users operated five times, each time with a different occlusion management technique. The answers from a previous session were not available to participants in order to prevent the subjects from recognizing the repeating pattern of trials. Nonetheless, our opinion is that no participant realized that the worlds were repeating. Moreover, there was no difference between the trials experienced by one user and the ones experienced by other users except the order in which techniques were encountered.

Each participant was asked to fill out an informal posttest questionnaire that inquired whether they preferred sitting on the chair or standing up and why, whether they modified the intensity of the technique frequently, and any other thoughts they had about the study.

4.5 Results

In total there were 1280 total trials recorded, but some of the timing (5 trials), and distance, and rotation (77 trials each) data were not usable.

1)Time: The average time spent flying through the virtual environment while performing a task is about 1 minute and 50 seconds, excluding answer recording and any breaks.

Fig. 6 depicts the relative timing of the five techniques. Surprisingly, overall “no technique” outperformed all other occlusion management paradigms, partly rejecting hypothesis H3. However, the second part of that figure also shows that “no

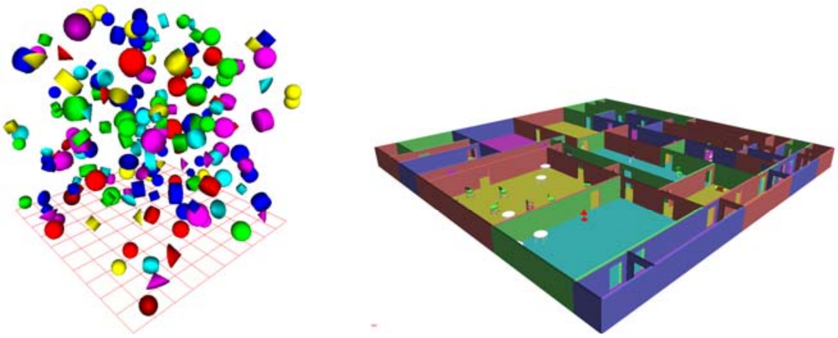
(a) (b)

Fig. 5. Example overview screenshots of the two scenarios implemented in the software platform. (a) Abstract 3D environment. (b) Architectural walkthrough application. See Color Plate 10.

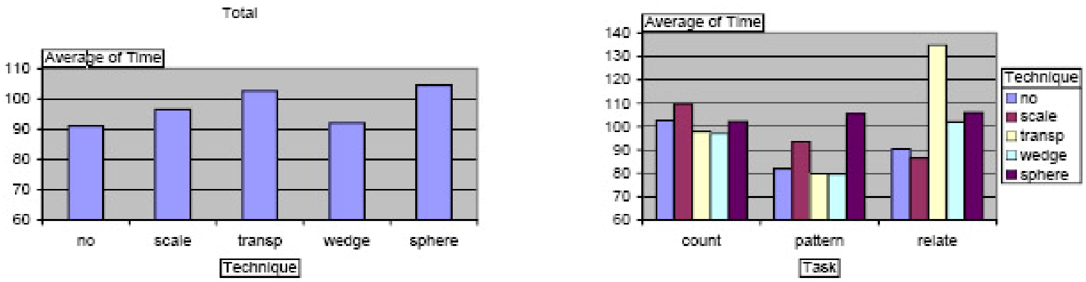

Fig. 6. Average time spent navigating the virtual world as a function of (a) technique; and (b) interaction between task and technique.

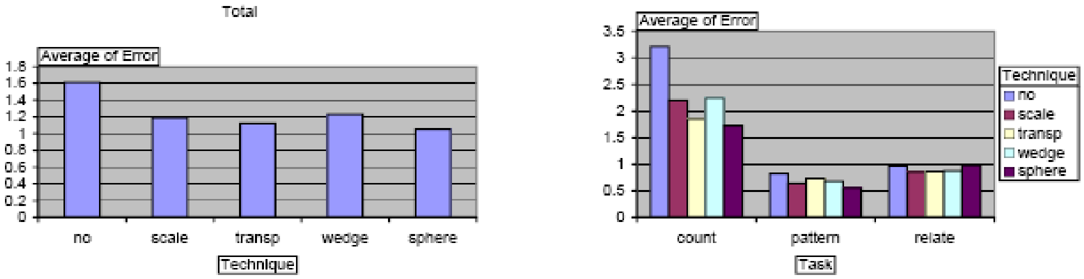

Fig. 7. Error depiction; the lower the error the more accurate the answer: (a) average per technique; and (b) interaction between task and technique.

