

The NAVE

Design and Implementation of a Non-Expensive Automatic Virtual Environment

Carlos Jensen ¹

carlosj@cc.gatech.edu

Jarrell Pair ¹

jarrell@acm.org

Julian Flores ²

eljulian@usc.es

Larry Hodges ¹

hodges@cc.gatech.edu

¹ College of Computing
Georgia Institute of Technology
Atlanta, GA 30313-0280, USA
+1-404-894-3152

² Instituto de Investigaciones Tecnológicas
Universidad de Santiago de Compostela
15706-Santiago de Compostela, Spain.
+34-981-563100

Abstract

This paper describes the NAVE, an affordable, immersive stereoscopic virtual reality display. The goal of the NAVE is to make key features of the CAVE available to a larger audience and introduce new and powerful features of its own. This paper describes the NAVE in detail, and offers diagrams and component information to allow others to build similar systems.

1. Introduction

As demonstrated by the popularity of the CAVE [1] and products such as the Virtual Workbench [3], there is great interest in projected stereoscopic environments as alternatives to head-mounted displays. The primary obstacle to widespread adoption has been their high cost. The goal of the NAVE project was to design a low-cost, PC-driven, multi-screen, multi-user, stereoscopic, multi-sensory virtual environment with many of the desirable elements of the CAVE at a fraction of its cost. The NAVE was built at a total cost of less than \$60,000.

The NAVE is a three-screen environment. Each screen is 2.4m wide and 1.8m high. The screens are positioned at 120° angles to each other, producing a three-sided display area 4.8m wide and approximately 2.1m deep.

The user is seated in a Thunderseat positioned at the center of the semi-circle formed by the three screens. Imagery for each screen is generated on a 500MHz Pentium III PC and back-projected in stereo. To

experience the stereoscopic effects, the user wears inexpensive, lightweight polarized glasses. A fourth PC coordinates the three screen-rendering machines and provides directional sound for the NAVE. Software support for the NAVE is based on the Simple Virtual Environments (SVE) Toolkit [2].




Figure 1: Computer-generated model of the NAVE

We chose the name NAVE for three reasons. In the tradition of the CAVE, the name is a recursive acronym (NAVE Automatic Virtual Environment). The name is also an acronym for its design goal: Non-expensive Automatic Virtual Environment. Finally, the name has an architectural context, a nave is the central part of a cruciform church building, appropriate for our first application, Santiago 2000.

Santiago de Compostela (Spain) has since the early middle ages been one of the most important pilgrimage sites in the Christian world. It is said to house the remains

of the Apostle Saint James, brought from Palestine after his death in 42 A.D. Since 1120 A.D., Santiago has enjoyed special privilege, first granted by Pope Alexander III in a Papal Bull (Bula Regis Aeterna).

This Papal Bull conferred the Grace of the Jubilee, that is the remission of all sins, even of those whose remission can only be conferred by the Pope himself. To obtain the Jubilee it is sufficient to visit the Cathedral during any Holy Year (when the 25th of July falls on a Sunday) and say a prayer for the intentions of the Pope. You must also have made confession and received Communion within the previous fifteen days or must do so within the next fifteen days. The last Holy Year of the millennium is 1999.



Figure 2: Santiago 2000 on the NAVE

Santiago has additional reason to celebrate, it has been chosen as one of nine cultural capitals for the European Union in the year 2000. To celebrate these two events, along with the general euphoria of the new millennium, the Santiago 2000 project recreates the Plaza de Obradoiros and the areas surrounding the Cathedral.

The environment is a detailed model of the city center surrounding the historic Cathedral. As users stroll through the virtual city, they are treated to the sights and sounds of Santiago, including church bells, bagpipes and singing troubadours. Santiago 2000 demonstrates that high-quality, stereoscopic virtual reality is possible on off-the-shelf PC's.

The NAVE and the Santiago 2000 projects are the result of a collaboration between the Virtual Environments group of the Georgia Institute of Technologies and the Technological Research Institute of the University of Santiago de Compostela, Spain. Santiago 2000 was first demonstrated at the VREX booth of the IEEE VR conference in March 1999. The NAVE was first demonstrated at the ACM Symposium on Interactive 3D Graphics in April 1999.

