Enhancing Fish Tank VR

Jurriaan D. Mulder, Robert van Liere
Center for Mathematics and Computer Science CWI
Amsterdam, the Netherlands
{mullie|robertl}@cwi.nl

Abstract

Fish tank VR systems provide head coupled perspective projected stereo images on a display device of limited dimensions that resides at a fixed location. Therefore, fish tank VR systems provide only a limited virtual workspace. As a result, such systems are less suited for displaying virtual worlds that extend beyond the available workspace and depth perception problems arise when objects that are (virtually) located on the edge of the workspace in between the viewer and the display screen are displayed. In this paper we present two techniques to reduce this disadvantage: cadre viewing *and* amplified head rotations*. Cadre viewing aims to eliminate the problems in depth perception for objects with negative parallax touching the screen surround. Subjective observations from an informal user study indicate a reduction of confusion in depth perception. Amplified head rotations provide a transparent navigation technique to allow users to view larger portions of the virtual world without the need for an additional input device to navigate. A user study shows it performs equally well when compared to a technique based on the use of an additional spatial input device.*

1 Introduction

The term "Fish Tank VR" was first introduced by Ware et al. [9] to characterize systems where a stereo image of a three dimensional (3D) scene is viewed on a monitor using a perspective projection coupled to the head position of the observer. Whereas Ware et al. speak about a monitor, we extend the definition by stating that fish tank VR systems are characterized by their limited physical display size that resides at a fixed location. Such display systems vary from ordinary CRT monitors to back-projection systems of sizes up to 1 by 1.5 meters. The viewer sits or stands in front of the display system typically wearing stereo shutter glasses and holding some type of (spatial) input device. A tracking device registers the viewer's head position and orientation such that the (stereo) images can be presented in a correct perspective.

Advantages of fish tank VR systems include:

- Simplicity. Compared to large surround-screen projection based systems like the CAVE [2], fish tank VR systems require less components, are easier to set up and calibrate, and can more easily be transported to a different location.
- High Resolution. Up till this date, fish tank VR systems can provide more pixels per degree in the viewer's field of view than is possible in head mounted display (HMD) systems and surround-screen projection based systems. In addition, the light intensity per degree can be higher.
- Low Costs. Because fish tank VR systems require less specialized hardware components, they are less expensive in both initial costs and maintenance.
- Versatility. Smaller fish tank VR systems can be used as ordinary desk top computer systems, larger ones as display media for presentations.

However, due to the limited display size and fixed location, a major disadvantage of fish tank VR systems is their limited workspace. This aspect affects the use of fish tank VR systems in two ways: Firstly, it limits the use of fish tank VR systems for walk-through type of VR applications. Users can only see what is in front of them, and need some type of a navigation method to be able to see what is to their left, right, top, or bottom. Secondly, it limits the use of fish tank VR systems for displaying objects in front of the display. Objects displayed with negative parallax on an edge of the viewing volume give rise to confusion in the depth perception of these objects. While the object is supposed to be located in front of the display screen, it is often perceived as if it resides behind the screen.

In this paper, we introduce two techniques to reduce the forementioned problems and thereby enhance fish tank VR systems. The first technique we call *cadre viewing*. It reduces the problem of conflicting depth cues for objects in front and on the edge of the display screen. By placing a virtual cadre in between the viewer and the display screen the depth cues for nearby objects are corrected. The second technique we call *amplified head rotations*. It makes fish tank VR systems more suitable for working with 3D scenes that extend beyond the display space provided by the system. It provides a transparent navigation method based on the viewer's head rotations. The tracked head rotations of the viewer are "amplified", such that he can look around in the virtual world.

In the next section we review previous work on fish tank VR systems, and some techniques that have been developed to improve them. In section 3 we describe the cadre viewing technique, and discuss its merits. In section 4 we describe our navigation technique based on the viewer's amplified head rotations, and compare it to a more conventional method of navigating with an additional spatial input device. Section 5 gives some conclusions and indicates areas for future research.

