


Visualization in Teleimmersive Environments



In teleimmersion, collaborators at remote sites share the details of a virtual world that can autonomously control computation, query databases, and gather results. They don't meet in a room to discuss a car engine. They meet in the engine itself.

Jason Leigh

Andrew E. Johnson

Maxine Brown

Daniel J. Sandin

Thomas A. DeFanti

Electronic Visualization Laboratory, University of Illinois at Chicago

In collaborative virtual reality (VR), the goal is to reproduce a face-to-face meeting in minute detail. Teleimmersion moves beyond this idea, integrating collaborative VR with audio- and video-conferencing that may involve data mining and heavy computation. When participants are teleimmersed, not only can they virtually see and interact with each other, but when they leave, their environment continues to evolve, autonomously controlling computation and data access.

The shared virtual environment can be many things—part of a car design; visualization of climatic data such as a storm; or even some three-dimensional environment that does not physically exist. Collaborators typically enter the environment as avatars—lifelike computer-generated representations. The environment transmits gestures as well as audio and video, so users have a greater sense of presence in the shared space than they would with other collaborative media.

This adds a new dimension to virtual collaboration. In teleimmersion, participants are not talking about a storm; they are standing inside it. They are not looking at a scale model of a car design; they are standing inside the engine block. They can then change parameters on a supercomputing simulation of the storm or engine block and collectively study the impact of those changes. With these enhanced capabilities, collaborators may choose to work virtually even if more traditional face-to-face meetings are possible.

The University of Illinois at Chicago's Electronic Visualization Laboratory (EVL) has hosted several applications that demonstrate rudimentary teleimmersion. All of our users are members of Cavern—the CAVE Research Network (<http://www.evl.uic.edu/cavern>)—a collection of participating industrial and research institutions equipped with CAVE (Cave Automated Virtual Environment), ImmersaDesk VR systems, and high-performance computing resources, including high-speed networks.

There are more than 100 CAVE and ImmersaDesk installations worldwide. The pressing challenge now is how to support collaborative work among Cavern users without having them worry about the details of sustaining a collaboration. Our goal is to free designers and scientists to do their real work within shared environments. To that end, we are building

- new display devices to serve as more convenient teleimmersion end points;
- an international networking infrastructure with sufficient bandwidth to support the needs of teleimmersive applications; and
- a software framework that researchers can use to rapidly create teleimmersive applications in which CAVEs, ImmersaDesks, other VR devices, and desktop workstations can work together seamlessly.

We are also tackling the problem of how to provide both synchronous and asynchronous collaboration. When collaborators are on the same continent, conducting synchronous sessions in the shared space is straightforward because time differences are minimal. When collaborators are across the globe, however, the greater time differences require asynchronous collaboration—collaborators working in the same virtual space at different times so they can work during their normal work hours. Supporting this type of collaboration is one of the main reasons we adopted the model of a continuing virtual world (one that persists even when collaborators are not present).

DISPLAY TECHNOLOGIES

Until 1992, when EVL developed the CAVE,¹ the VR community used mainly head-mounted displays—heavy and clumsy helmets with liquid crystal displays or cathode ray tubes (CRTs) mounted in front of eyepieces. As Figure 1 shows, the CAVE VR system is a 10 × 10 sq. ft. room in which CRT projectors project stereoscopic



Figure 1. The CAVE (Cave Automated Virtual Environment). CRT projectors project stereoscopic computer images onto the CAVE's translucent walls, giving the occupant (user), who is wearing special glasses, the illusion that objects surround him. The user interacts with the environment through a wand (not shown). —Image courtesy of the Electronic Visualization Laboratory

images. The images give the CAVE occupants, or users (there can be up to 10 users), the illusion that objects surround them. Users don lightweight liquid crystal shutter glasses to resolve the stereoscopic imagery and hold a three-button wand for 3D interaction with the virtual environment. An electromagnetic tracking system attached to the glasses and the wand lets the CAVE determine the location and orientation of the user's head and hand at any given moment. This information instructs the Silicon Graphics Onyx that drives the CAVE to render the images from the user's point of view.