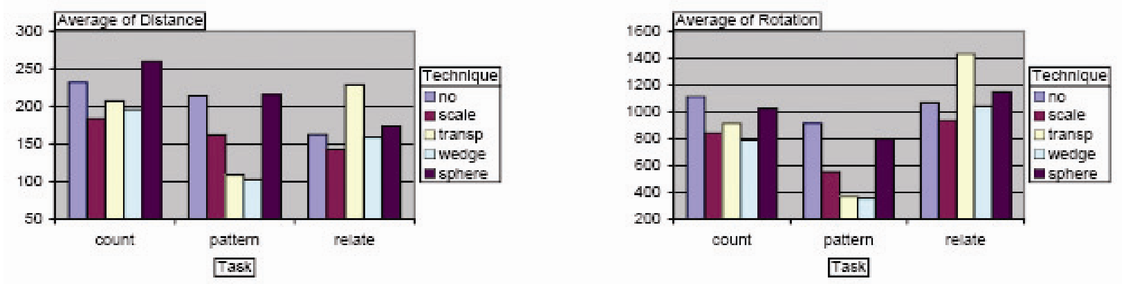

Fig. 8. Average (a) distance in feet and (b) rotation in degrees decomposed on task and technique. The two charts show a similar behavior of the techniques.

technique" was never the best in any of the individual tasks, and its over all performance is because it did not suffer any large penalty in any task. Wedge and transparency were the fastest for both "count" and "pattern" tasks, but they were slower in the "relate" task, where "scale" was the best. Scale in turn had the poorest performance in "count". The spherical BalloonProbe has a similar time performance regardless of task.

Statistical analysis of effects shows that TECHNIQUE is marginally significant (F4,52=2.51,p=.0528), while there is strong interaction between TECHNIQUE and TASK (F8,52=4.29,p=.0005). The technique does not seem to influence timing differently in different types of Worlds; no interaction was found between technique and world (F4,60=.67,p=.6178). The type of the world and density of objects are statistically significant for completion time.

2) Accuracy: The measure of error is presented graphically in Fig. 7. The average error is 2.247, 0.905, and 0.683 for the counting, relate, and pattern task respectively. Our results show that "no technique" is consistently less accurate than the occlusion management techniques. Sphere probe and transparency fisheye support better accuracy.

TECHNIQUE is a significant effect, and there is significant interaction between TECHNIQUE and TASK: (F4, 52=7.30, p<.0001) and (F8, 52=4.50, p=.0003), respectively. As for completion time, the experiment failed to find interaction between TECHNIQUE and WORLD for accuracy (F4, 60=.54, p=.7103). In addition, as for completion time, the world and density are significant factors.

3) Navigation: Navigation was measured along two dimensions, DISTANCE traveled and degrees of ROTATION. The two measurements mirror each other as shown in Fig. 8. It seems that more flying entails more rotation.

ANOVA shows that TECHNIQUE is a statistically significant factor for navigation (F4, 48=7.85, p<.0001for distance; F4, 48=6.10, p=.0005for rotation). There was strong interaction between TECHNIQUE and TASK (F8, 48=7.70, p<. 0001for distance; F8, 48=5. 50, p<. 0001 for rotation). The density of objects failed to register as a significant factor for distance (F1, 15=2.17, p=. 1613), while both density and world were significant effects for rotation.

4) Subjective Comments: A number of participants singled out BalloonProbe techniques, both the sphere and the wedge, in their comments about the study and in private conversations with the second author. A simple majority of participants preferred standing up because they felt they had more control over what they saw, and the ones that sat on the chair did so mainly because it was more comfortable. One person wrote that the chair reduced dizziness. Finally, most people said that they frequently changed the intensity of the occlusion management technique, especially in the architectural type of world (Fig. 5 (b)). From discussion with participants, it appears that the architectural worlds were viewed as more difficult by most people, especially for high object density.

4.6 Discussion

The results show a clear trade-off between speed and accuracy for the CAVE study. Fig. 6 and 7 indicate that the fastest technique ("no technique") resulted in the largest error. At the same time, spherical BalloonProbe, this was among the slowest techniques, proved to the most accurate. The implication for the design of interactive, immersive environments is that an occlusion management technique is appropriate when time is of no major concern, and mistakes incur a high cost, such as in the case of medical applications. For instances when time is important, it appears that simple navigation is more beneficial.