2 Related Work

In [9], Ware et al. show some merits of fish tank VR over traditional interactive 3D graphics. Through user tests they found that head coupled perspective rendering was a very important cue in obtaining a strong impression of a 3D scene, more important than stereo graphics. Combined stereo graphics and head coupling provided the best enhancement. In [3], Deering shows how to obtain accurate sub-centimeter virtual to physical mappings for the 3D graphics in fish tank VR systems. Aspects like predictive head tracking, dynamic optical location of the viewer's eye points, physically accurate stereo perspective viewing transformations, and corrections for refractive and curvature distortions of glass CRTs are addressed. However, the perceptual issues concerning the problem of confusing depth perception for objects with negative parallax on the edge of the display screen are not addressed.

Several techniques have also been developed that could be used to increase the usability of fish tank VR systems for larger virtual worlds. Slater et al. [7] developed an alternative viewing model for computer graphics that includes visual cues that should stimulate peripheral vision. Robertson et al. [6] have experimented with *peripheral lenses* as a navigation aid in desk-top VR systems. Their formal studies however, did not show an improvement over standard viewing in a user study concerning a search task in a virtual environment. Both these techniques have the disadvantage that portions of the display screen space, limited as it is, are being used for a distorted and unrealistic display of the peripheral.

A number of navigation techniques have been developed to navigate in virtual worlds or interactive 3D graphics applications. Ware et al. for instance, compare three different metaphors for exploration and virtual camera control in virtual environments [10]. The metaphors "eyeball in hand", "scene in hand", and "flying vehicle control" each have their own benefits and disadvantages depending on the type of task to be performed. Each of these metaphores however, is based on the use of an additional spatial input device and they are not transparent to the user.

3 Cadre Viewing

The confusion in depth perception for objects with negative parallax on the edge of the display screen is caused by what is often referred to as the "screen surround" [5], i.e., the horizontal and vertical edges of the display screen. When an object touches the screen surround, there will be a conflict in depth cues that many people find objectionable; they describe the effect as "blurry" or "out of focus" and often the object is perceived as located in the space behind the display screen instead of in front of the screen.

There are two causes for the paradoxical depth perception when an object with negative parallax touches the screen surround. Firstly, people often perceive the display as a window through which they look into the virtual world. Because the object is cut off by the window, people tend to think it has to be behind the window. The negative stereoscopic parallax however, indicates the object is in front of the window. Secondly, the display screen cuts off the object which causes unrealistic left and right images, see figure 1. The object on the left is located in front of the display screen and it is cut off due to the limited display size. Therefore, the left eye sees a larger portion of the object than the right eye does. This is a condition not encountered in real-world viewing.

Cadre viewing corrects these perception problems by placing a virtual cadre between the viewer and the display screen, as shown in figure 2. The effect of the cadre is that the viewer no longer looks at the 3D scene through the window defined by the display screen, but through a window defined by the virtual cadre. The cadre is rendered in the virtual world, in a color that closely matches the screen surround. Objects that (partially) reside behind the cadre will be obstructed and therefore clipped correctly.

The geometry of the cadre at distance d from the display screen is computed such that it allows for the maximum field of view for which objects behind the cadre are clipped correctly. For each frame to be rendered, the cadre is recomputed according to the viewer's eye positions, such that the viewer always has the maximum field of view from his current viewing position. Hence, if the interocular line is parallel to one of the display screen edges, the virtual cadre will not be visible along these parallel edges.

Placing the cadre between the viewer and the display

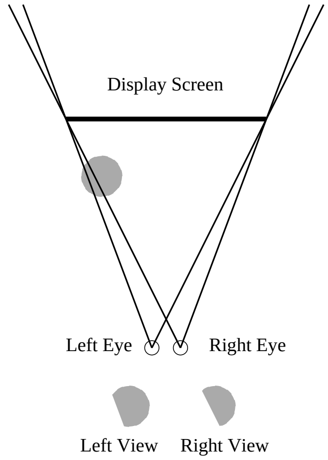

Figure 1. Unrealistic left and right images due to the limited display size.

