

# Next-Generation Tele-Immersive Devices for Desktop Trans-Oceanic Collaboration

A. Johnson, J. Leigh, T. DeFanti, D. Sandin, M. Brown, G. Dawe

Electronic Visualization Laboratory, University of Illinois at Chicago, Chicago IL, 60607, USA

## ABSTRACT

Tele-Immersion is the combination of collaborative virtual reality and audio / video teleconferencing. With a new generation of high-speed international networks and high-end virtual reality devices spread around the world, effective trans-oceanic tele-immersive collaboration is now possible. But in order to make these shared virtual environments more convenient workspaces, a new generation of desktop display technology is needed.

**Keywords:** Tele-Immersion, Virtual Reality, Trans-Oceanic, Collaboration, Display Devices, Tracking

## 1. INTRODUCTION


Tele-Immersion connects users of high-end virtual reality (VR) equipment together over high-speed high-bandwidth networks. Unlike 'traditional' collaborative VR, tele-immersion focuses on the integration of high quality audio and video into the computer generated environments. While the goal of audio and video teleconferencing is to allow distributed participants to interact as though they are in the same physical location, tele-immersion allows them to interact as though they are the same immersive virtual environment, allowing them to interact with each other and their shared environment at the same time.

This shared environment may be the design of a new car, a visualization of climatological data, or other three-dimensional environments that do not physically exist, or can not be physically visited. The participants are not just talking about a thunderstorm, they are talking about a thunderstorm while standing inside it; they are not looking at a scale model of a new car design, they are standing inside the full size engine block. By transmitting gestures as well as audio and video between the collaborators, these shared virtual environments give their users a greater sense of presence in the shared space than other collaborative mediums. By encouraging collaboration and conversation within the data, these environments may become the preferred place to work and interact even if more traditional face-to-face meetings are possible.

When the collaborators are distributed around the world these tele-immersive sessions become more challenging as this involves multiple networks, multiple time zones, and multiple cultures. Asynchronous collaboration, where the collaborators share the same virtual space at different times, is attractive in this situation as it reduces the need to arrange synchronous meeting times and allows work to be handed off from one site to another at the end of the day. For this to be effective, however, tele-immersive environments need to actively support both synchronous and asynchronous forms of collaborative work.

Our goal is not just to make these kinds of collaborations possible, but to make them convenient and routine. Tele-Immersive collaborations are taking place now in select research labs and major corporations. Making them convenient and routine will require a new generation of desktop display devices, but creating a good physical tele-immersion device involves much more than just the display. High quality audio and video capture is necessary within the physical tele-immersive environment. The system must be able to track the position and orientation of the user, and mediate the stereo imagery. Currently we are using magnetic tracking systems and infrared emitters that drive the stereo glasses. Within this 'hostile' electromagnetic environment we want to give the user a large, bright, crisp stereo display with high refresh and fast decay rates.

In Sect. 2 we discuss tele-immersion in more detail. In Sect. 3 we discuss several tele-immersive application domains and how those application domains drive our tele-immersive research. In Sect. 4 we discuss our current display hardware and plans for future devices. Finally, in Sect. 5 we discuss our plans for future work.



**Figure 1.** Remote participants in tele-immersive sessions can be seen in the shared virtual space in several ways. The user on the left is interacting with the articulated computer generated avatar of a remote participant in the CAVE while the user on the right is interacting with the live video avatar of a remote participant on the ImmersaDesk2