2. Design Criteria

The NAVE design had two major constraints imposed on it, those of budget and space. The primary objective was to make immersive multi-screen virtual environments accessible, and this meant bringing the cost down as far as possible. Our current design cost us under \$60,000 to implement, significantly less than a CAVE. We hope that the NAVE with its smaller price tag will be far more accessible than the CAVE. The second major design constraint was imposed by the lack of large lab-spaces at Georgia Tech. We were fortunate to get an 8.5x6.4m room for this project.

The NAVE was built with the Santiago 2000 project in mind, and so had additional requirements imposed. Santiago 2000 is designed like a theme-park ride, which means large volumes of naïve users. Every effort had to be made to minimize wear and tear of the environment and equipment. This excluded the use of expensive and fragile/temperamental technologies like shutter-glasses and trackers.

Within these constraints, we wanted to use the NAVE to explore ways to improve on the CAVE design. Among the aspects that we explored were alternative screen configurations. Most CAVE systems are designed as boxes with screens placed perpendicular to each other, which introduces sharp, visible edges. Ideally, the environment would be a perfectly smooth sphere, with every point equidistant to the users eyes. Given that this would require a very complex and expensive manufacturing, rendering and projection system, we decided instead to approximate a half circle with flat screens. Given our budgetary constraints and our desire for a full 180° field of view (horizontal), we designed a 3 screen environment where the screens are placed at a 120° angle of to each other. Though far from perfect, this configuration reduces the visibility of screen edges and produces an overall higher sense of immersion.

Another important design decision was to use off-the-shelf, sub-\$2,000 desktop PC's to render the virtual world. With the rapid growth of the PC graphics industry, largely driven by the games, it is now possible to use such inexpensive systems for these complex and demanding tasks. Having used both SGI and PC system for our applications, it is our experience that PC's often outperform SGI's. This is especially true for texture-rich environments like Santiago 2000. With the NAVE we hope to show that the PC platform is ready for use in VR applications.

3. System Overview

In an effort to help others to design and build their own NAVE's, we have included names of vendors and products as well as prices on components. This information is U.S. based, and may differ elsewhere. This information is provided for illustration purposes only, does not constitute an offer, and will likely be outdated by the time this of printing. Please contact the individual vendor for current information.

3.1 NAVE Layout

The NAVE was built in an 8.5 by 6.4m room, which restricted our design choices. The main part of the NAVE is the platform on which the three screens are mounted, and on which the user is seated. This platform, measuring 4.9x3.7m, was raised 46cm off the floor. This was done to gain additional throw distance for the projectors, as well as allow us to mount bass elements under the user.

The specific design of an environment such as a NAVE hinges on a large number of factors. Chief among these are the physical space available, the size and arrangement of the screens, throw distance of the projectors, and desired user environment. Therefore we cannot give a universally applicable design, but rather offer ours as an example. The information in the sub-sections should address most issues, and serve as starting points for alternative designs.

We have provided diagrams of the different elements of the NAVE. These diagrams are for illustration purposes only, and are not drawn to scale. The platform, screen, mirror and projector layout is shown in figure 3.

3.1.1 Projection system. On a traditional CAVE, the expense of the projection system is usually overshadowed by that of the computer system. Because of our decision to use off-the-shelf PC's for the NAVE, this dubious honor went to the projection system. Consequently a great deal of research went into making this system as simple an inexpensive as possible.

Though we could have made significant savings by going with a non-stereoscopic system, we felt that it was a tradeoff worth making. The total cost could be reduced by an additional \$5,000+ per screen by using mono projectors.

We decided early to use a passive, polarized light system for the NAVE. This decision was based on our wish to make as robust a system as possible. Santiago 2000 is designed as an amusement-park style environment, having to stand the wear and tear of thousands of naïve users. Consequently, the NAVE needs to be as robust, and "user-proof" as possible. By eliminating fragile and

expensive shutter-glasses, we significantly reduce our operating costs.




