

Global Tele-Immersion: Better Than Being There

Jason Leigh, Thomas A. DeFanti, Andrew E. Johnson, Maxine D. Brown, and Daniel J. Sandin

University of Illinois at Chicago
Electrical Engineering and Computer Science Department
Electronic Visualization Laboratory
851 S. Morgan Street, Room 1120, Chicago, IL 60607-7053 USA
spiff@evl.uic.edu, tom@uic.edu
<http://www.evl.uic.edu>

Abstract

The term *Tele-Immersion* was first used in October 1996 as the title of a workshop the Electronic Visualization Laboratory (EVL) at University of Illinois at Chicago (UIC) organized in Chicago to bring together researchers in distributed computing, collaboration, virtual reality (VR), and networking. Workshop attendees paid specific attention to the future needs of applications in the sciences, engineering, and education. EVL defines Tele-Immersion as the union of networked VR and video in the context of significant computing and data mining. Some researchers use the terms Collaborative Virtual Environment (CVE) or Distributed Virtual Environment (DVE) to describe the field of networked VR. Tele-Immersion, as defined by the authors, encompasses more image-based technology, like video and bit maps, than CVE/DVE researchers typically envision, so that more “reality” (so to speak) is incorporated. It also has the benefit of being pronounceable. Global Tele-Immersion is “better than being there” because physical travel, especially international travel, is best devoted to vacation, not work, in the authors’ experienced opinion.

Key words: Virtual Reality, Tele-Immersion, Tele-Presence, Networking, Virtual Environment, CAVE, Spatially Immersive Display, Projection-Based Display Technologies.

1. From Demonstration to Discovery Mode

The CAVE™ [1], shown in Figure 1, which premiered in 1992 at the SIGGRAPH conference in Chicago, had its motivation rooted in scientific visualization—and the challenge became one of attracting leading-edge computational scientists and engineers to adopt it. VR, then a fledgling technology, had to help scientists get to scientific discoveries faster without compromising the color, resolution, and flicker-free qualities they had come to expect using workstations. Scientists had been doing single-screen stereo graphics for more than 25 years; any VR system had to successfully compete. Most important, the VR display had to couple remote data sources, supercomputers, and scientific instrumentation in a functional way. In total, the VR system had to offer a significant advantage to offset its packaging and cost.

In 1992, EVL’s goals were to provide an infrastructure for computer scientists to work with computational

scientists to collect, maintain, develop, distribute, and

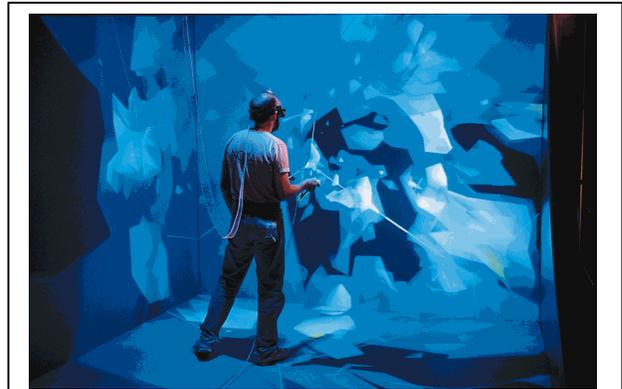


Figure 1: EVL has developed several VR projection-based displays, most notably the CAVE

evaluate VR tools and techniques for scientific computing. These included computer-based models, simulators, data libraries, programming libraries, and user interfaces. EVL also wanted to provide researchers with access to the best VR equipment to explore their data in highly immersive, highly interactive environments.

EVL partnered with the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign and the Mathematics and Computer Science Division of Argonne National Laboratory (Argonne) in a multi-year, ongoing effort to develop highly leveraged national collaborations at major professional conferences—notably SIGGRAPH and Supercomputing [2,3]. These events emphasized high performance computing and communication technologies, VR, and scientific visualization. The overall purpose was to encourage the development of teams, tools, hardware, system software, and human interface models on an accelerated schedule to enable national-scale, multi-site collaborations applied to National Challenge and Grand Challenge problems.

At least six companies now make projection-based VR displays (there were none in 1992), so now EVL turns its attention to developing the collaborative software tools needed for Tele-Immersion.