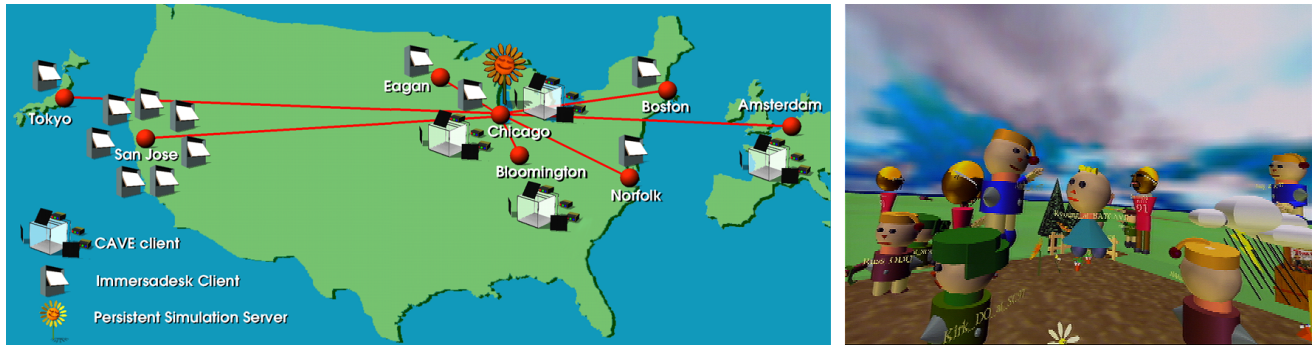
## 2. TELE-IMMERSION

In 1992 the Electronic Visualization Laboratory developed the CAVE. Now in 1998, with more than 90 CAVE and related projection-based VR environments around the world, there is a community that is eager to collaborate. Our focus is on high-end applications, helping to solve big problems across large distances, but it also involves connecting these collaborative environments to more commonly available technology.

A typical tele-immersive space will be persistent virtual environment maintained by a computer simulation that is constantly left running. The space exists and evolves over time. It may be the evolving design of a car, or the evolving simulation of climatological data. Users enter the space to check on the state of the simulated world, discuss the current situation with other collaborators in the space, make adjustments to the simulation, or leave messages for collaborators who are currently asleep on the far side of the planet. Since VR is not a common office commodity, agents may be left inside the virtual space to monitor and report back to the user via email or the WWW. When an interesting event occurs, or its time for a meeting, the user can walk over to their VR equipment and jump into the shared virtual world. Our focus is supporting high-quality interaction between small groups of tele-immersed participants. The scientists we work with suggest that rarely would there be more than five simultaneous collaborators working in the space at one time. However we also want to support a less interactive tutorial or presentation format where one expert tour guide leads tens of users through the shared virtual space.

Presence in the virtual world is typically maintained using an avatar, or a computer generated representation of a person. These avatars may be as simple as a pointer, but having physical body representations can be very helpful in aiding conversation and understanding in the virtual space as you can see where your collaborators are, and what they are looking at. Tracking simply a user's head and hand allows articulated avatars to transmit a decent amount of body language and are very useful in task oriented situations. Seeing high quality live video of a person's face can improve negotiation. Video avatars, full-motion full-body video of a user, allow very realistic looking collaborators in the space improving recognition. Articulated avatars require much lower network bandwidth than the video avatars making articulated avatars the norm today. Video avatars also require high quality cameras as the tele-immersed participants are typically in dimly lit rooms to improve the quality of the computer graphics imagery. As tele-immersive display technology improves this will be less of a problem, but currently this requires low-light CCD cameras with frame integration to build up a bright enough image of each participant. Figure 1 shows these two different kinds of avatars, with the user on the left interacting with an articulated avatar in the CAVE and the user on the right interacting with a video avatar on an ImmersaDesk2.

High quality audio is very important in maintaining the collaboration. If it is difficult to hear the other participants or if there is a large amount of delay in the audio then the collaboration quickly breaks down. Directional audio, where



**Figure 2.** Seventeen remote participants joined in a trans-global tele-immersive session centered at Supercomputing '97 in San Jose, CA. The image on the left shows the CAVE and ImmersaDesk sites that were connected in this session. These sites stretched from coast to coast across the United States and included sites in Europe and Asia. The image on the right shows a gathering of the participants during this session in the persistent space - each participant represented by a computer-generated articulated avatar.

the audio appears to come from a specific location in the space, can be very helpful in multi-user collaborations to tell exactly who is speaking as the voice appears to come from the speaker's avatar. Audio that is quieter with distance can be very helpful if there are multiple groups at different places in the space performing different tasks. Currently the telephone system offers the highest quality audio connection but does not allow directional audio. Digital networked-based audio allows directional audio but has lower quality, though that quality is rapidly improving and should become the norm within a couple years. To reduce the number of encumbrances on the user, we prefer to mount high-quality ambient microphones on the VR hardware itself, though these mikes can have problems in noisy environments.