Figure 3: Room layout

Within the polarized light systems, we had the option to go with either a CRT or LCD based system. The CRT systems typically offer higher refresh rates and resolution (up to 1600x1280) than the LCD systems, but at a higher cost (\$40,000+). These projectors are offered by among others Barco and Electrohome, popular suppliers for CAVE, Immersadesk, and Immersive Workbench systems. Though most CRT based systems are shutter-glass systems, they are available in a passive, polarized light version.

Passive LCD projectors typically have lower resolution and refresh rates than their CRT counterparts. On the other hand, they are cheaper and far more portable than the CRT systems, making them ideal for conference demonstrations and trade-shows. Because of their cost and smaller form-factor, they were ideal for our system.

There are two basic approaches to polarized light stereoscopic imagery on LCD projectors today. The first uses two projectors (used by Barco among others), one for each eye view. Polarizing and superimposing the two images achieves stereo. Existing LCD projectors can be retrofitted with filters and used in this manner. As an added benefit, only minimal software support is needed, as the two inputs are separate. Due to the high cost of each LCD projector, these systems are typically expensive.

The second approach uses only one projector, polarizing the individual rows of pixels in order to achieve an interleaved stereo projection. This approach, used by

* Diagram not drawn to scale, for illustration purposes only.

VREX, typically has a lower cost, as only one projector is needed. This approach also eliminates alignment problems common in the dual projector approach, but it effectively halves the vertical resolution. This approach also introduces the need for software interleaving of the two images, which can be an expensive process.

The NAVE uses three VREX 2210 stereoscopic projectors, one for each screen. These projectors have a mono resolution of 1024x768 at 400 ANSI lumens. Because the NAVE is in a dark room, 400 ANSI lumens is more than sufficient for a back-projected system. When in stereo mode, the effective resolution for each eye view is 1024x384, though the combined result appears as if at a higher resolution. The VREX stereoscopic projectors run from around \$8,000 to \$15,000.

3.1.2 Polarization. Within the polarized solutions, there are two different approaches: linear polarization and circular polarization. Linear polarization is the standard for stereoscopic LCD projectors. The light from the two views is polarized at an angle. The user wears polarized glasses, similar to regular sunglasses, where the lenses separate the two images.

This approach is very simple; the lenses needed, both for the projectors and the glasses are readily available and inexpensive. Glasses range from \$10+ to the \$0.25 range for cardboard 3D theater type glasses. The later has the advantage of providing a printable surface for advertisement, logos etc.

Because the light is polarized at a set angle, this scheme does have a disadvantage. Tilting ones head from side to side destroys the stereo effect as the glasses go out of phase with the light projected. In our seated environment, this is not a problem, but may be so for other types of applications. In general, we have noticed little head-tilting in VR environments, even with HMDs.

To counter this problem, circular polarizers may be used, or a 1/4 phase retardant may be fitted onto linearly polarized projectors. This allows users to tilt their head as much as they want without interference. This method is more expensive, adding \$500 to the cost of each projector. In addition, glasses are much harder to find, and may have to be custom made. As an additional disadvantage, this method blocks out a larger amount of light from the projector, resulting in a dimmer image.

For the NAVE we decided to use linear filters. Our users are seated and therefore less likely to move. Because of the problem of motion sickness in immersive environments, people tend to move very little or slowly in VR environments. We felt that the costs of the circular solution outweighed its benefits.

3.1.3 Mirrors. Unless you have access to an extraordinary large space or use very small screens, you will need to employ some form of folded optics to achieve the necessary projector throw-distance. The larger your projected surface is, the longer your throw-distance will be. For the VREX 2210 projecting onto a 2.4x1.8m screen, the needed throw-distance is 4.4m.

Initially, we wanted to use Mylar film as our mirror material. Mylar has the advantage of being very lightweight, which is important when dealing with big mirrors. Mylar film generally preserves the polarization of light very well, and is an excellent reflective material. Unfortunately we were unable to find any vendors.