The ImmersaDesk and its successor, the ImmersaDesk2, are smaller, drafting-table-like systems, as Figure 2 shows. They, too, can project stereoscopic images, but they fill different application needs. Whereas the CAVE is well suited for providing panoramic views (particularly useful for visualizing architectural walk-throughs), the ImmersaDesk is designed for displaying images that fit on a desktop, such as CAD models. Applications built for the CAVE are fully compatible with the ImmersaDesk.

The EVL is continuing its efforts to develop display technologies, aiming for more compact devices that



Figure 2. The ImmersaDesk. Like the CAVE, the ImmersaDesk projects stereoscopic images, but it is designed for images on a desktop scale, such as CAD models. —Image courtesy of the Electronic Visualization Laboratory

support teleimmersion in an office setting. We have recently developed prototypes of three new devices.

ImmersaDesk 3

The ImmersaDesk3 is an ImmersaDesk built from a 42-inch plasma screen mounted on a conventional office desk. The user can tilt the screen from fully horizontal to fully vertical, which significantly reduces the device's size and cost. Unfortunately, current plasma screen technology generates too much electromagnetic noise, making head tracking and stereo synchronization difficult.²

PARIS

As in the ImmersaDesk3, the Personal Augmented Reality Immersive System (PARIS), shown in Figure 3, is a desktop device that is placed around an office desk. The user sits at the desk and places his hands under a semitransparent mirror, which allows him to see his hands and the computer images simultaneously. We plan to place cameras under the mirrors to track hand location and gestures so that users can use natural hand motions to manipulate virtual 3D objects.





Figure 3. The Personal Augmented Reality Immersive System. PARIS lets users see their hands at the same time they see the computer imagery. —Image courtesy of Samroeng Thongrong of the Electronic Visualization Laboratory

Figure 4. The AccessBot teleconferencing device. The remotely controllable pan-tilt-zoom camera mounted above the plasma display lets users look around the meeting room. AccessBot gives the remote user an equal presence with local meeting participants. —Image courtesy of Jason Leigh of the Electronic Visualization Laboratory



PARIS is the first VR display system prototyped in VR before it was physically constructed. We wrote a PARIS simulation to run in the CAVE so that a designer can sit in the CAVE and observe the effects of the mirror's tilt angle on image placement. Past work in developing VR displays was based solely on mathematical models, which do not take into account user preferences and application-specific requirements. For example, a particular screen angle may be suitable when the user must navigate through a space, but another angle may be better when he must directly manipulate small virtual objects.

AccessBot

The AccessBot is a teleconferencing device that uses a life-sized plasma display screen, high-fidelity video and audioconferencing, a robotic camera, and high-speed networking to ensure that a remote participant commands an equal presence (telepresence) with the local participants in a meeting. This capability is especially important in providing accessibility for the handicapped. As Figure 4 shows, the AccessBot contains a pan-tilt-zoom camera, which lets the remote operator look around the meeting environment. The zoom capability further empowers users by letting them see at greater distances than is possible with normal vision.

A prototype AccessBot has been deployed at the National Center for Supercomputing Applications' Access Center in Washington, D.C. The next generation will include the ability to pan the entire plasma screen, hence simulating the turning of a human torso. It will be self-deploying and self-configuring so that it can be shipped to a meeting location and deployed with minimum human intervention. Finally, it will leverage high-bandwidth wireless networking as it becomes available.

NETWORKING

Connecting interactive visualization environments such as the CAVE and ImmersaDesk into teleimmersive visualization environments requires sophisticated networking. We view teleimmersion as an application driver for the networking community because it stimulates research into fine-tuning networks so they can handle the transmission of extremely large data sets or multimedia, where latency and bandwidth are of major importance.

Data distributed between teleimmersive applications is complex. It may consist of real-time audio and video streams, real-time tracker data to determine the position and orientation of participants in the space (so that participants can convey natural gestures), updates to shared information, large 3D models, and large multi-dimensional data sets. To deliver this variety of data types, teleimmersive applications must open multiple connections between collaborating clients. Thus, real-time data is transmitted over low-latency and often low-bandwidth links, and audio and video data is broadcast or multicast over high-bandwidth, low-latency links. Models and data sets require a reliable high-bandwidth link. For smooth interaction in all these scenarios, there must be some guaranteed quality of service.