Tele-Immersion also requires high-speed, high-bandwidth, low latency networks. Typically the computer graphics are generated locally at each site for the local VR display device and messages are passed between the sites to synchronize these worlds. The position and orientation information on each of the avatars is sent frequently (at least 20 times per second), and must have low latency, but the amount of information is small. It is also acceptable if some of these packets are lost as newer position and orientation information is only 0.05 seconds behind. Audio and video information may be streamed between the sites requiring large bandwidth and low latency. Occasionally complicated models or large data sets must be sent between the sites requiring large bandwidths and no packet loss.

Keeping these multiple data streams flowing without interruption requires sophisticated networking and software support and Kessler<sup>1</sup> presents an excellent review of recent work in this area. CAVERNsoft<sup>2</sup> is our software middleware for tele-immersion. It is the common collaborative software architecture for CAVERN, the CAVE Research Network, a global alliance of industrial and research institutions equipped with CAVEs, ImmersaDesks, and high-performance computing resources interconnected by high-speed high-bandwidth networks. It is also the one of the enabling technologies for the National Technology Grid - an NSF funded initiative to assemble computer science and computational science researchers to build and deploy powerful metacomputing systems and tools across the U.S to create prototypes for the 21st century's distributed computing infrastructure.<sup>3</sup> As part of this strategy, we have used major international conventions such as SIGGRAPH and Supercomputing over the last three years to field prototype tele-immersive devices and virtual environments in multiple domains to evaluate their effectiveness. Figure 2 shows a tele-immersive experiment conducted at Supercomputing '97<sup>4</sup> connecting 17 users on three continents.

Our tele-immersion research is not limited to CAVE-based devices. Other virtual reality equipment such as head-mounted displays and BOOMs can also participate. While tele-immersive environments will be centered around high end computing and visualization hardware, these worlds will be shared with more commonly available technology. Some activities are better done in an immersive environment while others are better done using familiar desktop applications. For example, a design reviewer in the CAVE may suggest changes to the user of a desktop CAD system, who then makes the changes and re-imports the model into the CAVE while the design review is going on. Tele-Immersion systems will typically require very high speed and low latency networks, but internet based tools such

as web browsers are also very valuable as they provide a convenient means to monitor the state of a tele-immersive virtual world.

### 3. APPLICATION DOMAINS

We have been working with VR application developers in several disciplines to see what features are important in a variety of tele-immersive workspaces. These disciplines include design, education, and scientific visualization.

#### 3.1. Design

VR can be very effective in visualizing 3D designs, and tele-immersion allows groups scattered across the planet to meet and discuss these 3D designs. However the design itself is often done using familiar desktop CAD packages, so it is very important to link desktop users into the collaboration and leverage the advantages of each platform.

General Motors developed VisualEyes, an application that allows designers to import 3D CAD models into the CAVE for quick visual inspection and design reviews. The CAVE was an ideal platform for GM as it allowed designers to see an automotive design at full size. This initial use of CAVE-based technology has generated considerable interest in other GM sites around the world, some of which are planning their own CAVE installations. This prompted GM to further extend VisualEyes to allow GM's trans-globally situated design and manufacturing teams to collaborate in remote design reviews. The goal is to allow designers to both synchronously and asynchronously access a design that persists and evolves over time. Adding in a tele-immersive component also allowed GM to link the desktop based modeling packages to the CAVE. Users in the CAVE seeing the design full size can give suggestions to a designer at a workstation who then quickly makes the required changes using a familiar interface and re-imports those changes into the CAVE for evaluation.