In the past two years, progress has been made toward linking CAVEs and ImmersaDesks™ with networks, from ISDN to ATM. CALVIN [4], a project first tested cross-country at Supercomputing '95, focused on collaborative VR. It led to NICE [5] (Narrative, Immersive, Constructionist/Collaborative Environments), which in turn led to a design for CAVERNsoft, the distributed memory software architecture for CAVERN, the CAVE Research Network. The CAVERN Users' Society (CAVERNUS) was formed (www.ncsa.uiuc.edu/VR/caverus) to give a voice to the scientists, engineers, and artists who have VR devices as they link their devices together over the vBNS and other networks.

The NICE project is a distributed collaborative virtual environment for K-6 students. It uses the concept of a continuously running simulation, in this case a garden, and allows children and teachers to interact with each other in planning and discussing the garden's evolution over weeks and months. NICE can automatically generate Web-based printable reports, in essence a narrative transcription, that documents the interactions. Students can check the garden's progress with a normal workstation if they are away from the VR devices. NICE incorporates many features that scientists desire in a Tele-Immersive intelligent data mining application, but in a sufficiently limited context to facilitate formal evaluation.

The concept of a substantial simulation running continuously that adult scientists can enter and explore alone or in collaborative sessions is very appealing. EVL aims to lower the barriers to this kind of computer usage by extending the NICE approach to test domains where significant computation and data mining are needed. EVL intends to provide the computation required by adopting Condor, software from the University of Wisconsin, that allows unused processors to be allocated automatically to background jobs [11]. At EVL, this means dozens of Onyx processors that are idle when the CAVEs are not being used. This will allow codes to be run that can benefit from computing, disk, and memory on the order of 20 CRAY Y-MPs operating nightly. EVL hopes scientists subsequently adapt these techniques on computing and visualization resources they have available locally or at supercomputing centers. Such intensive use of CAVE time is difficult to accommodate on EVL's research and development CAVE, which evolves rather often, so in the past year, UIC built a second CAVE for applications development and testing.

2. The Heart of Tele-Immersion: Advanced Networking

The networking research community considers Tele-Immersion to be an "application driver" to help fine-tune networks for large datasets and/or multimedia, where latency and bandwidth are of major importance.

UIC has participated in building the following networks:

- The Metropolitan Research and Education Network (MREN) is a 155Mb regional network connecting

Chicago-area research institutions, is managed by Ameritech, and is a model for "GigaPoPs," or other regional networks. It came into being as part of the I-WAY experiment and a desire for the University of Chicago, UIC, Argonne, Northwestern University, and Ameritech to create a regional high-bandwidth ATM network. MREN now connects the Universities of Wisconsin, Michigan, Michigan State, Minnesota, Illinois at Urbana-Champaign, and Indiana. (www.mren.org)

- UIC was one of the first vBNS Connections Program grantees to achieve a vBNS connection. (www.vbns.net) The vBNS is anticipated to be the first national network linking regional GigaPoPs as part of Internet2.
- STAR TAP (the Science, Technology And Research Transit Access Point) anchors the vBNS International Connections Program. Connectivity to STAR TAP is through the same Ameritech facility as MREN and the vBNS. The Canadian Network for the Advancement of Research, Industry, and Education (CANARIE) and the Singapore Internet Next Generation Research and Education Network (SINGAREN) are connected to STAR TAP. Connectivity with universities in Japan and Taiwan is imminent. Other consortia and countries have expressed interest in connecting. The DOE's ESNET and NASA's NREN networks connect to STAR TAP and vBNS authorized institutions. STAR TAP allows U.S. institutions to link with CAVEs in Europe, Japan, and Singapore, as well as anticipated ones in Canada, Korea and Taiwan, motivating research into latency-tolerant and asynchronous (recorded) Tele-Immersion, because the distance and time shifts are so great. (www.startap.net)
- EVL is deeply involved with the network research community as a partner in the National Computational Science Alliance, NCSA's new organization as part of the NSF Partnerships for Advanced Computational Infrastructure (PACI) initiative (alliance.ncsa.uiuc.edu).

EVL has been very fortunate that many of the early adapters of CAVE and ImmersaDesk technologies, introduced to it by responding to national Calls for Participation as part of SIGGRAPH and Supercomputing conference events organized by EVL, wanted to continue their collaborations afterwards. These computational scientists and engineers were gaining experience with VR technologies and wanted to turn anecdotal testimonies of their benefits into publishable scientific discoveries. Subsequently, to handle the needs of these "friends," EVL started what has become a modestly successful commercialization effort.