To satisfy these requirements, the EVL has participated in building three networks:

- the Metropolitan Research and Education Network (MREN), a high-performance network for advanced applications (<http://www.mren.org>) centered in the midwestern US;
- the Very High Performance Backbone Network Service (vBNS), an NSF-funded nationwide network that supports high-performance, high-bandwidth research applications (<http://www.vbns.net>); and
- the Science, Technology, and Research Transit Access Point (Star Tap), an international infrastructure that facilitates the long-term interconnection and interoperability of advanced networking (<http://www.startap.net>). Connections include Russia, the Asian Pacific Advanced Network (APAN) consortium, the Nordic countries, and CERN.

Some of these networks, such as the vBNS, have been operating for several years. Others, such as Star Tap, are just forming. Teleimmersive applications are only beginning to use these networks, and we will be evaluating network use in depth during the next few years. The results should help us fine-tune the networks to better support teleimmersive applications.

SOFTWARE FRAMEWORKS

High-speed low-latency networks are only half the

requirement to produce teleimmersive applications. Application developers also need a high-level software framework. Most existing efforts to develop software frameworks for collaborative virtual environments³ focus on existing low-bandwidth Internet infrastructures or massive connectivity involving thousands of simultaneous participants (as in military simulations or computer games). The use of VR for manufacturing scientific and information visualization has a different set of requirements. Groups tend to be smaller and more focused, and they have more demanding data distribution requirements because they need high-fidelity audio and video communications and must share large engineering and scientific data sets.

At the EVL, we use the Cavernsoft software framework to build teleimmersive applications.⁴ At its core is the Information Resource Broker (IRB), which was inspired by the CORBA object request broker (ORB), but has lower overhead. Like the ORB, it emulates both a distributed shared-memory system and a message-passing system (a useful model for data sharing in teleimmersion), but unlike the ORB and traditional systems, it does not unify the two systems. Instead, data distribution can occur over either reliable TCP, UDP, or multicast, which teleimmersive applications usually need to distribute their many data classes. The IRB also allows shared-memory segments to persist so that client programs can cache frequently used data at local sites.

Applications built with the IRB are client-server-symmetric—all client programs are automatically server programs without the need for additional programming. This symmetry means that a collection of clients can easily form arbitrary collaborative topologies and thus that collaborators can initiate spontaneous sessions.

Cavernsoft also provides a continually growing suite of IRB-based libraries that support basic teleimmersion needs, such as avatars, audioconferencing, collaborative interfaces, and virtual e-mail. However, from almost a decade of experience with technology transfer, we have learned that it is not enough to simply provide a suite of tools. We must also provide a higher level tool organization as an application framework so that users can conceptualize how to use the tools in a specific scenario. To provide that support, we created Limbo, a teleimmersion application framework. Limbo and its successor, Tandem, support audioconferencing, avatar rendering, model importing, and data distribution and manipulation. Collaborators use a basic Limbo space to quickly begin working in a shared virtual space without programming. As users import 3D models (car designs, scientific data sets, and so on) into the space, Limbo distributes them to all remote participants, who can begin to collectively modify them. We have also made the source code available so that application devel-



Figure 5. General Motors' VisualEyes. One of VisualEyes' creators, Randall C. Smith, evaluates an interior design with a remote collaborator represented as an avatar. —Image courtesy of Randall C. Smith of the General Motors Research and Development Center

opers can jump-start the development of their own domain-specific teleimmersive applications.

SAMPLE APPLICATIONS


Cavern users' application requirements motivated the development of much of the technology we describe. Candidate teleimmersion problems benefit, first, from visualization in an immersive environment and, second, from a collaborative solution. The seven Cavernsoft teleimmersive applications we describe here meet these criteria.⁵ They are compelling because they not only solve problems in their respective domains, but also provide real-world test cases that show us how collaborators work in high-end visualization environments and reveal the demands being placed on networks.