#### 3.2. Education

Virtual Reality offers great promise as a learning tool. Tele-Immersive learning environments allow an expert to remotely take a group of students through an immersive learning environment. This may be better than having the expert and the students in the same physical classroom as the teacher and students can share the space with the topic being discussed. Tele-Immersion also allows groups of students to collaboratively learn in an environment that reinforces heterogeneous roles or views. For evaluation studies, the ability to record these experiences through audio, video, and the recording of the virtual space itself is crucial.

The Virtual Reality in Medicine Lab at the University of Illinois at Chicago has developed the Virtual Temporal Bone,<sup>5</sup> a tele-immersive education program to allow a remotely located physician to teach medical students about the three-dimensional structure and function of the inner ear. In this environment the students and instructor may point at and rotate the ear to view it from various perspectives. They may also strip away the surrounding outer ear and temporal bone to clearly view the inner anatomy. Audio from the voice conference is used to modify the flapping of the eardrum to illustrate its function. Since the ear is shown at ten times its normal size in 3D these structures are much easier to see than at real scale or as 2D illustrations in a textbook. As with other medical applications, high resolution is a must.

The NICE project<sup>6</sup> and the Round Earth project are focusing on collaborative conceptual learning in children. Collaboration encourages conversation, and conversation serves learning by presenting each learner with a slightly different view of the subject matter. An individual is forced to enrich her own representation to understand her partner's discourse. Conversation also improves evaluation. Rather than 'thinking aloud' the participants are talking to each other. Both NICE and the Round Earth project used persistent virtual worlds. In NICE, this persistent world is a garden that has been running and evolving since the spring of 1996. The garden was designed as an environment for young children to learn about the effects of sunlight and rainfall on plants, the 'spontaneous' growth of weeds, the ability to recycle dead vegetation, and similar simple biological concepts. The Round Earth project is investigating the use of VR in learning deep conceptual ideas. VR can provide alternative cognitive starting points for learning these ideas that do not conflict with incorrect prior knowledge. In this specific research the deep idea is the sphericity of the Earth, where everyday experience tells a child that the Earth is flat. Tele-Immersion allows a pair of children to learn about a small spherical asteroid independent of their Earth-bound biases. The major hindrance in this work is that the VR equipment is not designed with small users in mind.



### 3.3. Scientific Visualization

Tele-Immersion allows scattered groups of researchers to come together in VR to visualize their multi-dimensional data. Live links to supercomputing resources can update these simulations, or the simulations can generate viewable models which are then shared.

Old Dominion University and the University of Wisconsin-Madison developed CAVE5D,<sup>7</sup> a configurable virtual reality application frame-work supported by Vis5D, a very powerful graphics library that provides visualization techniques to display multi-dimensional numerical data from atmospheric, oceanographic, and other similar models, including isosurfaces, contour slices, volume visualization, wind/trajectory vectors, and various image projection formats. CAVE6D, a tele-immersive version of CAVE5D, allows multiple users of CAVE5D to jointly visualize, discuss and interact with multi-dimensional data sets in the environment. Avatars can point at features of interest in the data set while they collaboratively modify visualization parameters such as salinity, circulation vectors, and temperature. Participants may turn any of these parameters on or off, globally (affecting everyone's view) or locally (affecting only one's own view.) This affords each participant the ability to individually customize his or her visualization. Time on the other-hand is globally shared and hence participants view all time-varying data synchronously. Participants may travel forward and backward in time as well as stop time in order to discuss what they see with their collaborators.

Motorola is working with the Institute for High Performance Computing in Singapore to collaboratively visualize impact simulations of their products. Like General Motors, Motorola has sites scattered around the world, and being able to collaboratively visualize these 3D simulations and discuss the results within the virtual space allows these distributed groups to work more closely together.