While searching for a Mylar supplier, we obtained a sample glass mirror from a local hardware store. The tile was 1/4 in. thick, rear surfaced. Though much heavier than Mylar, it showed no polarization loss with the VREX projector. We decided to give up on Mylar and got custom cut mirrors from the local hardware store. For \$300 we got a total of four mirrors, one measuring 0.46x0.46m, two measuring 1.4x2.0m, and a final mirror measuring 1.2x1.5m. The first and last pieces were for the double folded front screen optics, while the other two pieces were used for the single-folded side-screens.




Figure 4: Front Screen Double-Folded Optics *




Figure 5: Center Mirror 2 *

* Diagram not drawn to scale, for illustration purposes only

Because of space limitations, we had to use a double-folded optics scheme for the front screen. Figure 4 illustrates our layout. The design and specifications for the two front mirrors are given in figures 5 and 6. Figure 3 gives a birds-eye view of the system.




Figure 6: Center Mirror 1*

The two side-screens use a simpler single-folded design, and therefore require only one mirror. The mirror stands vertical, requiring the projector to be raised 94cm off the floor. The projector and mirror placements are shown in figure 3. Figure 7 shows the side mirror design.




Figure 7: Side Mirrors *

For the less adventurous, a company named Da-Lite makes custom folded optics systems. For a few thousand dollars per screen they will make a compact folded optic system to your specifications.

3.1.4 Screens. The NAVE uses special plastic rear projection screens that preserve the polarization of light. Our three 2.4x1.8m screens were purchased through VREX for roughly \$800 each. In Santiago, screens were obtained through a local plastics company. Sample materials were obtained and tested with the VREX projector. Once a material was selected, the rear projection screens were cut to our specifications. The company charged approximately \$300 for each screen. Though the price may seem somewhat high, the money is well spent. The difference between a good and a bad screen is significant.

A good screen should not have to reflective a surface to avoid glare and reflections. It is also important that it retains the polarization of light well, or your stereoscopic projectors will be useless. A good screen should also be fairly scratch resistant and strong to ensure a long life

The screens were mounted between two pieces of wood at the bottom and top. On the bottom edge, the screens are lightly squeezed between two pieces of wood to prevent it from shifting. On the top we carefully drilled holes through the screen, sandwiched between two thin strips of wood. The far sides of the two side screens are sandwiched between 5x10cm posts. This provides sufficient stability to prevent the screens from buckling under their own weight.



Figure 8: Mounting the Screens

The center screen is largely supported by the side-screens. The three screens are joined and secured through the use of silicon, minimizing the appearance of a seam. To provide additional stability and support, two posts are placed behind the screen (this is visible in figure 8). These two posts are placed 30cm away from the edges of the center screen to avoid blocking the projectors. By attaching the posts to the top mounts of the screen via a bridging piece they provide additional support. The bridge pieces also serve as mounts for the front speakers.

* Diagram not drawn to scale, for illustration purposes only

3.1.5 Miscellaneous Expenses. Approximately \$2,000 was spent on lumber and hardware supplies, including tools. A strobe light for lightning effects cost \$30 at a party supply store. Four fans used for wind effects cost less than \$15 each. Another \$1,500 was spent on miscellaneous video, audio, and power cables.

3.2 Computer System

3.2.1 Hardware. A 500MHz Pentium III PC running Windows 98 is responsible for each of the three screens. A fourth computer, a 450MHz Pentium II acts as the audio server and master simulation controller. The four systems are linked through an Ethernet hub, and cost less than \$2,000 each from Dell. Common features for these systems include 128MB of RAM and a 12GB 7200rpm Ultra ATA Hard drive. All these machines use the Windows 98 operating system in order to gain full DirectX support.

The lucrative PC game market has resulted in an incredible price/performance ratio for PC graphic cards. For our applications, we have found that PC's with these cards often match or outperform high-end SGI machines and expensive specialized OpenGL accelerators. Currently the NAVE uses graphic cards based on the popular NVIDIA RIVA TNT2 chipset.