VisualEyes

VisualEyes was created by Randall C. Smith and David A. Brown of the General Motors Research and Development Center. Designers import 3D CAD models into the CAVE for quick visual inspection and can then do a life-sized design review. 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. GM has since extended VisualEyes so that its transglobal design and manufacturing teams can collaborate on remote design reviews, as shown in Figure 5. The goal is to allow designers to both synchronously and asynchronously access a design that persists and evolves.⁶

In a typical VisualEyes working scenario, a designer modifies the design on a workstation in a 3D modeler such as Alias. GM wrote a plug-in for Alias that allows direct communication between Alias and VisualEyes. Changes in lighting and materials are propagated automatically to the networked virtual environment, and all collaborators see the changes

Figure 6. The Virtual Temporal Bone. Three users discuss the structure of the inner ear using pointers to identify and move its various components. The Virtual Temporal Bone makes it much easier to see the spatial layout of the inner ear structures, compared with the flat images in medical texts. — Image courtesy of Chris Scharver of the Electronic Visualization Laboratory



simultaneously. They can critique the design and suggest changes to the designer, who can make them immediately at the CAD workstation.

Virtual Temporal Bone

Virtual Temporal Bone, shown in Figure 6, was created by Mary Rasmussen, Theodore Mason, and Alan Millman of the University of Illinois at Chicago's Virtual Reality Medicine Laboratory and Chris Scharver and Mohammed Dastagir Ali of the EVL. It is a teleimmersive education program that lets a remote physician teach medical students about the 3D structure and function of the inner ear. The students and instructor can point at and rotate the ear to view it from various perspectives. They can also strip away the surrounding outer ear and temporal bone to view the inner anatomy more clearly. They use audio from the voice conference to modify the eardrum flapping and thus illustrate its function. Virtual Temporal Bone leverages the CAVE's stereoscopic capabilities and the ImmersaDesk to make the spatial layout of the inner ear structures clear—something difficult to do using the flat images in standard medical texts.

Alive

The Architectural Linked Immersive Virtual Environment was created by Edward J. Breedveld, Jorrit Adriaanse, Anton Koning, and Paul Wielinga of Stichting Academisch Rekencentrum Amsterdam (Academic Computing Services Amsterdam) and Jason Leigh and Dana Plepys of the EVL. The idea for Alive dates back to 1998, when Dutch architect Rem Koolhaas won an award for his design of the new Campus Center at the Illinois Institute of Technology in Chicago. A year later, the Netherlands joined the Star Tap network, which created an excellent opportunity for architects in Amsterdam and their clients in Chicago to use teleimmersion technology to collaboratively review the design. EVL was again able to monitor network use in a real-world teleimmersion application.

The EVL and SARA collaboration yielded Alive, a collaborative edition of SARA's software—an appli-

cation that lets viewers navigate through 3D CAD models to create and replay prerecorded animation paths. EVL used Cavernsoft to provide Saranav with avatars, audioconferencing, and network instruments that monitor bandwidth use, latency, and jitter during collaborative sessions.

CAVE6D

CAVE6D, a tool for the teleimmersive visualization of environmental data, was created by Glen H. Wheless and Cathy Lascara of the Center for Coastal and Physical Oceanography and Abhinav Kapoor of the EVL.

CAVE6D evolved from CAVE5D,⁷ a configurable VR application framework. CAVE5D was supported by Vis5D, a powerful graphics library that provides visualization techniques to display multidimensional numerical data from atmospheric, oceanographic, and other similar models, including isosurfaces, contour slices, volume visualization, wind/trajectory vectors, and a variety of image-projection formats.

CAVE6D integrates CAVE5D and Cavernsoft to produce a teleimmersive environment in which multiple CAVE5D users jointly visualize, discuss, and interact with data sets. Avatars have long pointing rays that collaborators can use to point at features of interest during a conference call. CAVE6D extends CAVE5D visualization parameters such as salinity, circulation vectors, temperature, and wind to allow collaborators to collectively operate them. Collaborators can also operate visualization tools independently. Because it distributes multiple observable dimensions among collaborators, CAVE6D effectively reduces the amount of data a single user must interpret. We are now studying the conditions under which collaborative views—*independent or coordinated*—affect the speed and quality of data interpretation.