There are CAVEs and ImmersaDesks worldwide, as detailed in Table 1. Most buyers expressed interest in either continuing joint research projects or in starting new ones. This presents a wonderful opportunity to solve real-world problems, not only locally at UIC, but

internationally, where other problems—notably networking latency and bandwidth considerations

become important.

Sites Owning Projection-Based CAVEs and ImmersaDesks or their Derivatives							
United States Sites							
Site	Likely Network Connectivity	CAVE	Immersa-Desk	Site	Likely Network Connectivity	CAVE	Immersa-Desk
1. Boston U	vBNS		2	9. U Illinois-NCSA	vBNS, DREN, MREN	1	2
2. Indiana U	vBNS, MREN	1	1	10. U Illinois	vBNS, DREN, MREN		1
3. Iowa State U	vBNS	1 (C2)		11. U Iowa	vBNS		2
4. Michigan State U	vBNS		1	12. U Michigan	vBNS, MREN	1	
5. Old Dominion U	vBNS		1	13. U Utah	vBNS		1
6. Purdue U	vBNS		1 (clone)	14. U Wisconsin	vBNS, MREN		1
7. U Chicago	vBNS, MREN		1	15. Virginia Tech	vBNS	1	1
8. U Illinois Chicago	vBNS, MREN	2	6	16. Wright State U	vBNS	1	
United States Government Research Laboratories							
Site	Likely Network Connectivity	CAVE	Immersa-Desk	Site	Likely Network Connectivity	CAVE	Immersa-Desk
1. Argonne National Laboratory	vBNS, MREN, ESNET CAIRN	1	4	6. Naval Research Lab Washington DC	DREN	1 (C2)	
2. Army Research Lab	DREN		1	7. NIST, Gaithersburg,	Internet		1
3. DARPA, Washington DC	DREN	1		8. U.S. Army Corps Engineers Waterway Exp. Station, MS	DREN		1
4. ISI, Marina del Rey	DREN, CAIRN		1	9. WP AFB	DREN		1
5. NASA/ U Houston	NREN	1					
United States Industry Sites							
Site	Likely Network Connectivity	CAVE	Immersa-Desk	Site	Likely Network Connectivity	CAVE	Immersa-Desk
1. CTC, PA		1		4. SGI/Cray, MN	Internet		1
2. EDS, Detroit, MI		1		5. SGI, Mountain View	Internet		1
3. General Motors	Internet	2					
International Sites							
(Note: Singapore vBNS approved, other connectivity proposed)							
Site	Likely Network Connectivity	CAVE	IDesk	Site	Likely Network Connectivity	CAVE	Immersa-Desk
1. Holland – SARA	STAR TAP	1		10. Japan – Sega		1 (clone)	
2. Austria – Ars Electronica		1	1	11. Japan – Tohwa U	STAR TAP (APAN)		1
3. Germany – GMD	STAR TAP	1 (C2)		12. Japan – Tokyo Inst Tech	STAR TAP (APAN)	1 (clone)	
4. Germany – VISLAB FhG IAO	STAR TAP	1 (C2)		13. Japan – U Tokyo	STAR TAP (APAN)	1 (clone)	
5. Germany – VVR GhG IDG		1 (C2)		14. Singapore – NUS	STAR TAP (direct connection)	1	
6. Japan – NIFS		1		15. Sweden – Helsingborg	STAR TAP	1	1
7. Japan – Nippon Inst Tech	STAR TAP (APAN)	1 (clone)		16. Sweden – RIT	STAR TAP		1
8. Japan – Nissho		1	1	17. Sweden – SICS	STAR TAP	1 (clone)	
9. Japan – NTT/ICC		1					

C2 = Iowa State designed CAVE clone/GMD produced Clone = Self-built CAVE-like or Immersa-Desk-like setup

Table 1: Sites Owning Projection-Based CAVEs and ImmersaDesks or their Derivatives.

3. CAVERN: A Research Agenda for Distributed Virtual Environments

EVL's current research focus is the design and implementation of CAVERN, the CAVE Research Network. CAVERN will be comprised of dozens of network-connected CAVEs, ImmersaDesks, and other VR devices, like Head-Mounted Displays, Responsive Workbenches, and so on. CAVERN will be managed by the CAVE libraries and CAVERNsoft, a distributed shared memory software package optimized for collaboration [6,7].

A truly excellent collection of papers on DVEs can be found in the March 1997 issue of IEEE Spectrum, one that provides an overview and critique of the current research in the field. Jason Leigh's Web site (www.evl.uic.edu/spiff/ti) has an excellent bibliography and papers that discuss how CAVERNsoft differs from the Distributed Interactive Simulation (DIS) implementation and other CVE approaches.

