Editors: Lawrence J. Rosenblum and
Michael R. Macedonia

Multiway Tele-Immersion at Supercomputing 97

Or, Why We Used $6,000,000 Worth of VR Equipment to do the Hokey Pokey

Andy Johnson, Jason Leigh, and Jim Costigan

University of Illinois at Chicago

Tele-Immersion—the union of networked virtual reality (VR) and video to support collaboration among scientists, engineers, and educators—is an important element in the computing information infrastructure envisioned by the National Computational Science Alliance. Tele-Immersion will let people from around the world casually enter a shared virtual environment (VE), manipulate that environment—whether a scientific simulation or a design space—and engage in discourse with their collaborators.

The Alliance consists of more than 50 universities and government institutions led by the National Center for Supercomputing Applications. Supercomputing 97 provided an excellent opportunity to connect several of the Alliance partners together and experiment with the kind of shared interaction my colleagues and I will be focusing on in the next few years. Rather than show a Tele-Immersion demo at Supercomputing, we wanted to use the conference as an opportunity to do research on the show floor.

At previous conferences, such as Supercomputing 95 and Siggraph 96, we shared VEs by using a handful of Cave Automatic Virtual Environments (CAVEs) and ImmersaDesks from the east coast to the west coast in the US and from the Gulf of Mexico to Chicago. At Supercomputing 97 we wanted to do more, both in terms of the number of participating sites and the geographic distances between them. In the past we also had a great deal of control over the hardware and software at the various sites, and over the networks that connected the sites together. This time we would rely on our collaborators.

Setting up the experiment

Several Alliance partners had limited experience with collaborative VR applications, so we wanted to get them immersed in a shared environment where they could interact with the virtual space and other users with as little hand-holding as possible. The software had to be easy to set up and run, and it had to work over the partners' existing networking. This would allow us to investigate the following:

- How easy is it to get the partners working together in an immersive collaborative VE? Will the participants be sociable?
- Can we coordinate efforts within that shared space using avatar representations and a simple audio communications link?
- What is the effect on the networking and graphics when we have multiple transglobal users interacting simultaneously over existing networks?

We decided to use the Narrative Immersive Constructionist/Collaborative Environments (NICE)[1] as our testbed. This environment was designed as a collaborative learning environment for young children. Given that 6-year-olds learned how to work in the space in about 10 to 15 minutes, we figured it would be simple for adults to use as well. However, while the interaction is simple, NICE follows the model that we believe will become the standard in future Alliance teleimmersive applications: a VE supported by a persistent computer simulation that collaborators can synchronously or asynchronously enter into and work with.

NICE consists of an island that children can explore. In the center of this island is a persistent garden. A simulation program, which has been continuously running since the summer of 1996, controls the garden's growth. Children can enter NICE, check on the progress of the garden (simulation), tend the garden (alter the parameters of the simulation), and discuss their work with other children in the space, represented via avatars. Another reason we chose NICE was that it is a "fun" environment that encourages social interaction, and we wanted the partners to have fun during this experiment.

One month before Supercomputing, we contacted the members of the CAVE Research Network User's Society (Cavernus), a user group of CAVE-based and CAVE-like VR systems, to see who wanted to participate. Since more than 50 CAVEs and ImmersaDesks exist worldwide, we didn't have any trouble finding collaborators. We gave the NICE application to the participants and encouraged them to run the application, enter the garden, and interact with it. At this point, we found to our surprise that the various sites were setting the viewing parameters of their ImmersaDesks differently. This meant that if two collaborators stood at exactly the same place in the virtual space, they would not see exactly the same thing. We also found that the avatars helped diagnose problems at a remote site. If we looked at a remote user's avatar and saw that her head was tilted at an unnatural angle or that her body was in an odd position, then the remote site might have problems with tracking that were not visible locally. This prompted us to think about prescribing a set of protocols to guarantee proper CAVE/ImmersaDesk calibration.

At Supercomputing

At the conference, we ran our experiment twice, for one hour each time. On the exhibit floor we used six ImmersaDesks and one ImmersaDesk2. Remote participants included four CAVEs, five ImmersaDesks, several CAVE simulators running on workstations, and one Web-based Java interface. The collaborators ranged from the west coast to the east coast of the US and included sites in the Netherlands and Japan. These collaborators also used a variety of Silicon Graphics hardware to run their VR systems. There were Maximum Impact-based CAVE simulators, ImmersaDesks run from Octanes, deskside Onyxs and Onyx 2s, and CAVEs run from rack Onyxs and Onyx 2s.