### 3.4. General Tools

Our colleagues are primarily scientists, and when scientists first see tele-immersive spaces they find the idea interesting, but often have a hard time seeing how they can make use of it. When those same scientists see their own data visualized in a tele-immersive space they get very excited. A major focus of our work is providing tools that allow domain scientists to quickly set up and use tele-immersive spaces for their own work. Part of this is deploying the tele-immersion hardware, but it is also important too have a software library that encourages new users to start getting immersed. Building upon CAVERNsoft we have built LIMBO, a generic collaborative space allowing collaborators to work in a shared space quickly.


Its very important for scientists to be able to record their work. These recordings also become critical when asynchronous collaboration is the norm as it allows the collaborators to leave recordings for their collaborators. As part of the evaluation tools for NICE, the environment recorded all activity and interaction in the space as a 3D movie, allowing a future user to come back into the CAVE and replay the actions while viewing the scene from any perspective. The virtual director<sup>8</sup> allows users to create 2D flythrough movies of their work in the 3D virtual environments, allowing this information to be sent to colleagues outside of the virtual space. V-Mail<sup>9</sup> allows users to record messages for other users in the virtual space with a combination of audio and gestures.

Another goal is to incorporate new tools for diagnostics, usage tracking, and direct feedback on performance. Initial experiments are allowing us to gather preliminary statistics which will allow us to predict the kinds of human tasks in tele-immersion that are possible across similar trans-oceanic links.<sup>10</sup> These kinds of experiments will also allow us to isolate bottlenecks in the networks and the tele-immersive systems, and make appropriate recommendations for improvements. Finally, it will allow us to prescribe an experimental suite that can be deployed at other tele-immersion sites around the world for testing and comparing the quality of their trans-oceanic links.

Figure 3 shows a selection of tele-immersive virtual environments, including the collaborative version of GM's visualEyes in design, CAVE6D in scientific visualization, the Round Earth project in learning, and the generic tele-immersive LIMBO space.

## 4. DISPLAY HARDWARE

The Electronic Visualization Laboratory at the University of Illinois at Chicago has developed several projection-based virtual reality displays to date, most notably the CAVE and the ImmersaDesk, and are currently developing several new devices.




**Figure 3.** This figure shows several current tele-immersive virtual environments. Clockwise from upper left: A full size car interior from the tele-immersive version of GM's VisualEyes, four participants exploring a simulation of the Chesapeake Bay in CAVE6D, two users sharing a generic LIMBO space, the two heterogeneous views of an avatar on the asteroid in the Round Earth project

#### 4.1. Current VR Devices

The CAVE is a multi-person room-sized 3D video and audio environment. It is a 10' by 10' by 10' room constructed of three translucent walls. 1280 x 1024 pixel stereoscopic images are rear-projected onto the walls and front projected onto the floor, and viewed with light-weight LCD stereo glasses to mediate the stereoscopic imagery. Attached to the glasses is a location sensor. As the viewer moves within the confines of the CAVE, the correct perspective and stereo projection of the environment are updated, and the user may walk around or through virtual objects. Speakers are mounted to the top four corners of the CAVE structure to provide environmental sounds and audio from the remote participants. The CAVE's room-sized structure allows for multiple users to move around freely, both physically and virtually. While as many as ten people can stand comfortably in the CAVE as well as the tracked user, five people is a reasonable maximum when doing real work. An ambient microphone mounted on top of the front wall of the CAVE allows all of the people in the CAVE to converse with remote participants. The user interacts with the environment using "the wand," a simple tracked input device containing a joystick and 3 buttons. It is used to navigate around the virtual world, and to manipulate virtual objects within that world. We are currently moving from two to four tracked position and orientation sensors as the standard configuration, allowing us to track the head, 2 hands, and a body sensor.