Tele-Immersion has entered the Next Generation Internet (NGI) (www.ngi.gov) and Internet2 (www.Internet2.edu) vocabulary.

In the applications section of the Computing Research Association's "Research Challenges for the Next Generation Internet," five key enabling technologies were identified as common to the future use of the NGI [10]:

- **Database Access.** Gigabyte, terabyte, and petabyte databases, perhaps distributed, will need to be mined by sophisticated search and analysis algorithms.
- **Audio and Video.** Streams of radically different bandwidths, from compressed voice to high-definition television, will be needed to achieve comfortable communication. Low latency/jitter is also important.
- **Real-Time and Delayed Collaboration.** Collaborative technology is essential to support virtual enterprises, collaboratories, desktop videoconferencing, and distance-independent learning. Collaboration requires managed interactions, maintenance of history and audit trails, and proper support of distributed protocols.
- **Distributed Computing.** In the future, more computing will be done using geographically separated elements, either for capacity or access to unique resources. Computational applications will increasingly depend on high bandwidth connections among the computing elements.
- **Tele-Immersion.** Tele-Immersion will enable users in different locations to collaborate in a shared, virtual, or simulated environment as if they were in the same room. It is the ultimate synthesis of networking and media technologies to enhance collaborative environments. Tele-Immersive applications must combine audio, video, virtual worlds, simulations, and many other complex technologies. They will require huge bandwidth, very fast responses, and guarantees of delivery.

3.1 Requirements of Persistent Tele-Immersion

Persistent Tele-Immersion is defined as collaborative computer-generated worlds that remain extant (i.e., continue to exist and evolve) even when there are no human participants. The "world" is a visual representation of a computer simulation.

The requirements to support persistent Tele-Immersion may be viewed at three levels: the end user level, the application developer level and the core technology developer level.

End user issues concern the high-level support of communication, collaborative interaction/manipulation, persistence/recording, and connectivity:

- **Communication in a Tele-Immersive environment.** This typically involves the use of video and, more importantly, audio conferencing. It also includes the use of avatars to provide a representation of the participant in the environment and to convey gestural cues (pointing at objects, waving of the hand, nodding and shaking of the head) that augment audio communication.
- **Synchronous and asynchronous collaboration in the environment.** Persistent environments assume one's ability to work in an environment at the same time with other participants (synchronously), in addition to being able to work in the same environment at different times (asynchronously). Any changes made to the environment must persist for when the other participants return.
- **Recordability.** This denotes the ability to record virtual experiences for later playback or continuation. This also allows participants to place annotations in the environment so that other participants may view them at a later time.
- **Collaborative interaction and manipulation.** Tools to enable participants to interact and manipulate shared objects and datasets being visualized in the environment are needed. These tools must guarantee that proper consistency is maintained across each participant's view of the environment, especially when multiple participants attempt to modify the same shared object.
- **Connectivity with external computational and database resources.** The data that end users want to access in a complex Tele-Immersive environment will likely reside on external databases or be generated by simulation codes on supercomputers or workstation farms. The data may consist of 3D models, collected field data, computed data, persistent recordings of virtual environments, and video and audio streams. Although some simulations may not be initiated in the virtual environment but on a desktop workstation, the simulations will still need to direct the result of computations to a data store that the virtual environment can access. Hence, end users will need

tools to enable them to easily interface their existing non-VR tools/applications with the virtual environment.

Application developer issues revolve mainly around the integration of software tools into applications that meet end-user needs. In the past, we used diverse technologies to build new collaborative applications (such as CALVIN and NICE), which proved highly successful but the underlying toolsets had limited reuse. From these experiences, however, needed software patterns began to emerge. These include:

- Modules for audio and video compression.
- Modules for the rapid construction of a variety of networking topologies and protocols (e.g., TCP, UDP, multicast, and HTTP).
- Control of networking bandwidth, latency, and jitter, referred to as network Quality of Service (QoS).
- Modules for avatar rendering and articulation.
- Connection scalability to allow access to workstation compute clusters, and connectivity of small to large numbers of participants.
- Data scalability to provide homogeneous accessibility to small-to-enormous scientific databases, and the ability to choose the database performance needed to

support the particular data storage task at hand; in essence, the notion of a “database quality of service.”