The full list of collaborators included

- Virtual Environments Lab at the Center for Coastal Physical Oceanography, Old Dominion University Norfolk, Virginia (ImmersaDesk)
- Electronic Visualization Laboratory, University of Illinois at Chicago, (two CAVEs, an ImmersaDesk, and a simulator)
- Virtual Reality/Virtual Environments Lab, Indiana University, Bloomington, Indiana (CAVE)
- SARA (Stichting Academisch Rekencentrum Amsterdam), Amsterdam, Netherlands (CAVE and simulator)
- Cray Research, Eagan, Minnesota (ImmersaDesk)
- The Scientific Computing and Visualization Group's Laboratory for Immersive Virtual Environments, Boston University, Boston (ImmersaDesk)
- Towa University, Fukuoka, Japan (ImmersaDesk)

Participants on the exhibit floor at Supercomputing 97 included

- Army Research Laboratory (ImmersaDesk)
- Laboratory for Advanced Computing, University of Illinois at Chicago (ImmersaDesk and Web-based Java client)
- Argonne National Laboratory (two ImmersaDesks)
- National Computational Science Alliance (ImmersaDesk2)
- Department of Computer Science, University of Illinois at Urbana-Champaign (ImmersaDesk)
- Virtual Reality/Virtual Environments Lab, Indiana University (ImmersaDesk)

All these collaborators were connected to the simulation that maintained the state of the garden, and to a central user datagram protocol (UDP) reflector, both located at the Interactive Computing Environments Laboratory of the University of Illinois at Chicago, as shown in Figure 1.

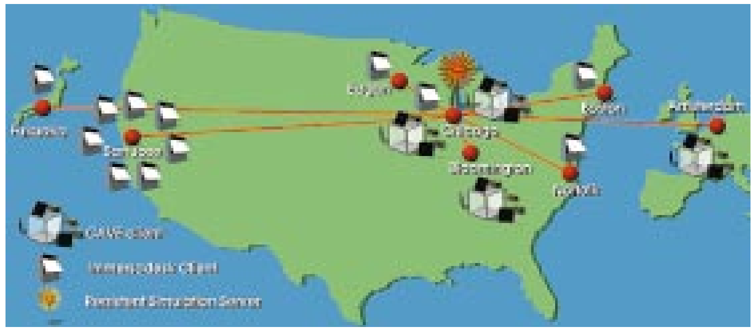

1 **A map of all of the sites involved in the experiment.**

2 **The popular "Dennis Rodman" avatar. When we participate in transcontinental or transglobal collaborations from our lab, it helps to have a recognizable Chicago personality as the avatar from Chicago.**

During our experiment, we had a few networking problems, partly due to intermittent networking at the conference itself and to running both the server for the garden and UDP reflector on the same Indy (which also acted as an active Web server). We intentionally placed the server on an active Web site, since this made it easy to distribute updated models for the garden. We were also curious to see how well a single server would perform under these conditions. From discussions with our partners, it seemed that there would typically be at most a handful of people collaborating in such a world, though there could be more for a distributed presentation within the space. As the number of clients increased, the server became overloaded, but it recovered well, and all the client software recovered dynamically to keep the environment going. This kind of fault tolerance meant that we didn't have to continually restart every client when the network went down. Another problem we encountered was with remote clients who tried to connect through their local firewalls. The other collaborators could see them, but they could not see anyone else.

Within the space itself we were partly a victim of our own success. We had designed several avatars for the original NICE application and added several more for this experiment, but not enough to give everyone a unique character (more than a few people wanted to be the Dennis Rodman avatar shown in Figure 2). The plan

3 **Photos taken by Tom Coffin at Supercomputing during the collaboration.**

Images courtesy of Tom Coffin

4 **Snapshots taken within the virtual space during the collaboration. Clockwise from top left: SARA in the Netherlands offering a flower to one of the participants in Chicago; the Army Research Lab claiming the high ground; a small gathering in the catacombs; and a crowd in the garden.**

was that each site would develop its own Virtual Reality Modeling Language (VRML)-based avatar body parts, but there was not enough time to do this. While each participant could choose a name for their chosen avatar and have that name appear on their character's chest, a certain amount of confusion still existed. We encouraged names such as "Kukimoto_in_Japan" or "Jim_at_SC97" so that we could not only tell to whom we were talking to, but how far away they were.

We set up a conference call so participants could talk to each other in real time. Audio amplitude was sent over the network along with other positional information to move each avatar's mouth and body. The mouth movement allowed remote users to correlate a voice with an avatar. This technique has worked well in past small group collaborations—we wanted to see how it worked with 10 or more people. It worked better than expected when a single person presented or lead and others responded or asked questions. Communication started to break down when multiple tasks or groups began to form, as the audio had no correlation to the virtual space. Remote users who were nearby in the virtual space had the same volume as those far away. We have not found an inexpensive solution to this problem that still delivers the real-time audio necessary to sustain this level of interaction.

Differences between the various hardware platforms quickly became noticeable as the number of users increased beyond 10. The slower SGI workstations (that is, anything less than an Onyx) could not keep up with the rendering. This was a larger hindrance to collaboration than network lag, though the guaranteed real-time audio link kept the situation from getting too chaotic.

We were interested in whether people would be sociable in the space and whether they could coordinate their efforts and work together. Part of the demo was devoted to letting various subgroups form and explore parts of the island such as in a VR multiuser dungeon (MUD). Another part focused on "centralized" activity such as clearing the garden of weeds.