A number of 32MB TNT2 based graphic cards are currently available. These cards provide excellent OpenGL and Direct3D performance coupled with a \$150 price-tag. Using these cards, Santiago 2000 runs at 15-25 fps in stereo at a resolution of 1024x768. To produce interleaved stereo images for the VREX projectors, it is important that these cards feature a 32-bit stencil buffer. Most new video cards support this feature in hardware.

For those with deeper pockets, Evans and Sutherland offers high quality graphics accelerators supporting dual screens and hardware stereo interleaving, though these cards are priced as high as \$2,000 each. These cards supposedly eliminate the need for any special software with the VREX projectors.

3D sound can be implemented using Microsoft's DirectSound SDK. Our audio application-programming interface (API) requires a sound card with four-speaker output and DirectSound3D support, which both the Diamond Monster MX300, and the Soundblaster LIVE! provide. These cards can be found for as little as \$100.

The NAVE audio environment is driven by two independent speaker systems in the master simulation controller. This machine has two sound cards; the primary system drives the directional audio system (using the

Soundblaster LIVE! Card) while the other drives the bass system (using a Diamond Monster Sound MX200 card).

3.2.2 Software. Santiago 2000 is built on the SVE Toolkit. SVE is a graphics and sound library developed by the Georgia Tech Virtual Environments Group. Built upon OpenGL, the library supports the rapid implementation of interactive 3D worlds. SVE allows applications to selectively alter, enhance, or replace components such as user interactions, animations, rendering, and input device polling. Multiple system configurations are supported through an initialization file. This file allows the modification of hardware and software options at runtime. SVE runs on both PC and SGI platforms.

A custom real-time audio API allows us to attach sounds to graphic entities. The spatial position of sounds is synchronized with the graphical representation of the object. The API is also capable of producing audio effects such as reverberation and doppler effects.



Figure 9: The Cathedral door in Santiago 2000

The Santiago 2000 system is a detailed model of the center of Santiago de Compostela. In order to optimize performance, the SVE toolkit uses an assortment of optimization and polygon reduction techniques. As a consequence, only 2000-3000 polygons are visible at any time. The system relies on the heavy use of textures, over 20MB, to generate an almost photo-realistic scene.

Various techniques such as MIP mapping and multiple textures for different distances and resolutions are used to optimize performance and the visual appearance of the model. We also had to add stereo-interlacing support for use with the VREX projector. The technique we used generates both eye views, before using the stencil buffer to ignore every other line of these views and superimposing the result. More efficient solutions can of course be found.

3.3 User Environment

The NAVE is a seated, immersive environment (see figure10). In it, two users sit in Thunderseats while having a 180° horizontal field of view of the environment. The environment is controlled with a Microsoft Force Feedback Pro joystick or a Microsoft Sidewinder Force Feedback Wheel. Through the DirectInput API (part of DirectX), these provide programmable tactile feedback while the user navigates through the environment.



Figure 10: NAVE User Environment

3.4 Audio System

Low quality speakers can significantly reduce the benefits of audio cues in a virtual environment. We found the best speaker solution in terms of size and versatility to be the Bose Acoustimass-6 system. The system includes 5 small cube speakers and a passive subwoofer. Two of the cube speakers are mounted on the top corners of the center screen. These speakers are oriented downward and toward the two seats. Two other cube speakers are mounted behind the seats at approximately ear height (see figure 10). The subwoofer is placed in a corner of the room to maximize its effectiveness. The fifth cube speaker is not used in the NAVE. The Bose Acoustimass-6 system retails for approximately \$700. Figure 11 illustrates the audio layout.

Four Pioneer SX-205 100 watt amplifiers power the NAVE audio system, two for the primary system and two for the bass system. These amplifiers cost \$100 each. The primary system connects to the Bose speakers, creating a 3D sound field. The secondary system steers audio across four bass-zones embedded in the floor.