- Imbedded capabilities for performance monitoring so that an application can query itself from time to time on its performance. For example, if networking latency for a networking channel dedicated to transmitting avatar data begins to increase, the application can be notified so that interpolative techniques may be engaged to maintain the illusion of smooth motion.
- A high-level interface that combines networking with databases so that different Application Programmer Interfaces (APIs) will not have to be learned and made to inter-operate.
- High-level modules containing pre-built environments equipped with minimum Tele-Immersion capabilities, which application programmers may extend to build domain-specific applications.

Core development issues. Application programmers should not have to worry about core technology development; that is, the software and hardware infrastructure to support the expected patterns of resource access in Tele-Immersion. Table 2 summarizes the patterns for networking and databases for a variety of

Patterns for Networking and Databases for a Variety of Tele-Immersion Data Types							
Real-Time Interaction	Volume	Rate of Change	Persistent	Complexity of Data	Database Integrity	Type of Database Transaction	Network QoS
Raw scientific dataset	L	SD	Y	C	R	WO,SA,SU	R
Derived data	MA-L	SCD	Y and N	V	R	SA,SU	R
Avatar data	S	RCD	N	S	-	-	U,L,J
Virtual world state and metadata	S	SCD or RCD	N	V	-	-	R,L,J
3D models	MA-L	SD	Y	V	R	SU	R
Video conferencing data	MA	RCD	N	V	-	-	U,B,L,J
Continuously Recorded	Volume	Rate of Change	Persistent	Complexity of Data	Database Integrity	Type of Database Transaction	Network QoS
Raw scientific dataset	R	SD	Y	C	-	-	R
Derived data	MA	SCD	Y	V	R	SA	R
Avatar data	S	RCD	Y	S	D	RA	U,L,J
Virtual world state and metadata	S	IRCD	Y	V	D-R	RA-SA	R,L,J
3D models	R	SD	Y	V	-	-	R
Video conferencing data	MA	RCD	Y	V	D	RA	U,B,L,J
Snapshot	Volume	Rate of Change	Persistent	Complexity of Data	Database Integrity	Type of Database Transaction	Network QoS
Raw scientific dataset	R	SD	Y	C	-	-	R
Derived data	MA	SCD	Y	V	R	SU	R
Avatar data	S	SCD	Y	S	R	SU	U,L,J
Virtual world state and metadata	S	SCD	Y	V	-	SU	R,L,J
3D models	R	SD	Y	V	-	-	R
Video conferencing data	MA	SCD	Y	V	R	SU	U,B,L,J

Derived data = Data derived from raw datasets
 Metadata = Miscellaneous data coordinating all the other data

L=Large
 MA=Medium atomic
 S=Small
 R=Referenced
 SD=Static data
 SCD=Slow changing data
 RCD=Rapid changing data
 IRCD=Intermittent RCD

Y=Yes
 N=No
 C=Complex
 V=Varies
 S=Simple

R=Required
 D=Desired

SA=Safe appends
 SU=Safe updates
 RU=Rapid updates
 WO=Write once

B=Bandwidth guarantee desired
 R=Reliable
 U=Unreliable
 L=Latency guarantee desired

Table 2: Patterns for networking and databases for a variety of Tele-Immersion data.

Tele-Immersion data:

- Raw data (such as scientific datasets).
- Derived data (data generated from processing and filtering raw data).
- Avatar data (data reflecting position and orientation of the participant's hands and heads in the environment).
- Virtual world state/meta data (data reflecting the state of objects in the virtual world and house-keeping data in general).
- 3D models (derived from polygonizing or rasterizing derived or raw data); video conferencing data (compressed audio and video streams).
- Each of these data categories is further grouped into three access categories (real-time operation, intermittent snapshots, and continuous recording) that reflect the differing database requirements needed to support them. For example, real-time avatar data does not require database operations, snapshots of avatar data require safe writes, and continuous recordings require rapid updates even if it means a reduction in database integrity.

The general conclusion that can be derived from Table 2 is that persistent Tele-Immersion consumes networking and database resources with a number of real-time demands for potentially long periods of time. Patterns for graphics access are not included in the Table, as 30 frames per second is the constant desired frame rate. However, even with the fastest graphics processors currently available (such as the SGI Onyx2 Reality Monster) only a fraction of the data typically present in scientific or engineering datasets can be visualized in real time. Filtering of the data is needed to create meaningful abstractions that can be rendered in near real time.