Another design application of these results is that a spherical BalloonProbe may be more appropriate for a wider range of tasks than any other technique, including no technique. That is because, as shown in Fig. 6 (b), the time required for different tasks incurs little variation in speed and provides consistent accuracy (Fig. 7). These attributes, consistency and predictability across various tasks, are highly valued by a designer in a general purpose system.

V. USER STUDY (WORK STATION)

The main question to be answered by the complementary study presented in this section is whether the occlusion management techniques behave in the same manner in an immersive environment as they do when using a regular mouse-and-keyboard PC. We sought to preserve as many similarities between the two studies as possible to make a comparison of the two sets of results straightforward. The main differences were in the manner in which the user navigates and controls the technique. In the description below, we focus on outlining these differences. All other parameters of the experiment can be assumed to be identical to the CAVE user study.

Note that this study may be too limited in scope on its own, but it still provides important and statistically valid insight in combination with the data gathered in the immersive experiment.

- The hypotheses listed in the previous section are applicable to this study, but in order to compare the two studies, we assume a meta-hypothesis MH1: there is interaction between technique and display type, i.e., some techniques perform better in an immersive VR device and others performed better on a desktop-based VR device.
- The participants in this study were recruited from the same student population as the CAVE study. Eight volunteers were paid $40 each with three of them also receiving a $30 prize for being the most accurate at each of the three tasks. No participants of the CAVE study also participated in the desktop study.
- The equipment used was the same the type of workstation as the CAVE computers, but the graphics card, 3Dlabs Wildcat4 7110, while having virtually the same performance as the CAVE GPUs, was not capable of genlock. A single 19" Dell CRT monitor displayed the virtual world.
- The sixteen worlds used in the PC study were generated by different random seeds than the ones used in the immersive CAVE. The same world type was presented in the same sequence.
- The intensity of a technique was changed using two keys on the keyboard. The cursor (also the focus point of each technique) was always directly in front of the user's virtual eye and could only be moved closer to or further from the viewpoint. No up-down or left-right movement of the cursor in relation to the virtual eye was possible.

The PC controls were by necessity a little more unwieldy than the CAVE due to the lack of a true 3D input device. The expressive power was similar for both platforms, however.

5.1 Results

This analysis was performed by grouping the two data sets from CAVE and PC experiments together and by introducing DISPLAY as an additional between-subject factor. The close design of the two experiments justifies this type of analysis.

The PC experiment brought an additional 640 data entries, each containing time, error rate, distance traveled, and total angle of rotation. The average time for PC was 2 minutes and 20 seconds per trial (a 27% increase from the immersive

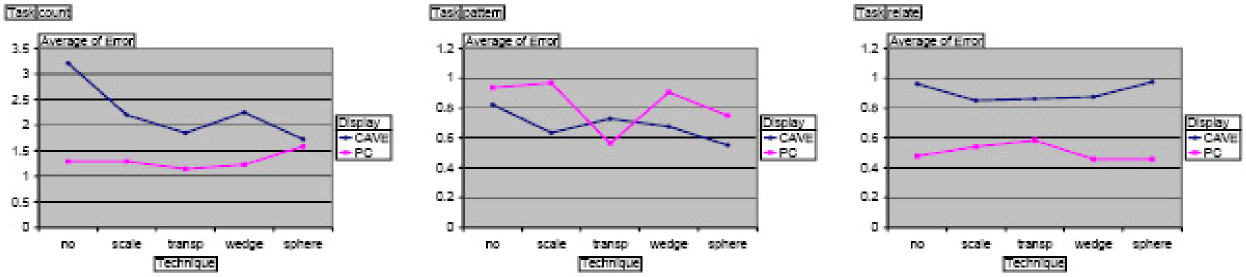

(a) (b) (c)

Fig. 9. Average error per technique and display type (a) Counting task; (b) Pattern task; (c) Relate task. There is obvious interaction between technique and display.