Since the audio link was a real-time link between all the sites, this allowed us to look at the lag in the transmission of avatar position data. To get a sense of this within the environment, we gathered most of the collaborators together and did the Hokey Pokey. The Hokey Pokey is a rather silly American dance where the instructions and lyrics follow the general pattern of "You put your right hand in, you take your right hand out, you put your right hand in and you shake it all about, you do the Hokey Pokey and you turn yourself around, that's what it's all about." Since our avatars had a head, a body, and one arm, we modified the dance slightly, but you could get a very good idea what each user was doing (at least when the server wasn't overloaded). With everyone listening and several people singing the instructions with minimal lag, it allowed us to see how much time it took for the avatars to go through the motions.

Figures 3 and 4 show several photos taken during the multiway demo, both live shots taken on the floor of Supercomputing 97 and shots taken by a virtual camera within the virtual space.

Future work

We found that it was possible to quickly field a collaborative VR environment over existing networks, though we need to work on ensuring that all our collaborators have properly calibrated equipment. We also observed that the participants were very sociable and they could coordinate action within the space. We found that building in automatic recovery into the server and clients was very important, that the server should reside on a decent machine, and that guaranteed quality-of-service is a necessity for the next generation of these kind of VEs. And we discovered that our collaborators had a lot of fun participating in our experiment.

We believe it's important to build these collaborative worlds and let domain scientists and educators play with them to see where we should focus our efforts. This experiment showed us that a lot of stuff works, but a lot can be improved. The venue for those improvements will be CavernSoft[2,3]—the hybrid real-time networking and persistent datastore library we're currently developing to better support data distribution in Tele-Immersion. CavernSoft uses Nexus from Argonne National Labs for its networking support and PTool from the Laboratory for Advanced Computing at the University of Illinois at Chicago for its database support. Also, collaborative spaces such as these bring up many interesting communications issues, which we're continuing to investigate.

So what's in store for us at Supercomputing 98? Our experiment allowed us to look at logistical issues, both outside and inside the virtual space. Supercomputing 98 may be a good place to look more in depth at the communication issues within a couple of specific collaborative visualization and design spaces. Plus, we still have several more continents to connect to. . . and then, of course, there's Mir. ■

Related Web Pages

Here are some related Web pages of projects or tools mentioned in this article.

Tele-Immersion at EVL: http://www.evl.uic.edu/spiff/ti/
NICE: http://www.ice.eecs.uic.edu/~nice
Cavernus: http://www.ncsa.uiuc.edu/VR/cavernus/
Nexus: http://www.mcs.anl.gov/nexus/
PTool: http://www.lac.uic.edu/

Acknowledgments

We would like to thank the following people for helping to make our experiment possible at the various sites: Joe Alexander, Rick Angelini, Stuart Bailey, Lilac Berniker, Edward Berrdveld, Hisham Bizri, Glenn Bresnahan, Erik Brisson, Maxine Brown, Russell Burgett, Tom Coffin, Junna Cui, Kathleen Curry, Marek Czernuszenko, Fred Dech, Tom DeFanti, Ralph De Stefano, Margaret Dolinsky, Ian Foster, Tom Frisch, Sara Graffunder, Karen Green, Bob Grossman, John Hicks, Tomoko Imai, Joe Insley, Wilfred Janssen, Ted Jordan, Hank Kaczmarski, Abhinav Kapoor, Kelly Kirk, Cal Kirchhof, Anton Koning, Dick LaRoche, Cathy Lascara, Johnny Lawson, Gary Lindahl, Nobuyuki Kukimoto, John Melchi, Bill Nickless, Dave Pape, Mike Papka, Kyoung Park, Dana Plepys, Robert Putnam, Larry Ramey, Dan Reed, Paul Rossman, Nancy Rowe, Bill Sherman, Shuji Suma, Larry Smarr, Craig Stewart, Rick Stevens, Matt Szymanski, Samroeng Thongrong, Jeroen van Hoof, Rene van der Moolen, Luc Verburgh, Alan Verlo, Fang Wang, Eric Wernert, and Glenn Wheless.

References

1. A. Johnson et al., "The NICE Project: Learning Together in a Virtual World," *Proc. VRAIS 98*, IEEE Computer Society Press, Los Alamitos, Calif., 1998, pp. 176-183.
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*, Vol. 2, No. 2, Dec. 1997, pp. 217-237.
3. J. Leigh, A. Johnson, and T. DeFanti, "Issues in the Design of a Flexible Distributed Architecture for Supporting Persistence and Interoperability in Collaborative Virtual Environments," *Proc. Supercomputing 97*, 1997, http://www.supercomp.org/sc97/proceedings.

Contact Johnson at the Electronic Visualization Lab, University of Illinois at Chicago, 851 S. Morgan St., Room 1120 SEO, Chicago, IL 60607-7053, e-mail ajohnson@eecs.uic.edu.