3.2 CAVERNsoft

In light of the complex interaction of computer graphics, networking, databases, and human factors in Tele-Immersion, developing these applications can be a daunting task. The temptation and common mistake, made by application developers building Tele-Immersive applications for the first time, is to first build a non-collaborative application and then attempt to retro-fit it for Tele-Immersive capabilities.

It is important to provide tools that encourage application developers to envision Tele-Immersive scenarios at a high-level so that they can determine how such capabilities would be most useful in their own applications. A high-level set of tools, however, does not help those trying to retrofit existing applications. A high-level library of well-integrated tools often assumes a specific software design that may be incompatible with the software that is being retrofitted.

To address this issue, a software infrastructure (called CAVERNsoft) that supports the rapid creation of new Tele-Immersive applications and eases the retrofitting of previously non-collaborative VR applications with Tele-

Immersive capabilities, is currently under development.

CAVERNsoft, diagrammed in Figure 2, consists of a central structure called the Information Resource Broker (IRB) surrounded by layers of support software. Although these layers appear to increasingly hide the lower layers from the main application, they are accessible at every level. The lower levels facilitate the construction of new components, as well as the retrofitting of existing applications. The higher levels facilitate the rapid development of new Tele-Immersive applications.

The IRB is a relatively low-level merging of networking and database capabilities that is completely separate from graphics. Hence the basic IRB core can be placed in any software application regardless of its graphics capabilities. This permits graphical applications to communicate with non-graphical applications and it also allows existing non-collaborative applications to achieve networking capabilities with minimal disturbance to their existing graphics.

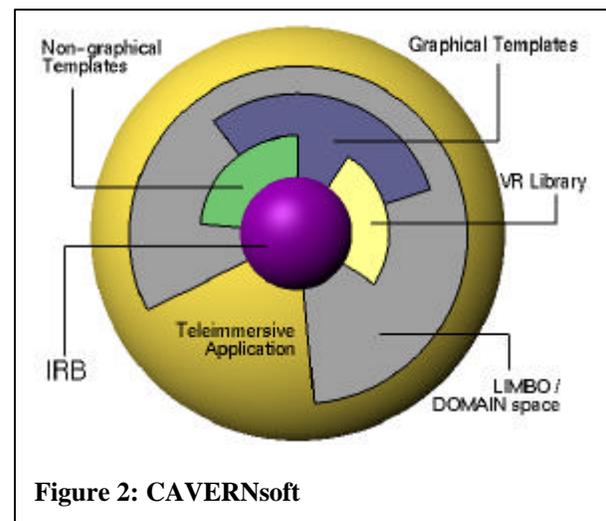


Figure 2: CAVERNsoft

At a layer above the IRB, non-graphical template libraries support the coordination of avatars, as well as audio and video data compression algorithms. On top of this is a higher-level layer that consists of graphical versions of the previous layer, like OpenGL, Performer, and Video avatar templates. These higher level templates can be gathered into even higher-level, fully functional Tele-Immersion spaces called LIMBO spaces.

LIMBO spaces provide varying degrees of avatar rendering and recording; model importing, distribution, manipulation and version control; and audio/video teleconferencing. By using a basic LIMBO space, collaborators can work in a virtual space immediately. They may start with an avatar of their choosing and import 3D models (e.g., car designs, scientific datasets). The LIMBO space will ensure proper distribution of the model to all remote participants. Once objects are distributed, participants may collectively modify them. In addition, application developers may use the well-documented source code of LIMBO space to jump start

the development of their own domain-specific Tele-Immersive applications.

As more domain-specific applications are developed, a growing library of CAVERNsoft-based reusable components (such as collaborative visualization tools) will emerge. These can be added to the library of existing templates and collected to build DOMAIN spaces that are specializations of LIMBO spaces. For example, this will enable a designer to build a Tele-Immersive design application by starting with an existing DOMAIN space equipped with collaborative tools specifically for collaborative design, rather than starting from the basic LIMBO space.

3.3 The Information Resource Broker

The IRB is a small, highly configurable nucleus that can be imbedded into every distributed software system that is expected to interoperate. That is, every client or server application will possess an IRB that it will use to communicate with other compatible applications.

The IRB is spawned as a lightweight thread, with both integrated networking and database capabilities, so that it can serve both client and server needs. Any application using an IRB is automatically a client and a server at the same time. As a client, the application may connect to other servers or clients to access their information resources. As a server, the application may accept connections from other clients. This dual capability is transparent to the application developer but allows the construction of a wide variety of distributed topologies. Its symmetry allows distributed applications to treat each other as information resources.