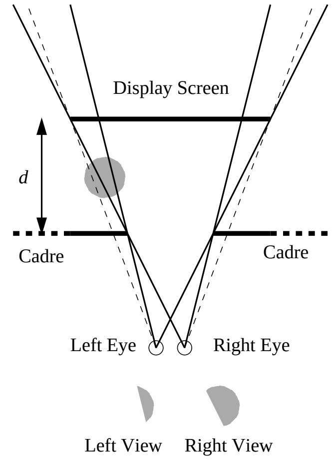

Figure 2. Corrected left and right images by insertion of a cadre.

screen reduces the field of view for each eye, as can be seen in figure 2. The amount of reduction depends on a number of factors such as the interocular distance of the viewer, the distance from the viewer to the screen and the viewer's position with respect to the position of the display's edges, the distance from the cadre to the screen, and the size of the display screen. As an example, consider a display of 1.5 meters wide where a viewer with an interocular distance of 0.065 meters sits in front of at a 1 meter distance facing the center of the display. The (horizontal) field of view for each eye is then about 73.7 degrees when no cadre is placed. When a cadre is placed between the user and the display at a distance of 0.2 meters away from the display, the field of view for each eye is reduced by 0.6 degrees to about 73.1 degrees. Placing the cadre at 0.4 and 0.6 meters from the display reduces the field of view with respectively 1.7 and 3.9 degrees to respectively 72.0 and 69.8 degrees.

Color plates 1 and 2 illustrate the viewing of a scene in stereo with and without a cadre.

3.1 Evaluation

We have performed an informal preliminary user study to investigate the effect of cadre viewing. Our fish tank VR system consists of a vertical Electrohome Retrographics projection screen with an Electrohome Marquee 8500 projector, a Polhemus Fastrack tracking system for head tracking and the tracking of a Virtual Presence Space Stick spatial input device, Crystal Eyes shutter glasses for stereo viewing, and a two processor SGI Onyx2 with InfiniteReality2 graphics board for the rendering. The projection screen is 1.025 meters high and 1.365 meters wide and its center resides 0.95 meters above the ground. Cadre viewing was implemented in our software architecture for portable VR applications, PVR [4].

3.1.1 Set-Up

Subjects were seated in front of the projection screen, such that their eyes were at approximately 1.45 meters above the ground and they were facing the center of the screen from about 1 meter away. The 3D scene consisted of two red spheres with a radius of 0.15 meters. One located in the center of the screen, and one on the left or right edge. The background was colored black. The center sphere could be moved towards and away from the viewer by control of the joystick on top of the wand (the Space Stick).

The subject's task was to position the center sphere at the same depth plane as the sphere on the edge. The depth position of the sphere on the edge was either 0.2 meters behind the display screen, on the display screen, or 0.2 meters in front of the display screen. The height position of the sphere on the edge was either 0.2 meters above or below the horizontal center line of the display screen. The left and right position of the sphere on the edge was chosen such that the center of the sphere was on the edge of the display screen when viewed from 1 meter facing the center of the display screen. The initial depth position of the center sphere also varied between 0.2 in front, exactly on, and 0.2 meters behind the display screen, but it always differed from the depth position of the sphere on the edge. The height position of the center sphere was kept the same as the sphere on the edge, i.e., 0.2 meters above or below the horizontal center line. The left-right position of the center sphere was kept at the vertical center line of the display. When cadre viewing was enabled, the cadre was located 0.4 meters in front of the display screen. The set-up is schematically depicted in figures 3 and 4.

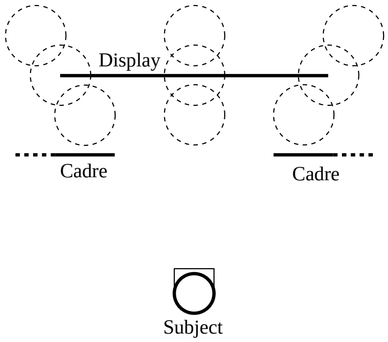

Figure 3. Schematic top view of the sphere positions in the cadre viewing test.