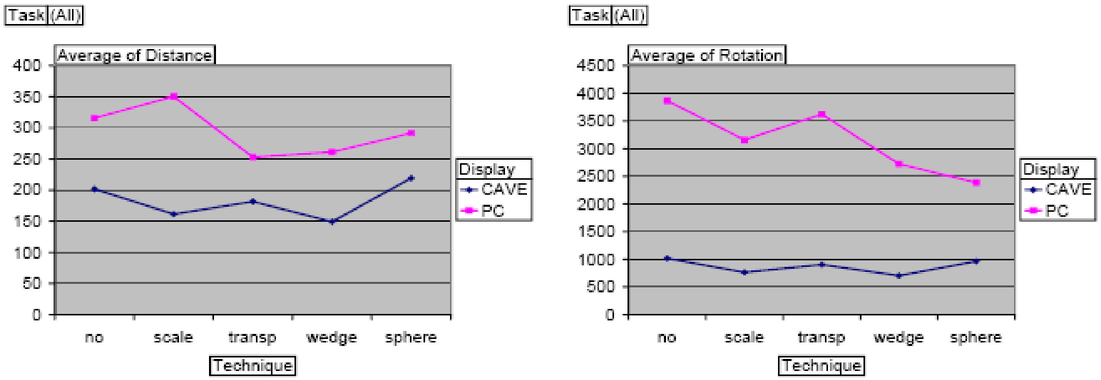

(a) (b)

Fig. 10. Navigation as it relates to interaction between technique and display type (a) Distance (b) Rotation in degrees.

condition). The error rate was lower for the counting task (1.308 PC average) and the relate task (0.704 PC average), yet higher for pattern (0.825 PC average) than in the CAVE.

There is interaction between display and technique on all measured variables except time (F4, 88=1.14, p=.3428). The statistics are: for ERROR F4, 88=2.79, p=.0310, for DISTANCE F4, 84=5.18, p=.0009, and for ROTATION F4, 84=8.83, p<.0001. Therefore, our meta-hypothesis MH holds for accuracy and navigation measures, but not for time. Fig. 9 presents the comparison of error rates for PC and CAVE for each individual task and technique. Fig. 10 shows the strong interaction between technique and display for navigation data: distance and rotation. Overall the average distance navigated on a PC is about double the distance for CAVE, and the rotation is almost three times as much on a desktop than in an immersive environment. The lack of interaction between display and technique is apparent in Fig. 11 which shows that the time spent performing a task was uniformly higher in the PC study.

Across the two experiments, TECHNIQUE is shown to be a significant factor for all dependent variables: for time (F4, 88=3.79, p=.0068), for error rate(F4, 88=2.88, p=.0272), for distance traveled (F4, 84=7.07, p<.0001), and for rotation (F4, 84=11.93, p<.0001).

Display type is a significant factor for time (F1, 18=7.58, p=.0131), distance (F1, 18=7.87, p=.0117), and rotation (F1, 18=55.01, p<.0001). Error is not influenced by display (F1, 18=2.13, p=.1613).

5.2 Discussion

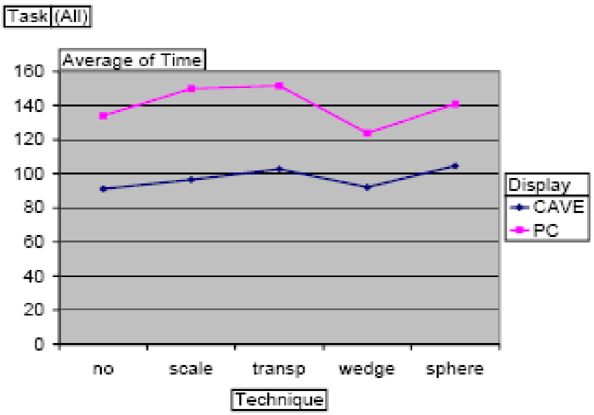

Fig. 11. Time recorded for the five techniques per each display type.

The collected data shows that people tend to spend more time in front of a desktop PC, and there is a well-known trade-off between accuracy and time. Thus, the accuracy of the tasks in the PC condition is higher. It may also be the case that the PC environment enhances accuracy, but this explanation contradicts the findings of Bayyari and Tudoreanu [24] who found that for a counting task in which the time was limited to two minutes, user accuracy was significantly better in the CAVE than on a PC. Nonetheless, tasks that require a global understanding of the world, such as the pattern task, tend to be more accurately performed in an immersive environment and in a shorter amount of time (see Fig. 9 (b) for accuracy of the pattern task and Fig. 11 for the time required in the CAVE and on a desktop computer).