3.4 Customized Networking in the IRB

IRB-based applications communicate with each other by establishing single or multiple communications channels. Each channel is individually customized to meet the specific needs of the Tele-Immersive data being transmitted. The application may negotiate with the IRB interface for the desired networking protocol (reliable TCP, unreliable UDP, unreliable multicast) and the desired bandwidth, latency, and jitter. These QoS capabilities will be provided by Nexus, a multi-threaded communications library developed by Argonne National Laboratory [8].

3.5 Customized Database in the IRB

After communications channels have been established, each application may create an arena of data that can be linked to one another so as to emulate a distributed shared memory (DSM) segment between the connected applications. Using the same channel or different channels, many arenas can be created, each of which can be used to store different forms of data (avatar tracking, 3D models, etc.) As a default capability, any time one arena is updated, any other applications that have arenas linked to the modified arena will also see their arenas updated. This default coherence protocol can be altered

so that updates occur only when requested. A locking mechanism is provided to support entry consistency.

Besides providing customized communications channels between IRBs, another key feature that distinguishes the IRB from traditional notions of DSM is that each arena can be made persistent; i.e., the IRB may commit an arena to a persistent store. This is useful for Tele-Immersive applications; e.g., the arena could hold a 3D model that was downloaded from a remote application, and rather than downloading the model each time the application is restarted, the model could be retrieved from the persistent arena. Since each commit of an arena is marked with a time stamp, data modified by a remote application prior to a re-start would be downloaded once the local and remote arenas were re-linked.

The persistent capabilities of the IRB are provided by a database interface layer. This layer will eventually allow the applications to negotiate the QoS of the underlying database system. That is, an arena may choose to have safe updates or less safe rapid updates; or, it may choose to relax database integrity requirements to achieve speed. For example, in order to record a continuous Tele-Immersive session, rapid updates at the expense of losing database integrity are needed, whereas for storing static snapshots, performance takes second priority over database integrity. Given the desired "database QoS," the database interface will either assign the arena to be managed by the database system best able to meet the requirements of the arena, or re-negotiate with the client-application for a slightly lower QoS.

The development of CAVERNsoft is currently in its infancy. A proof-of-concept prototype of the IRB has been developed that:

- Facilitates the rapid construction of arbitrary distributed topologies;
- Provides support for various networking protocols (including reliable and unreliable, unicast, broadcast and multicast)
- Provides facilities for supporting concurrent programming (both a message-passing and a distributed-shared-memory model)
- Provides support for persistence of small-to-medium-sized data (data that will fit in the main memory of a computer.)

This current alpha release version of the IRB can be downloaded from www.evl.uic.edu/spiff/cavernsoft.

Additional background information can be found in www.evl.uic.edu/spiff/ti.

3.6 Using CAVERNsoft with the CAVE library

The official CAVE library and its variations use similar models for separating computation from graphics rendering. When a CAVE application is launched, the CAVE library first forks off a number of rendering processes (one for each wall of the CAVE) and a tracker

process. These processes, including the main process, communicate with each other using shared memory. This model is somewhat different from the threaded model of CAVERNsoft. CAVERNsoft, when launched, threads off a number of concurrent threads. Threads are similar to separate forked processes in that they allow multiple tasks to operate concurrently. Threads are different from forked processes in that they take less processing time and resources to create. Also, they share the same address space as the main program, so there is no need to use shared memory for inter-process communication.

For the CAVE forked model and the CAVERNsoft threaded model to work together, CAVERNsoft can only be initiated after the CAVE library has forked off its many processes (after the CAVEInit() call).

CAVERNsoft can only be used in the CAVE's main processing loop, as it has no scope in the separate rendering processes. This is acceptable because the rendering processes should be dedicated solely to rendering tasks anyway.

As CAVERNsoft is a threaded library (it uses Posix Threads for its implementation), the mechanisms for mutual exclusion are incompatible with those used by the CAVE library. The CAVE library uses shared memory locks that are part of SGI's shared memory library. The CAVERNsoft library uses the Pthreads mutual exclusion mechanism. The CAVE library provides a lock management API to allow applications to create read and write locks that are shared across processes (including the rendering processes). CAVERNsoft provides a lock management API that is a simplification of the Posix standard that allows locking within the main process and between threads. Hence, as long as the application uses CAVERNsoft's locking mechanism to lock data within the scope of the CAVE main process but uses the CAVE's locking mechanism to lock data between the main process and the rendering processes, there should be no conflict.