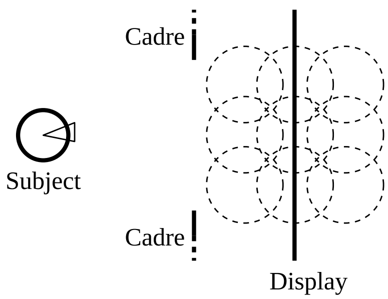

Figure 4. Schematic side view of the sphere positions in the cadre viewing test.

In total, 12 different position combinations were formed, and the subjects were asked to complete the task for each of these positions twice; one series without the cadre, and one series with cadre viewing. Half the subjects performed the series with the cadre enabled first, the other half without cadre viewing first. Before each series, 4 additional practice positions were performed. These positions were not included in the results. Furthermore, the subjects were not informed in advance nor during the experiments about the different techniques (with or without cadre) that were used to display the 3D scene nor whether they performed practice or registered tasks.

3.1.2 Results

For six subjects the results were logged (all males, age range from 26 to 33, all had normal vision or corrected to normal with contact lenses). Figure 5 shows *box and whisker* plots [8] for the differences between the depth positions at which the subjects set the center sphere and the depth positions of the sphere on the edge. A box and whisker plot shows a central box from the lower quartile to the upper quartile (i.e., the box covers the center half of the data), the mean (a +), and the median (the horizontal line inside a box). Points more than 1.5 times the interquartile range above or below the box (i.e., outside points) are depicted with a □. Points more than 3 times the interquartile range above or below the box (i.e., far outside points) are depicted as a □ with a + inside. The vertical lines extending above and below the box are called whiskers. They extend to the largest and smallest values which are not classified as outside points.

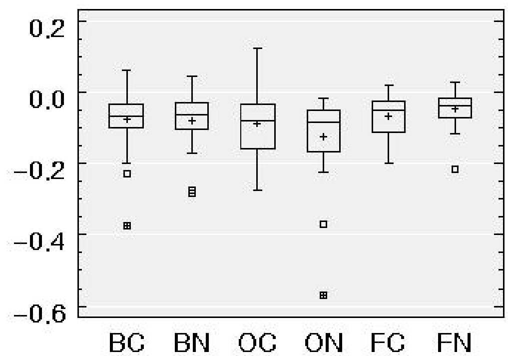

Figure 5. Box and whisker plots for the error rates obtained in the cadre viewing tests. The first letter declares the position of the sphere on the edge: behind (B), on (O), or in front (F) of the display. The second letter denotes whether cadre viewing is enabled (C), or disabled (N). The plots denote the difference in meters between the obtained and the correct positions.

The plots show the error rates with the cadre enabled and disabled. The results are split up into error values for the positions were the sphere on the edge was behind the display screen, on the display screen, and in front of the display screen. It was expected that there would be no difference between cadre viewing and traditional viewing in case the sphere on the edge was located behind the screen. A positive effect on the results was expected for the cases where the sphere on the edge was located in front of the display screen. However, as can be seen in figure 5, no significant differences where obtained in either of the three cases. In fact, the obtained error rates are pretty high for all cases. Apparently, the task was quite difficult, and the subjects were not able to perform the tests with reasonable results.

However, even though the results are not conclusive, one very interesting observation was made. Of the subjects that first performed the series with the cadre enabled, several commented on “blurry”, “incorrect”, and “difficult” depth perception for sphere positions in front of the display screen as soon as the cadre was disabled (the subjects were not notified when this happened). This observation encourages us to design and conduct a different, more formal user study to investigate the possible benefit of cadre viewing for quantitative depth perception in fish tank VR systems.

3.2 Extensions

In an attempt to further improve the depth perception in fish tank VR systems, we have extended cadre viewing to *tunnel* viewing. This is illustrated in figure 6. Here, the cadre has been extended into a tunnel. This tunnel extends from the inside corners of the cadre to as far as the far clipping plane. As a result, the tunnel limits the viewable objects to those objects that can be seen with both eyes. It clips objects that fall behind the cadre or outside the display screen area for either of both eyes. On the walls of the tunnel, a grid can be superimposed to enhance the depth perception of the tunnel and the scene. To ensure that the grid and tunnel walls are always visible, the end of the tunnel is shrunken a little. Now even if the interocular line is parallel with one of the display screen borders the tunnel and grid will still be visible. Color plate 3 shows an example view of a 3D scene with the tunnel activated.