Round Earth

Round Earth, shown in Figure 7, was created by Andrew Johnson of the EVL, Thomas G. Moher of UIC's Interactive Computing Environments Laboratory, and Stellan Ohlsson of UIC's psychology department. This educational project uses teleimmersion to teach young children the concept of a spherical Earth.⁸ As two children collaborate in exploring a small spherical asteroid, one child, the astronaut, explores the surface of the asteroid; the other child, mission control, guides the astronaut from an orbital (spherical) view. VR helps situate the astronaut child on the surface of the asteroid, where he can experience circling the globe and coming back to the same place without falling off at the bottom. He can also see objects appear over the horizon top first. VR gives the mission control child an obviously spherical world to monitor. The children share the




Figure 7. The Round Earth. Two children share a virtual asteroid from different perspectives. The child acting as an astronaut views the environment from the surface (left). The child acting as mission control views the entire spherical asteroid (right). The two children must integrate their views, and by doing so, learn how to map the apparently flat surface to its true spherical shape.—Image courtesy of Andrew E. Johnson of the Electronic Visualization Laboratory

virtual environment but see it in different ways. They must integrate these views to complete their mission. Through integrating, they learn to map from the apparently flat surface of the sphere to its true spherical shape.

This work is part of a larger study to investigate how elementary school science teachers should use VR. As part of this study, we deployed an ImmersaDesk in a local elementary school, where it will remain for the next two years. We hope to gain more insights into how the additional representational bandwidth of VR-based visualizations can support the construction of mental models.

Silk Road Cave Shrine

The Silk Road Cave Shrine was created by Samroeng Thongrong, Jason Leigh, Maggie Rawlings, Mohammed Ali Khan, and Yalu Lin of the EVL and Sarah Fraser, Bill Parod, Jim Chen, Dennis Glenn, and Harlan Wallach of Northwestern University. It is a collaboration among historians, artists, and computer scientists to create a virtual cultural and artistic exhibit of the Mogao Grottoes of Dunhuang. Dunhuang, one of western China's ancient cultural sites, is considered the gateway to the well-known Silk Road—the East-West trade route between Asia and Europe. The Mogao Grottoes consist of 492 caves with murals covering 25,000 square meters (approx. 269,097 sq. ft.), wall fresco paintings, and more than 3,000 painted sculptures. The caves were built over a thousand years, from 300 to 1300 A.D. Access to these caves is limited, and photographs do not accurately convey the relationships among the murals that surround the visitor. Teleimmersion lets an historian take remote visitors on tours of the caves.



Figure 8. The Teleimmersive Data Explorer. TIDE lets collaborators explore terascale data sets. Here, TIDE is compressing a sponge lattice, which it generated from a super-computer simulation involving terabytes of data.—Image courtesy of Jason Leigh of the Electronic Visualization Laboratory

The virtual exhibit of the Grottoes has been under development for a year. The EVL has created 3D models using data from researchers at NWU and the Silk Shrine Foundation in China. The models have been imported into Limbo as a rapid collaborative evaluation and brainstorming tool. When the exhibit is finished, it will be the first collaborative virtual exhibit of an historic site.

TIDE

The Teleimmersive Data Explorer (TIDE), shown in Figure 8, was created by Nikita Sawant, Chris Scharver, and Jason Leigh of EVL and Emory Creel, Georg Reinhart, Suma Batchu, Stuart M. Bailey, and Robert Grossman of UIC's National Center for Data Mining. It is designed to let collaborators query and access terascale data sets.⁹

Multiple remote users—for example, one in a CAVE, another on an ImmersaDesk, and yet another

Teleimmersion currently provides the closest emulation of physical colocation possible.

on a desktop workstation—see each other as avatars. TIDE streams digital audio between sites, which may be hundreds of miles apart, allowing the participants to speak to each other naturally. The desktop workstation displays a dataflow model that an operator can use to construct the visualization the three devices share. The VR display participants can use 3D tools to directly manipulate the visualization.