Even with folded optics a CAVE takes up a 30' by 20' by 15' space and requires a couple days to setup and align. The ImmersaDesk is a 1-screen version of the CAVE that looks like a drafting table with a 6' by 4' angled rear-projected screen comfortably supporting up to three viewers. When folded up it fits through a standard institutional door, and deploys onto a 6' by 8' footprint. Unlike the CAVE it requires no architectural modification to the



**Figure 4.** This figure shows EVL’s current VR devices being used in tele-immersive sessions: Eddie interacting with the NICE environment in a CAVE on the left, and Dr. Mason teaching with the Virtual Temporal Bone on an ImmersaDesk on the right (Virtual Temporal Bone image courtesy of the Virtual Reality in Medicine Laboratory)

surrounding space and can be set up in a several hours. The ImmersaDesk uses tracked glasses like the CAVE to mediate the stereo images, and the same wand to interact with the virtual world. The ImmersaDesk2 is a ruggedized version of the ImmersaDesk designed to be easily shipped via air or truck. It also features a movable display screen with variable angles of view and height adjustment, and improved access for disabled persons. Speakers mounted to the side of the desk provide audio, and as with the CAVE an ambient microphone at the top of the desk allows a group of users to converse in a tele-immersive session. The top of the ImmersaDesk also provides a convenient location to mount a camera for transmitting video between the participants.

While initially envisioned as a device to prototype CAVE applications, the ImmersaDesks have been very popular in smaller laboratories and for ‘taking VR on the road’. These road trips to conferences, schools and museums have encouraged the development of casual tele-immersion. The ImmersaDesks can be set up virtually anywhere with an internet connection and people there are joining into a tele-immersive space within a couple of hours.

Figure 4 shows users interacting with tele-immersive applications running in the CAVE and on the ImmersaDesk.

The ImmersaDesks were designed for laboratories, and take up a large amount of space in a typical office. The ImmersaDesk screen is large to present a wide field of view, but the desk structure itself is large because the rear-projection distances are significant, even with folded optics, and the desk uses a large and heavy projector. Our current work is focusing on desktop systems<sup>11</sup> that are more suited to an office. Fish-tank VR using CRT monitor technology has been used for several years, but these systems typically have very small fields of view. Increasing the field of view without filling the office with equipment means an increased reliance on flat-panel display technology. Instead of leaving your office to walk over to a CAVE or an ImmersaDesk to enter a tele-immersive space, the goal is that the surface of your desk becomes a display device, the walls become display devices, perhaps even the floor and the ceiling. In effect, your office becomes a small CAVE.

The CAVE, ImmersaDesk, and ImmersaDesk2 are commercial products sold by Pyramid Systems. Several companies, such as Panoram and VRex, offer well-designed, non-tracked displays for the office and showroom. Trimension, MechDyne, Tan, and Fakespace have products similar to the ImmersaDesk; MechDyne and Tan sell CAVE-like systems. The goal of EVL’s research is not to compete with the commercial sector, but to investigate and inspire new display and tracker technologies for tele-immersion. Given that affordable bright wall-sized high-resolution border-less displays with high refresh rates and fast decay rates do not yet exist, we are prototyping these systems using available components: existing flat panel technologies, or simulating flat panel systems with rear-projection hardware.



**Figure 5.** This figure shows the concept sketch of the ImmersaDesk3 on the left and a photograph of the current prototype of the ImmersaDesk3 on the right (concept sketch courtesy of Sam Thongrong, photograph courtesy of Dave Pape)

## 4.2. Prototyping the ImmersaDesk3

Our initial work in this area has focused on developing a prototype of the ImmersaDesk3. The major design criteria for the ImmersaDesk3 was that the display should be mountable on a conventional office desk, feature a bright wide field of view flat panel stereo display, and include head and hand tracking. For convenience, the user should be able to position the screen at any angle from horizontal to vertical, forward or back, on the desk. Figure 5 shows the concept sketch and the current prototype of the ImmersaDesk3.

Several different flat panel technologies are currently available including liquid crystal displays, Plasma panel displays, Light Emitting Diode displays, and Ferro-electric Liquid Crystal displays. Plasma panel displays were the only ones that met our needs. We chose a Fujitsu Plasmavision display panel which is relatively lightweight for its size (80 lbs. for a 42" diagonal panel), and of moderate cost (\$10,000). Unfortunately its resolution was low (800 x 600). Of greater concern was that no matter what video frequency was input, the panel displayed a 30 Hz interlaced NTSC quality image, much lower than the 96 or 120 Hz video frequencies we commonly use. This resulted in a very noticeable flickering when viewed in stereo. The red and green phosphor decay rate was also not fast enough, resulting in cross talk between the images for the left and right eye, making it very difficult to fuse the stereo images.