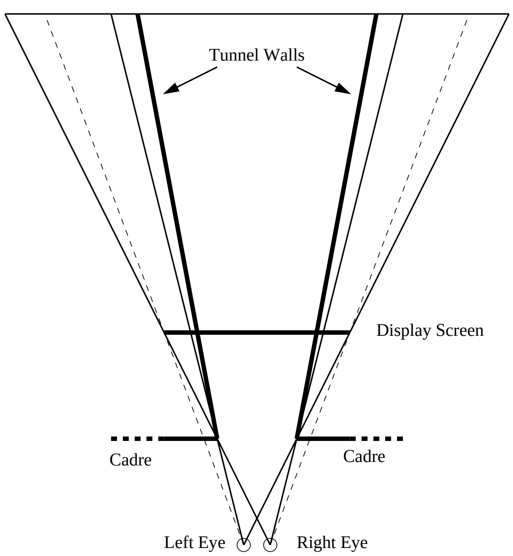

Figure 6. Extending the cadre into a tunnel.

4 Amplified Head Rotations

The second disadvantage of the small workspace of fish tank VR systems is that it makes them less suited for applications where the virtual world extends beyond the available display space. The viewer can only see what is in front of him and not easily look beside or above or below him. Traditionally, navigation methods would be provided to the viewer that enabled him to navigate in the virtual world with the use of some type of (spatial) input device, e.g. the “eye ball in hand” or “scene in hand” metaphores. Disadvantages of this approach are that the interaction is not transparent, an aspect that is considered one of the major benefits of VR over interactive 3D graphics in the first place [1], and it overloads the input device which could otherwise be used for other tasks. It is for these reasons that we developed a new navigation method for fish tank VR systems.

Our method is based on the viewer’s head rotations. Intuitively, people who want to see what is to their left or right of them turn their heads in that direction. In fish tank VR systems, the viewer can only look a limited number of degrees to his left or right because of the limited display size. In our method, the viewer’s head rotations as obtained by the tracking system are reversely applied to the virtual world: the 3D scene is rotated in the counter direction about the viewer’s head, see figure 7. In effect, we simply “amplify” or “exaggerate” the viewer’s head rotations thereby allowing the viewer to look further beside, above, and below him than the display system would allow.

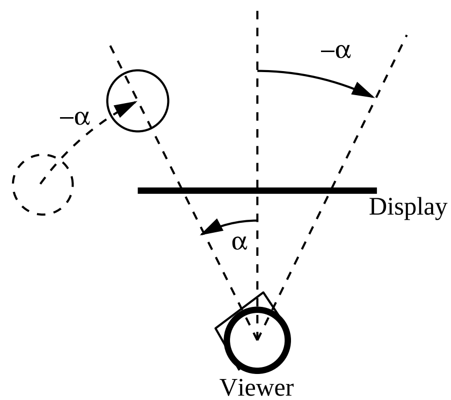

Figure 7. Amplifying head rotations. The viewer rotates his head α degrees to the left. The virtual world is rotated α to the right about the viewer’s head. As a result the sphere becomes visible.

The viewer’s head rotations can be split into panning α (looking left or right), tilting β (looking up or down), and rolling γ (rotating about one's line of sight). Only the viewer's pan and tilt rotations are amplified, as these are the rotations used to look around whereas roll is not. The amplification factors for α and β can be varied independently, depending on how much of the virtual world is to be seen and on the size of the display screen.

4.1 Evaluation

We have implemented the amplified head rotations technique in our fish tank VR environment and conducted a user experiment to evaluate the technique. As stated before, traditionally navigation is performed with an additional input device. Therefore, we have compared our technique to a similar technique where navigation is performed with the wand: instead of rotating the head to look around in the virtual world, the viewer can look around by rotating the wand. Pointing the wand to the left or right rotates the world to the right or left, and pointing up or down rotates the world down or up. Roll rotations of the wand are not taken into account.

