TELE-IMMERSIVE VIRTUAL ENVIRONMENTS FOR COLLABORATIVE KNOWLEDGE DISCOVERY

Cathy M. Lascara and Glen H. Wheless
Virtual Environments Laboratory
Old Dominion University
768 W. 52nd St.
Norfolk, VA 23529
lascara@ccpo.odu.edu

Donna Cox, Robert Patterson, Stuart Levy
National Center for Supercomputer Applications
University of Illinois, Urbana-Champaign
405 N. Mathews Ave.
Urbana, IL 61801
cox@ncsa.uiuc.edu

Andrew Johnson, Jason Leigh, Ahbinov Kapoor
Electronic Visualization Laboratory
University of Illinois at Chicago
851 S. Morgan St., Rm. 1120
Chicago IL 60607-7053
ajohnson@eecs.uic.edu

Keywords: Virtual Reality, Tele-Immersion, Interactive Systems, Environmental Data

ABSTRACT

This paper describes the design and implementation of two tele-immersive applications, CVD and Cave6D, both designed to support collaborative knowledge discovery from large multidimensional datasets. CVD integrates the capabilities of two existing VR applications, Cave5D and Virtual Director, in order to provide immersive experiences of distributed data using high performance networks and interactive hardware and software. Cave6D is similar in function yet is tightly integrated with the CAVERNSoft toolkit so as to provide access to remote computational platforms and databases.

INTRODUCTION

In recent years there has been a rapid increase in the capability of environmental observation and modeling systems to provide high resolution spatial and temporal data about the world around us. It is not enough, however, to be able to collect, generate or share large amounts of data. Scientists, educators, students and managers must have the ability to collaboratively view, analyze and interact with the data in a way that is understandable, repeatable and intuitive.

Concurrent advances in visualization and computational capabilities have spurred the creation and use of realistic, three-dimensional virtual environments for interaction with and analysis of this data. Viewing and interacting with these large multivariate datasets in a georeferenced virtual environment provides an actual sense of presence that inherently changes the way the data is analyzed (Wheless *et al.*, 1995) thereby aiding in the mental process of assimilating complex information. Users are able to navigate through, view and interact with the data in a fully three-dimensional context, thus preserving necessary geospatial relationships crucial for intuitive analysis.

High performance networks, such as the NSF-funded very High Bandwidth Network Service (vBNS), now provide an infrastructure for low latency multi-user access to these virtual environments (DeFanti *et al.*, 1996). The use of these Collaborative Virtual Environments (CVEs) enable users at many distributed sites to interact with each other and with the data in a many-to-many session from within a common virtual world. Linking remote users into a shared virtual space enables a high level mode of collaboration that is conducive to knowledge discovery (NSTC, 1998).

We describe in this paper our efforts to integrate this collaborative capability into an existing VR-based scientific visualization application, Cave5D (Wheless *et al.*, 1996). Our objective is to enable multi-user immersive visualization of large multi-dimensional datasets in support of oceanographic, atmospheric and terrestrial scientific studies. Two separate implementation paths were followed, each resulting in a distinct prototype. In one case, Cave5D was augmented with remote interaction techniques and camera choreography capabilities provided by the VR application Virtual Director. In the other case, Cave5D was retrofitted with collaborative features provided by the CAVERNSoft toolkit (Leigh *et al.*, 1997), a software library that is designed to allow for easy development of similar applications.

TELE-IMMERSION

We use the concept of *Tele-Immersion* as the fundamental underpinnings of our work. We define tele-immersion as the union of audio and video conferencing, networked collaborative VR and image-based modeling in the context of significant computing and data mining. Tele-immersion enables users in different locations to collaborate in a shared, virtual, or simulated environment as if they are in the same room. It is the ultimate synthesis of networking and media technologies to enhance CVEs.