Occlusion management techniques offer a better trade-off between speed and accuracy. The best example is the wedge-shaped BalloonProbe which although slightly faster than "no technique" produces more precise results. This can be observed by comparing "no technique" and "wedge" in Fig. 11 where wedge is clearly faster for PC and about as fast for CAVE, and then referencing the data in Fig. 9 where the error rate of wedge is at least equal, but often lower, than "no technique" in a similar experimental condition.

The common analysis of the two studies further confirms that the technique is an important factor that determines user performance and navigation behavior. The efficiency of different occlusion management techniques depend strongly on the type of task being performed. Furthermore, the manner in which the user navigates in the world changes not only with the technique, but also with the type of display and task (Fig. 10). A designer has to consider all these factors, and these two studies can provide a starting point for deciding what technique is best suited to a given situation.

The effectiveness of different occlusion management techniques varies between immersive VR and desktop VR. The interaction between technique and display holds for all variables except for time. An explanation may be that people get tired more easily in immersive VR and tend to shorten the time spent performing a task regardless of the actual technique used. This explanation is also supported by the finding that DISPLAY is a statistically significant factor for time, which means that an average figure for completion time can be estimated just by taking the display type into account.

The relation between immersive environments and desktop VR uncovered by the two studies presented here seems to be that user performance may not be directly influenced by environment, but navigation behavior and the time spent in the virtual world are strongly dependent on the display technology. The type of virtual environment (i.e., the type of display) cannot explain the variations in accuracy. Nonetheless, the other three dependent variables, time, distance, and rotation, are all significant influenced by display type.

VI. CONCLUSIONS

We have presented a comparative user evaluation of two different techniques for managing 3D environments with high object density. We performed the study both in an immersive VR environment as well as on a standard desktop computer. The study involved two variants of generalized fisheye views, one using object scale and the other object transparency as the degree-of-interest function, as well as two variants of the BalloonProbe technique, one involving a spherical 3D probe and the other a wedge-shaped 3D probe. Subjects were asked to perform three typical tasks in both abstract 3D environments akin to information visualizations, as well as in a more realistic architectural walkthrough application. We confirmed that the techniques have complementary properties; for example, simple navigation is fast, but the spherical probe is more accurate as well as more stable for a wide range of tasks. The experiment also shows that some techniques exhibit properties that make them desirable for situations where global relationships are important. These discoveries and the characterization of five occlusion management techniques are intended to guide designers of immersive, high-object density environments.

ACKNOWLEDGMENT

The authors would like to thank Christian Storey, Rionel Tee, and Patrick Foley for their help with administering the study. Thanks to Howard Goodell for help improving the paper. Many thanks to the anonymous test subjects who participated in the user studies.

REFERENCES

[1] G. W. Furnas. Generalized fisheye views, in Proceedings of *the ACM Conference on HumanFactors in Computer Systems*, pp. 16-23, 1986.

[2] N. Elmqvist. BalloonProbe: Reducing occlusion in 3D using interactive space distortion, in Proceedings of *the ACM Symposium on Virtual Reality Software and Technology*, pp. 134-137, 2005.

[3] N. Elmqvist and M. E. Tudoreanu. Evaluating the effectiveness of occlusion reduction techniques for 3D virtual environments, in Proceedings of *the ACM Symposium on Virtual Reality Software and Technology*, pp. 9-18, 2006.

[4] M. Q. W. Baldonado, A. Woodruff and A. Kuchinsky. Guidelines for using multiple views in information visualization, in Proceedings of *the ACM Conference on Advanced Visual Interfaces*, pp. 110-119, 2000.

[5] R. Stoakley, M. J. Conway and P. Pausch. Virtual reality on a WIM: Interactive worlds in miniature, in Proceedings of *the ACM Conference on Human Factors in Computing Systems*, pp. 265-272, 1995.

[6] T. T. Elvins, D. R. Nadeau and D. Kirsh. Worldlets-3D thumbnails for wayfinding in virtual environments, in Proceedings of *the ACM Symposium on User Interface Software and Technology*, pp. 21-30, 1997.

[7] S. Fukatsu, Y. Kitamura, T. Masaki and F. Kishino. Intuitive control of "bird's eye" overview images for navigation in an enormous virtual environment, in Proceedings of *the ACM Symposium on Virtual Reality Software and Technology*, pp. 67-76, 1998.