Suppose, for example, that the CAVE user changes an isosurface value in the data set. TIDE automatically propagates the change to the other visualization displays. Meanwhile, the ImmersaDesk user, noticing an anomaly in the data set, inserts an annotation as a reminder to examine the region more closely later. He then instructs a remote rendering server with gigabytes of memory and terabytes of disk space to render the images in full detail as a stereoscopic animation sequence. These animations will take some time to generate, so the users continue to examine other aspects of the data set. Eventually, the rendering is complete, and the remote server streams the animation to each collaborator for viewing.

Research continues on how to apply TIDE. The current focus is on visualizing massive data sets from the US Department of Energy's simulations of nuclear stockpiles and weather data from the National Oceanic and Atmospheric Administration. Other research will attempt to develop a range of tools and techniques. We need tools to manage and mine massive data sets for visualization and navigation. We also need tools that are network aware (can use network quality of service when available) and that can create 3D annotations and recordings of discoveries during collaborative sessions. We plan to study how users collaboratively explore data so that we can evaluate the tools' effectiveness.

In 1998, a joint US National Science Foundation and Department of Energy workshop report noted that the rate at which today's science problems are generating data far outpaces our current ability to visualize it.¹⁰ We need new techniques that combine the power of networking, data mining, and visualization to bridge that gap. Teleimmersion in this context is particularly challenging because multiple users must coordinate access to the volumes of data that a single user now views.

The obstacles to making *global* teleimmersion a stable, convenient, and persistent technology are formidable. We must build systems that can bridge cultural as well as language differences if these spaces are to support collaboration effectively. We also need broadly deployable high-bandwidth, low-latency net-

works that can provide the quality of service needed for real-time interactive applications.

Yet even with these challenges, the promise of teleimmersion is strong. Research in computer-supported cooperative work has found that the highest quality of collaboration occurs when collaborators are physically colocated.¹¹ Teleimmersion currently provides the closest emulation of physical colocation possible, relative to telecollaborative technologies such as videoconferencing. Persistent teleimmersion servers will provide a permanent environment to support the asynchronous and spontaneous sessions needed to encourage long-term collaboration. Teleimmersion can also enhance productivity by allowing a collaborator to hand off his work at the end of his day to those just starting theirs. ♦

Acknowledgments

The virtual reality research, collaborations, and outreach programs at the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago are made possible by major funding from the National Science Foundation (NSF), awards EIA-9802090, EIA-9871058, ANI-9712283, ANI-9730202, and ACI-9418068, as well as NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement ACI-9619019 to the National Computational Science Alliance. The EVL also receives major funding from the US Department of Energy (DOE), awards 99ER25388 and 99ER25405, and support from the DOE's Accelerated Strategic Computing Initiative (ASCI) Data and Visualization Corridor program. In addition, the EVL receives funding from Pacific Interface on behalf of NTT Optical Network Systems Laboratory in Japan. CAVE, ImmersaDesk, and ImmersaDesk2 are trademarks of the Board of Trustees of the University of Illinois.

References

1. C. Cruz-Neira, D. Sandin, and T. DeFanti, "Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE," *Proc. SIGGRAPH 1993*, ACM Press, New York, 1993, pp. 135-142.
2. D. Pape et al., "The ImmersaDesk3—Experiences with a Flat Panel Display for Virtual Reality," *Proc. 3rd Int'l Immersive Projection Technology Workshop*, ACM Press, New York, 1999, pp. 107-112.
3. G. Kessler, L. Hodges, and M. Ahamad, "RAVEL, a Support System for the Development of Distributed, Multi-user VE Applications," *Proc. Virtual Reality Ann. Int'l Symp. (VRAIS)*, IEEE CS Press, Los Alamitos, Calif., 1998, pp. 260-267.