Adding to the problem, the panel emitted infrared radiation in the same part of the spectrum that the emitters use to trigger the stereo shutter glasses, forcing us to hard-wire the glasses. The panel also generated enough electromagnetic noise to interfere with our normal tracking equipment, forcing us to switch to a different tracking system. On the other hand the display was very bright even in well lit rooms, and the size of the panel worked well for a desktop system.

We attached the ImmersaDesk3 to a standard office desk and shipped it and an ImmersaDesk2 to SIGGRAPH '98 in Orlando. The ImmersaDesk3 was rugged enough to survive shipment to and from the convention center, and moving around the conference several times a day. It was used in several tele-immersive sessions with the ImmersaDesk2.<sup>12</sup> Due to the previously mentioned problems with generating stereo imagery the panel was mostly used in head-tracked monoscopic mode, and it worked quite well in that mode. These initial experiments with the ImmersaDesk3 suggest that the concept of a desktop mounted flat panel display is sound, though commercial display technology is not up to the task at this time.

## 4.3. New Devices

EVL currently has several other desktop tele-immersion devices in the design stage, as shown in figure 6. These include:

- Totally Active Workspace (TAWS)



**Figure 6.** This figure shows initial design sketches for three of the desktop tele-immersive environments that EVL is currently developing. From left to right, the Totally Active Workspace (TAWS), the Personal Penta Panel (P3), and the Personal Augmented Reality Immersive System (PARIS)

- Personal Penta Panel (P3)
- Personal Augmented Reality Immersive System (PARIS)

The Totally Active WorkSpace (TAWS) is a cubicle-sized 7 x 7 x 7 CAVE-like structure where the user works on a glass desk surface. As there is no floor projection in this model and its size is much smaller than the CAVE, we are also free to add a top-projected ceiling. Ideally this concept would be realized using large flat panel displays eliminating the need for rear projection and its huge consumption of space in an office environment, though rear projection will be used in the prototype. The TAWS is large enough for two colleagues to share the workspace and the LCD glasses can be set up to allow each user to see the image in the correct perspective.

The Personal Penta Panel (P3) is an open box made out of five large (greater than 40") flat panels. The user places his / her tracked head and hands into the box of screens and is presented with a surround stereo view. Current flat panels have a frame, creating seams that are difficult to eliminate. There are, however, optical methods to relay an image a few inches forward, which could be used to mostly eliminate these effects. It will be interesting to monitor human / computer interaction problems such as claustrophobia and simulator sickness with users of the P3.

The Personal Augmented Reality Immersive System (PARIS) is a desktop augmented reality device. This device will be initially prototyped using two 1280x1024 LCD projectors with electronic shutters compatible with active glasses to achieve stereo separation, and will later move to flat panel displays. The PARIS can also be used to prototype passive (polarized) stereo since we can polarize the two projector outputs, allowing very inexpensive and lightweight glasses to be incorporated, an important feature for use in museums and schools.

It is easier to deal with mounting cameras and microphones in these desktop arrangements compared to the CAVE or the ImmersaDesk. For example in the PARIS, the video camera(s) can be mounted looking through the half-silvered mirror. As the user's workspace is smaller it should also be easier to integrate gesture tracking and recognition, and haptic devices. Since the user is working on their desktop, this allows easier access to a keyboard and other familiar desktop materials which are difficult to bring with you to a CAVE or an ImmersaDesk. Ideally we would like to move to cheaper stereo glasses, and preferably to not needing glasses to mediate the images. We would like to use a high quality tetherless tracking system. We would like significantly higher resolution displays that can still be run in stereo at high refresh rates.