Figure 11: Audio layout *

The bass shakers can be purchased for approximately \$150 per pair. Six bass shakers are installed in the NAVE floor, two per zone as shown in figure 11. The six bass shakers and the subwoofers mounted under the user seats are subdivided into four zones. The second audio card is dedicated to controlling these four discrete bass zones. Consequently, it is possible to convincingly create audio-tactile effects for simulating the vibration of vehicles, thunder, explosions and earthquakes

4. Future Work

In the near future we hope to implement tracking features, allowing users to move freely around the room. This will allow us to use the NAVE for more general VR applications. Along these same lines, we are also considering expanding the range of interface devices we support. We will also continue to revise plans and designs to help lower the cost, and update the system to leverage new technologies as they emerge.

One of the reasons for continued design work is the desire to develop a more portable system. This would be very desirable for seasonal exhibits, demos and conferences. Another interesting question is how to best expand the system through the use of more screens.

On the software side of things, we will continue to work with SVE to provide tighter coordination between multiple scene renderers. The current system works well, but network code could be optimized to allow greater scalability and lowered latency.

* Diagram not drawn to scale, for illustration purposes only

5. Conclusion

The NAVE succeeds in making many of the advantages of the CAVE available at a significantly lower price. Chief among these is the high sense of immersion that an environment with a large field of view provides. The added audio features serve to heighten the sense of immersion and introduce audio-tactile feedback possibilities.

The reduced price does come at a cost. Most noticeably affected is the visual quality of the 3D scene due to lower resolutions. We believe that the resolution is sufficient for most applications. On the positive side, we believe that the light and familiar passive glasses puts users at ease, and therefore better able to enjoy the environment.

The choice to make the NAVE a seated environment was a voluntary decision. We could easily extend it to include trackers, and do indeed plan to do so in the near future. Overall, having the user seated is a desirable feature as it reduces the degree of motion sickness.

The inclusion of force-feedback controls has been a popular feature. Working with DirectX has greatly simplified development. In general, our experiences with the PC platform have been overwhelmingly positive. Among the benefits are lower development time, higher availability of features and lowered costs.

The audio system itself is somewhat of a mixed blessing. Though it opens up new avenues of interaction and research, the low frequency sound causes problems. Apart from complaints from the neighbors, the vibrations move elements such as mirrors, projectors and even the platform itself. This means that adjustments and tune-ups need to be performed regularly.

The construction was done by a group of four students and a professor, and bears those marks (and so do we!). With a professional crew, a more polished and robust system could be built. On the other hand this demonstrates that such an environment can be built by non-professional labor in a matter of weeks.

Acknowledgements

Support for this work was provided in part by the Technological Research Institute of the University of Santiago de Compostela, Microsoft, and the National Science Foundation. For more information visit our website at: <http://www.gvu.gatech.edu/virtual/nave/>

The NAVE was designed and built by Carlos Jensen, Jeff Wilson, Jarrell Pair, Dave Gotz and Dr. Larry F. Hodges. Santiago 2000 was created by the Multimedia Group of the Systems Laboratory at the Technological Research Institute of the University of Santiago de Compostela. The lead 3D artist was Ramón L. Seco de Herrera. The audio system and API were designed and developed by Jarrell Pair and Jeff Wilson. The stereoscopic, networking, and control software was developed by Emma Varela Pet, José Manuel Ferro, Dr. Julian Flores, Jeff Wilson and Dave Gotz.

References

- [1] Cruz-Neira, C., D. Sandin, T. DeFanti. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. In *Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH*, New York 1993. pp.135-142
- [2] Kessler, D., R. Kooper, and L. Hodges. The Simple Virtual Environment Library User's Guide Version 2.0. *Graphics, Visualization, and Usability Center*, Georgia Institute of Technology, 1997. Available online at: <http://www.cc.gatech.edu/gvu/virtual/SVE/>
- [3] Krüger, W., Bohn, C-A., Fröhlich, B., Schüth, H., Strauss, W., and Wesche, G. The Responsive Workbench: A Virtual Work Environment. *Computer* 28, 7 (July 1995), pp. 42-47.
- [4] Sutherland, I. E. The Ultimate Display. *Proceedings of the IFIP Congress*, pp. 506-508