4.1.1 Set-Up

The experiment conducted involved a "search and select" task in a virtual world. The viewer was placed in the center of a very simple wire-frame virtual room of 20 meters wide by 20 meters long, and 8 meters high. Inside the room a red sphere was present with a radius of 0.3 meters. The task of the subject was to search for the sphere, aim the ray through it, and press the wand button located under the subject's index finger (i.e., find and select the sphere). If the subject successfully selected the sphere, it would disappear and reappear at a different location inside the virtual world.

The positions of the sphere were chosen such that the sphere was not visible when looking or pointing at the center of the display screen. Subjects were forced to search for them. A total of 12 different positions were selected, left-right at -3.0, -2.0, 2.0, and 3.0, up and down at -2.0, -1.5, 1.5, or 2.0 and the depth at -3.0. The different positions are schematically shown in figures 8 and 9. Each position occurred twice in a task series, and subjects had to perform two series, one with the amplified head rotations technique and one with the wand pointing technique. Half the subjects first performed the series with the wand pointing technique, the other half first performed the amplified head rotations technique. An amplification factor of 2 was used, i.e., if the viewer would look 10 degrees to the left, the 3D scene would be rotated 10 degrees to the right about the viewer's head. Each series was preceded with four practice positions with the appropriate technique. During each series, the time was measured between the moment that the sphere appeared in the scene, and the moment the subject selected the sphere.

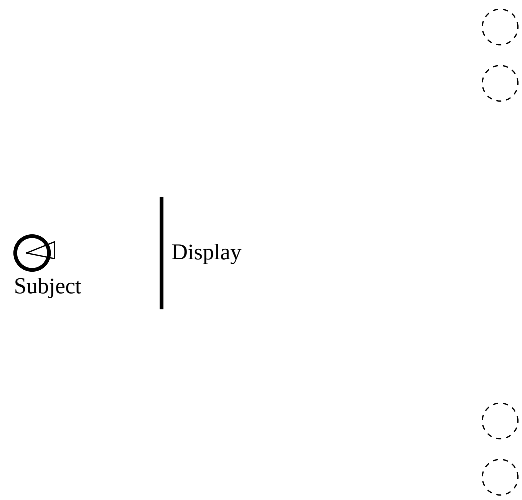

Figure 8. Schematic side view of the sphere positions in the amplified head rotations test.

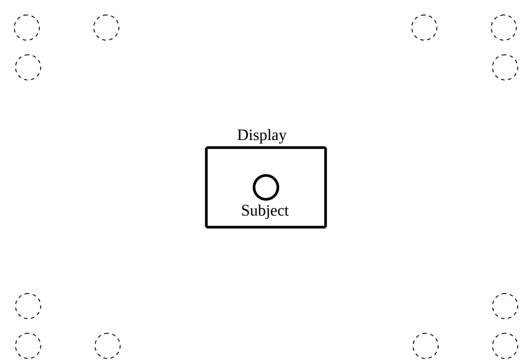

Figure 9. Schematic front view of the sphere positions in the amplified head rotations test.

4.1.2 Results

Sixteen subjects performed the tasks, 13 males, 3 females, age range from 27 to 51, who were recruited from the staff of CWI. All subjects had normal vision or corrected to normal by contact lenses. None of the subjects had ever participated in a similar experiment before. A total of 768 measurements where obtained, 384 for both the wand pointing technique and the amplified head rotation technique.

Figure 10 shows box and whisker plots of the obtained measurements. Clearly, both techniques performed equally well. Statistical analysis did not reveal any significant difference between the two methods. When the subjects were asked which method they preferred, diverse answers were obtained. There was no clear preference for either of the methods. Some subjects commented on the techniques that they found the head rotation technique to be more natural, while others stated that they thought they were able to find the sphere faster with the wand pointing technique.