Each heterogeneous VR device in a tele-immersive collaboration will have its own strengths and weaknesses. CAVEs and other large VR devices will remain valuable for walking around within virtual environments and experiencing them life-size. Desktop VR devices will allow more convenient access to tele-immersive sessions though with a lesser degree of immersion. As more devices are produced and field-tested, their salient properties will emerge.

## 5. FUTURE WORK

Our research in tele-immersion is proceeding on several fronts. We are continuing to investigate and evaluate new display devices. We are continuing our development of CAVERNsoft to support tele-immersive applications and are working with our partners to help develop these applications. Aside from the technological problems there are many interesting psychological, communication, and cultural issues involved in trans-oceanic collaboration which can be explored once these systems are operational and deployed on a long-term basis.

## ACKNOWLEDGMENTS

The virtual reality research, collaborations, and outreach programs at the Electronic Visualization Laboratory are made possible through major funding from the National Science Foundation and the US Department of Energy; specifically NSF awards CDA-9512272, ANI-9712283, CDA-9720351, EIA-9802090, EIA-9871058, and the NSF Partnerships for Advanced Computational Infrastructure program.

The CAVE and ImmersaDesk are trademarks of the Board of Trustees of the University of Illinois.

## REFERENCES

1. G. Kessler, L. Hodges, and M. Ahamad, "Ravel, a support system for the development of distributed multi-user ve applications," in *Proceedings of the Virtual Reality Annual Interational Symposium (VRAIS) 98*, pp. 260–267, 1998.
2. J. Leigh, A. Johnson, and T. DeFanti, "Cavern: A distributed architecture for supporting scalable persistence and interoperability in collaborative virtual environments," *Virtual Reality: Research, Development and Applications* **2**(2), pp. 217–237, 1997.
3. R. Stevens, P. Woodward, T. Defanti, and C. Catlett, "From the I-WAY to the national technology grid," *Communications of the ACM* **40**(11), pp. 50–60, 1997.
4. A. Johnson, J. Leigh, and J. Costigan, "Multiway tele-immersion at supercomputing '97," *IEEE Computer Graphics and Applications* **18**(4), pp. 6–9, 1998.
5. T. Mason, E. Applebaum, M. Rasmussen, A. Millman, R. Evenhouse, and W. Panko, "The virtual temporal bone," in *Proceedings of the Medicine Meets Virtual Reality Conference*, 1998.
6. A. Johnson, M. Roussos, J. Leigh, C. Barnes, C. Vasilakis, and T. Moher, "The NICE project: Learning together in a virtual world," in *Proceedings of the Virtual Reality Annual Interational Symposium (VRAIS) 98*, pp. 176–183, 1998.
7. G. Wheless, C. Lascara, A. Valle-Levinson, D. Brutzman, W. Sherman, W. Hibbard, and B. Paul, "Virtual chesapeake bay: Interacting with a coupled physical/biological model," *IEEE Computer Graphics and Applications* **16**(4), pp. 52–57, 1996.
8. M. Thiebaux, "The virtual director," Master's thesis, Electronic Visualization Laboratory, University of Illinois at Chicago, 1997.
9. T. Imai, "The virtual mail system," Master's thesis, Electronic Visualization Laboratory, University of Illinois at Chicago, 1998.
10. J. Leigh, K. Park, R. Kenyon, A. Johnson, T. DeFanti, and H. Wong, "Preliminary STAR TAP tele-immersion experiments between chicago and singapore," in *Proceedings of HPC'Asia '98*, pp. 687–693, 1998.
11. T. DeFanti, D. Sandin, G. Dawe, M. Brown, M. Rawlings, G. Lindahl, A. Johnson, and J. Leigh, "Personal tele-immersion devices," in *Proceedings of the Seventh IEEE International Symposium on High-Performance Distributed Computing*, pp. 198–205, 1998.
12. "Guerilla VR," in *Visual proceedings of SIGGRAPH 98*, pp. 128–129, 1998.