Ultimately, a solution is to move the CAVE library to a threaded model. Since operating systems such as Windows NT use a threaded model, creating a threaded CAVE library will also simplify the port across multiple platforms, especially when a uniform threaded API such as ACE [9] is used. At this time, however, the main computing platforms for the CAVE are SGI computers which are only beginning to provide stable Pthread implementations. The current operating systems offer little control for scheduling threads across multiple processors. This capability is crucial for the CAVE because the CAVE library tries to dedicate independent processors to each forked process.

Since CAVERNsoft does not depend on any part of the CAVE library, it can easily be run on VR devices that may not be running the CAVE library.

4. Conclusions

EVL develops software tools and hardware technologies

with the end user in mind—notably, the computational scientist and engineer. EVL's goal is to make its VR displays fully functioning “discovery” resources for the computational science and engineering community. Research activities, in addition to technology development, include: quantitative assessment, human-centered intelligent interfaces, improved real-time imagery, fast-access methods to very large databases, distributed computing, real-time navigation and successive refinement with large-scale computational models, and high-definition VR. Now that commercial vendors sell CAVEs and ImmersaDesks, as well as similar projection-based VR displays, EVL's has shifted its attention from technology development to Tele-Immersion.

5. Acknowledgments

EVL's VR research, collaborations, and outreach programs would not be possible without major funding from the National Science Foundation (NSF), the Defense Advanced Research Projects Agency (DARPA), and the U.S. Department of Energy. This material is based upon work supported by NSF under grants CDA-9303433, CDA-9512272, ASC-9418068, NCR-9712283, CDA-9720351, and the ASC Partnerships for an Advanced Computational Infrastructure (PACI) program.

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

6. References

- [1] C. Cruz-Neira, D.J. Sandin, T.A. DeFanti, “Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE,” *Computer Graphics (Proceedings of SIGGRAPH 93)*, ACM SIGGRAPH, August 1993, pp. 135-142.
- [2] T.A. DeFanti, I. Foster, M. Papka, R. Stevens, and T. Kuhfuss, “Overview of the I-WAY: Wide Area Visual Supercomputing,” *The International Journal of Supercomputer Applications and High Performance Computing*, Vol. 10, No. 2/3, Summer/Fall 1996, pp. 123-131.
- [3] T.A. DeFanti, M.D. Brown and R. Stevens (Guest Editors), “Virtual Reality Over High-Speed Networks,” *IEEE Computer Graphics & Applications*, July 1996, Vol. 16, No. 4, pp. 42-84.
- [4] J. Leigh and A. Johnson, “Supporting Transcontinental Collaborative Work in Persistent Virtual Environments,” *IEEE Computer Graphics and Applications*, Vol 16, No. 4, July 1996, pp. 47-51.
- [5] M. Roussos, A. Johnson, J. Leigh, C. Barnes, C. Vasilakis, and T. Moher, “NICE: Combining Constructionism, Narrative, and Collaboration in a Virtual Learning Environment,” *Proceedings of the 1997 SIGGRAPH Educator's Program*, *Computer Graphics*, August, 1997, pp. 62-63 (and cover image).
- [6] J. Leigh, A.E. Johnson, T.A. DeFanti, “Issues in the Design of a Flexible Distributed Architecture for Supporting Persistence and Interoperability in

Collaborative Virtual Environments,” Proceedings of SC’97, IEEE/ACM Supercomputing, November 1997, to be published.

- [7] J. Leigh, A. Johnson, and T. DeFanti, “CAVERN: Distributed Architecture for Supporting Scalable Persistence and Interoperability in Collaborative Virtual Environments,” *Virtual Reality: Research, Development and Applications*, Vol 2, No. 2, December 1997, pp 217-237.
- [8] I. Foster, C. Kesselman, and S. Tuecke, “The Nexus approach to integrating multithreading and communication,” *Journal of Parallel and Distributed Computing*, No. 37, 1996, p.70-82.
- [9] D.C. Schmidt, T. Harrison, and E. Al-Shaer, “Object-oriented components for high-speed network programming,” *Proceedings of USENIX Conference on Object-Oriented Technologies*, Monterey, CA, 1995
- [10] J.E. Smith and F.W. Weingarten (eds.), *Research Challenges for the Next Generation Internet*, Computing Research Association, 1997, p. 20.
- [11] M. Livny, et al., A worldwide flock of Condors: Load sharing among workstation clusters, *Journal Future Generations of Computer Systems*, Volume 12, 1996.