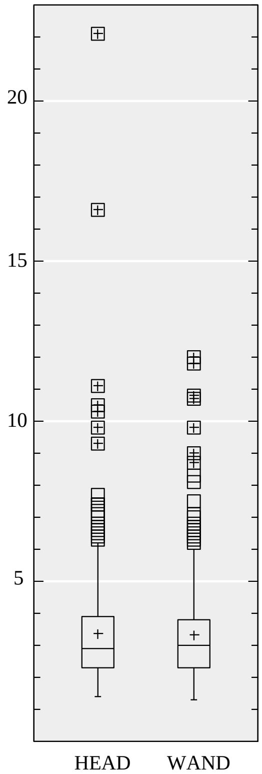

Figure 10. Box and whisker plots for the amplified head rotations technique and the wand pointing technique of the time (in seconds) it took the subjects to find and select the sphere.

5 Conclusion and Future Work

Fish tank VR systems offer a number of advantages over other VR systems such as surround-screen or head mounted display systems, in that they are versatile, low-cost and simple to construct and maintain, yet provide good quality 3D images. However, the main disadvantage of fish tank VR systems is that they only have a limited virtual workspace. In this paper, we have presented two techniques to reduce this disadvantage: cadre viewing and amplified head rotations.

Cadre viewing intends to improve the depth perception of objects with negative parallax that touch the edge of the display screen used in the system. Although no empirical evidence is yet available to support this hypothesis, subjective observations tend to indicate that viewing these objects with the cadre enabled is more comfortable than with the cadre disabled. We intend to design and perform a thorough user study to investigate this formally.

The amplified head rotations technique was designed to provide the user with a larger useful workspace. It provides a transparent interaction technique to inspect virtual worlds that extend beyond the viewing volume of the fish tank VR system, yet it does not rely on additional input devices such that those can be used to perform other manipulative tasks in the virtual world. An interesting area for future research would be to investigate to what extend fish tank VR systems equipped with techniques such as the amplified head rotations technique can concur with full-surround VR systems like the CAVE.

Acknowledgements

The authors would like to thank all the subjects who voluntarily participated in the user studies.

References

[1] C. Cruz-Neira. Making VR a useful technology: From immersive displays to applications. Invited talk at the 1997 EURO-VR Mini Conference, Amsterdam, the Netherlands, November 10–11 1997.

[2] C. Cruz-Neira, D.J. Sandin, and T.A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 135–142, 1993.

[3] M. Deering. High resolution virtual reality. In E.E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 195–202, 1992.

[4] R. van Liere and J.D. Mulder. PVR - an architecture for portable vr applications. In M. Gervautz, A. Hildebrand, and D. Schmalstieg, editors, *Virtual Environments '99, Proceedings of the Virtual Environments Conference & 5th Eurographics Workshop*, pages 125–135. Springer Verlag, 1999.

[5] D.F. MacAllister, editor. *Stereo Computer Graphics and Other True 3D Technologies*. Princeton University Press, 1993.

[6] G. Robertson, M. Czerwinski, and M. van Dantzich. Immersion in dektop virtual reality. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '97)*, pages 11–19, 1997.

[7] M. Slater and M. Usoh. Simulating peripheral vision in immersive virtual environments. *Computers & Graphics*, 17(6):643–653, 1993.

[8] J.W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, 1977.

[9] C. Ware, K. Arthur, and K.S. Booth. Fisk tank virtual reality. In S. Ashlund, K. Mullet, A. Henderson, E. Hollnagel, and T. White, editors, *INTERCHI '93 Conference Proceedings*, pages 37–42, 1993.

[10] C. Ware and S. Osborne. Exploration and virtual camera control in virtual three dimensional environments. In R. Riesenfeld and C. Sequin, editors, *Computer Graphics (1990 Symposium on Interactive 3D Graphics)*, pages 175–183, 1990.

Color Plates

Color Plate 1. Viewing a 3D scene without the cadre.

Color Plate 2. Viewing a 3D scene with the cadre.

Color Plate 3. Example of viewing a 3D scene with the cadre tunnel.