The development of tele-immersive CVEs is currently one of the most challenging areas of research in Virtual Reality (VR) because new dimensions to human-factors, networking, and database issues must be explored. For example, human-factors research in VR has traditionally focused on the development of natural interfaces for manipulating virtual objects and traversing virtual landscapes. Collaborative manipulation, on the other hand, requires the consideration of how participants should interact with each other in a shared space, in addition to how co-manipulated objects should behave. Other issues include: how participants should be represented in the collaborative environment; how to effectively transmit non-verbal cues that real-world collaborators so casually and effectively use; how to best transmit video and audio via a channel that allows both public addressing as well as private conversations to occur; and how to sustain a virtual environment even when all its participants have left. Many of these issues were explored in our development work.

Although there are a broad range of devices that support tele-immersion activities, our work focuses on projection-based graphical devices such as the Immersadesk or CAVE™ (Cruz-Neira *et al.*, 1993). Growing in number since 1992 from just a few at selected research institutions, the number of Immersadesks, CAVEs™ and CAVE-like devices now number well into the hundreds. Stereographic LCD shutter glasses and head/hand tracking is often used to augment the experience. The CAVELibrary software API (please see http://www.ncsa.uiuc.edu/VR/cavernus/) is an example of a common underlying interface between CAVE-specific applications, serving to coordinate all ancillary devices (eg: LCD glasses, navigational wand), stereo transformations and program synchronization.

The Immersadesk, a drafting-table format, projection-based device with a single 4` by 5` angled screen, enables users to view and interact with a CVE in a semi-immersive fashion. This type of semi-immersive VR is useful for applications that do not require full immersion of the user into the virtual environment, such as a 3D model of a bio-molecular simulation. However, full immersion is more useful for realistic portrayal and useful interaction in very large-scale virtual environments that also include small man-made objects and fine-scale environmental features.

The CAVE™ allows for such a full immersion. The CAVE™ is a 10x10x10-foot structure consisting of rear-projected screen walls and a front-projected floor. Images appear to float in space, with the user free to "walk" around them, yet maintain a proper perspective. This information can be much richer and more accurate, with finer granularity of many more variables, and with sound as well as visual depictions.

TELE-IMMERSIVE APPLICATIONS

Linking these VR devices together in dedicated tele-immersive environments is an area of active research. Tele-immersive applications not only combine audio, video, virtual worlds, simulations, and many other complex technologies, but they also require huge bandwidth, very fast responses, and guarantees of delivery. We now describe two tele-immersive applications, CVD and Cave6D, both stemming from one of the first VR applications for interaction with large datasets in the CAVE™, Cave5D.

Cave5D

Cave5D is a configurable VR application framework that integrates the CAVELibrary VR software API with the Vis5D visualization API (Hibbard *et al*., 1992) in order to visualize large multivariate data sets in the CAVE or on the ImmersaDesk and to enable user interaction with the data.

The representation of data within a virtual world involves three primary steps: the conversion of data into a graphical objects, the mapping of the graphical objects into the physical space of the virtual world, and defining interactions with the graphical objects (and thus, by association, with the data). Cave5D performs these functions in addition to managing the data sets, displaying the graphical objects, and specifying the space-time definition of the virtual environment.

Large multi-dimensional numerical data sets from atmospheric, oceanographic, and other similar models, are visualized in Cave5D using isosurfaces, contour slices, volume visualization, wind/trajectory vectors, and various image projection formats. Users are able to navigate through the data, change the view of the entire data volume and interactively slice through the data set of interest while within the virtual environment. Navigation of a data set, as defined here, is the steering of the VR user's view through the graphical display of multi-dimensional data.

Virtual Director

Managing and documenting visualizations from large scientific datasets can be an unwieldy and difficult task. When attempting to spatially navigate and temporally record a large-scale data set, time and space become non-trivial to the VR applications manager. The tracking of simulation time (runtime, interpolated, or non-linear), real clock time, VR display/recording/playback time, and digital frame time, is a fundamental problem in the design of a general navigation and recording system for VR scientific applications. A simulation often produces large data sets that are stored, and then post-processed and visualized with interactive VR visualization packages. To add to this complexity is the effort to convert an interactive VR exploration from a real-time experience into a movie or video.

The development of Virtual Director partially solved these time management, documentation and non-linearity problems. Virtual Director is a software framework that enables real-time VR data exploration and camera choreography for creating animations. Current functionality to interactively construct and preview visualizations include voice and gesture input, navigation and scale, camera control, camera path creation, playback and display controls, and session configuration and file management. Extremely large scientific data sets are sub-sampled for efficient real-time interaction.