[8] J. D. Mackinlay, G. G. Robertson and S. K. Card. The Perspective Wall: Detail and context smoothly integrated, in Proceedings of *the ACM Conference on Human Factors in Computing Systems*, pp. 173-179, 1991.

[9] R. Rao and S. K. Card. The Table Lens: Merging graphical and symbolic representations in an interactive focus+context visualization for tabular information, in Proceedings of *the ACM Conference on Human Factors in Computing Systems*, pp. 318-322, 1994.

[10] J. Lamping and R. Rao. The Hyperbolic Browser: A focus + context technique for visualizing large hierarchies, *Journal of Visual Languages and Computing*, vol. 7, no. 1, pp. 33-35, 1996.

[11] N. Elmqvist and P. Tsigas. View projection animation for occlusion reduction, in Proceedings of *the ACM Conference on Advanced Visual Interfaces*, pp. 471-475, 2006.

[12] M. Agrawala, D. Zorin and T. Munzner. Artistic multiprojection rendering, in Proceedings of *the Eurographics Workshop on Rendering Techniques*, pp. 125-136, 2000.

[13] K. Singh. A fresh perspective, in Proceedings of *Graphics Interface*, pp. 17-24, 2002.

[14] K. Singh and R. Balakrishnan. Visualizing 3D scenes using non-linear projections and data mining of previous camera movements, in Proceedings of *AFRIGRAPH*, pp. 41-48, 2004.

[15] N. Wong, M. S. T. Carpendale and S. Greenberg. EdgeLens: An interactive method for managing edge congestion in graphs, in Proceedings of *the IEEE Symposium on Information Visualization*, pp. 51-58, 2003.

[16] H. Sonnet, M. S. T. Carpendale, and T. Strothotte. Integratingexpanding annotations with a 3D explosion probe, in Proceedings of *the ACM Conference on Advanced Visual Interfaces*, pp. 63-70, 2004.

[17] M. C. Chuah, S. F. Roth, J. Mattis and J. Kolojejchick. SDM: Selective dynamic manipulation of visualizations, in Proceedings of *the ACM Symposium on User Interface Software and Technology*, pp. 61-70, 1995.

[18] L. Chittaro and I. Scagnetto. Is semitransparencyuseful for navigating virtual environments?, in Proceedings of *the ACM Symposium on Virtual Reality Software and Technology*, pp. 159-166, 2001.

[19] J. Diepstraten, D. Weiskopf and T. Ertl. Transparency in interactive technical illustrations, *Computer Graphics Forum*, vol. 21, no. 3, pp. 317-325, 2002.

[20] Niklas Elmqvist and M. EduardTudoreanu. Interactive cutaway rendering, in Proceedings of *EUROGRAPHICS*, pp. 523-532, 2003.

[21] C. Coffin and T. Hollerer. Interactive perspective cut away views for general 3D scenes, in Proceedings of *the IEEE Symposium on 3D User Interfaces*, pp. 25-28, 2006.

[22] I. Viola, A. Kanitsar and E. Groller. Importance driven volume rendering, in Proceedings of *IEEE Visualization*, pp. 139-145, 2004.

[23] N. Elmqvist and P. Tsigas. A taxonomy of 3D occlusion management techniques, in Proceedings of *the IEEE Conference on Virtual Reality*, pp. 51-58, 2007.

[24] A. Bayyari and M. E. Tudoreanu. The impact of immersive virtual reality displays on the understanding of data visualization, in Proceedings of *the ACM Symposium on Virtual Reality Software and Technology*, pp. 368-371, 2006.

Niklas Elmqvist is a postdoctoral researcher at the InSitu Lab of INRIA Futurs located at Universit’e Paris-Sud XI in Paris, France. Previously he was a PhD student at the Department of Computer Science and Engineering at Chalmers University of Technology in Gothenburg, Sweden, where he received his PhD in December 2006. Active in the research fields of visualization and HCI in general, he specializes in topics such as 3D occlusion management, causality visualization, and general information visualization.

M. Eduard Tudoreanu is an Assistant Professor in the Department of Information Science at the University of Arkansas at Little Rock. His areas of research are information visualization and Virtual Reality. More specifically, he is interested in employing animation to uncover the underlying processes that are responsible for changes observed in various types of dynamic data.