4. J. Leigh, A. Johnson, and T. DeFanti, "Cavern: A Distributed Architecture for Supporting Scalable Persistence and Interoperability in Collaborative Virtual Environments," *J. Virtual Reality Research, Development and Applications*, Vol. 2, No. 2, 1997, pp. 217-237.
5. J. Leigh et al., "A Review of Tele-Immersive Applications in the CAVE Research Network," *Proc. IEEE Virtual Reality 99*, IEEE CS Press, Los Alamitos, Calif., 1999, pp. 180-187.
6. A. Johnson et al., "Next-Generation Teleimmersive Devices for Desktop Trans-Oceanic Collaboration," *Proc. IS&T/SPIE Conf. Visual Communications and Image Processing*, 1999, pp. 1,420-1,429.
7. G.H. Wheless et al., "Virtual Chesapeake Bay: Interacting with a Coupled Physical/Biological Model," *IEEE Computer Graphics and Applications*, July 1996, pp. 52-57.
8. A. Johnson et al., "The Round Earth Project: Deep Learning in a Collaborative Virtual World," *Proc. IEEE Virtual Reality Ann. Int'l Symp. (VRAIS)*, IEEE CS Press, Los Alamitos, Calif., 1999, pp. 164-171.
9. J. Leigh et al., "A Methodology for Supporting Collaborative Exploratory Analysis of Massive Data Sets in Tele-Immersive Environments," *Proc. 8th Int'l Symp. High-Performance and Distributed Computing*, IEEE CS Press, Los Alamitos, Calif., 1999, CD-ROM.
10. P. Smith and J. Van Rosendale, eds., *Data and Visualization Corridors: Report on the 1998 DVC Workshop Series*, Tech. Report CACR-164, Center for Advanced Computing Research, California Inst. Technology, Pasadena, Calif., 1998; <http://www.cacr.caltech.edu/publications/DVC/index.html>.
11. J. Olson and G. Olson, "What Mix of Video and Audio Is Useful for Small Groups Doing Remote Real-Time Design Work?" *Proc. SIGCHI 95*, ACM Press, New York, 1995, pp. 362-368.

Jason Leigh is a research scientist in the Electronic Visualization Laboratory at the University of Illinois at Chicago. His research interests include the development of teleimmersion tools to explore massive data sets, the study of computer-mediated human interaction, and the application of advanced networking technology to support wide-area collaboration. Leigh received a BS in computer science from the University of Utah and a PhD in computer science from the University of Illinois at Chicago. Contact him at spiff@uic.edu.

Andrew E. Johnson is an assistant professor in the Department of Electrical Engineering and Computer

Science at the University of Illinois at Chicago. His current research focus is on teleimmersive virtual environments—in particular, their application in education. He received a BS in computer engineering from the University of Michigan and an MS and a PhD in computer science from Wayne State University.

Maxine Brown is an associate director of the Electronic Visualization Laboratory, where she works closely with EVL's collaborators, the National Computational Science Alliance and Argonne National Laboratory, to facilitate research activities. She is also responsible for the funding, documentation, and promotion of the EVL's research activities. She has worked in computer graphics and supercomputing for several decades. She received a BA in mathematics from Temple University and an MSE in computer science from the University of Pennsylvania.

Daniel J. Sandin is a director of the Electronic Visualization Laboratory and a professor in the School of Art and Design at the University of Illinois at Chicago. He is responsible for EVL's research activities in virtual environments, teleimmersion, scientific visualization, computer-based interactive systems, computer animation, distributed computing, sonification, and abstract math visualization. Sandin was also a principal collaborator in conceiving the CAVE virtual reality theater in 1991. He has received grants from the Rockefeller Foundation, the Guggenheim Foundation, the National Science Foundation, and the National Endowment for the Arts.

Thomas A. DeFanti is a professor in the Department of Electrical Engineering and Computer Science at the University of Illinois at Chicago, a director of the Electronic Visualization Laboratory, and director of UIC's Software Technologies Research Center. He is also a principal investigator in the National Computational Science Alliance at the University of Illinois at Urbana-Champaign. DeFanti has strong working collaborations with partners of the Alliance and with Argonne National Laboratory. He is a member of the Board of Trustees of the University Corporation for Advanced Internet Development (UCAID—the parent organization of Internet2) and chair of the UCAID Applications Strategy Council. He is also an ACM Fellow. DeFanti was a principal collaborator in conceiving the CAVE virtual reality theater.