Video or higher-resolution rendering of full, non-sampled data sets is often required by scientists to record or "document" the simulation or interactive session. To interface real-time interaction with animation production, Virtual Director provides standard format camera output that is compatible with several renderers (e.g., Pixar's Particle Renderer, Wavefront). The current API reads and interactively renders particle data; the interactive particle renderer can output animation frames directly to disk from Virtual Director.

CVD

The first tele-immersive prototype application we describe, CVD, is the result of an effort to lessen the effort required to explore large datasets in a collaborative fashion. CVD, or Cave5D/Virtual Director, integrates the capabilities of both existing VR applications, Cave5D and Virtual Director, in order to enable the user to view and interact with the data from within the data set, visualize the data in real time and easily record the experience. CVD has a powerful remote collaboration capability based on the CAVE-to-CAVE networking functions of the CAVELibrary API so that multiple users at many remote locations are able to participate in useful immersive virtual experiences. Remote interaction is enabled by transmitting tracker data, voice, video, and navigation information, thereby enabling scientists to collaboratively work with other scientists or visualization professionals using CVD at other locations.

Avatars representing each user show the location and actions of the user within the CVE. Wand position and orientation is displayed as a virtual hand, thereby providing gestural signals to all other users. All avatars in the CVE may have personalized avatar shape and color, and a digital image of the user overlaid on the avatar shape provides more personalization. The use of multicast audio with the

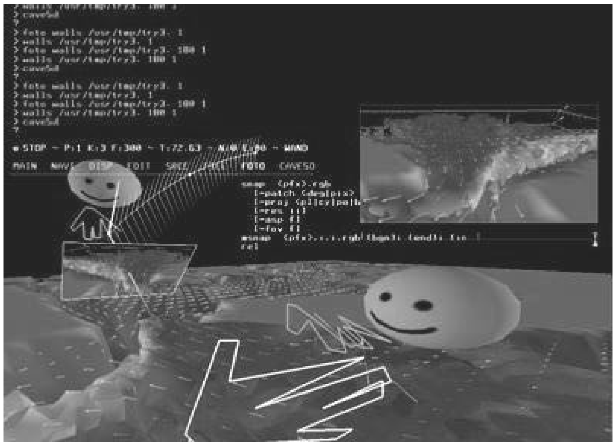

Figure 1 View into CVD showing data, avatars, virtual hand, menu items and camera spline path.

freely available via application permits verbal communication among all CVE users. Techniques have been developed to share CVD current status information and camera spline path information so that users at remote sites can navigate and control the virtual camera.

The power of Cave5D using Vis5D graphical display techniques in the immersive environment of the CAVE or Immersadesk has thus been greatly leveraged by the functionality of Virtual Director. Animations and images of temporally varying behaviors in several different environmental datasets made from Cave5D/Virtual Director may be viewed at http://www.ncsa.uiuc.edu/VR/grants/virdir. CVD has been demonstrated several times, most recently at Supercomputing'98 in Orlando, FL. (Wheless *et al*., 1998b) in which several users spread over thousands of miles were able to interact in a shared virtual environment.

CAVERNSoft And Tele-Immersion

Naturally, tele-immersive CVEs pose new challenges as they require an unconventionally broad range of networking, database and graphics capabilities. This vast range makes the rapid construction of complex CVEs difficult. Most attempts at building networking and database architectures for CVEs have resulted in ad-hoc solutions that are specifically designed to solve a small range of problems. The CAVERNSoft programming environment was designed to provide a systematic solution to these challenges (Leigh *et al*., 1997).

The CAVERNSoft framework employs an Information Request Broker (IRB) as the nucleus of all CAVERN-based client and server applications. An IRB is an autonomous repository of persistent data driven by a database and accessible by a variety of networking interfaces in order to provide a tightly coupled mechanism to automatically transfer messages. Combining the distributed shared memory model with active database technology and real-time networking technology, this hybrid system allows the rapid construction of arbitrary topologies of interconnected, tele-immersed participants.

CAVERNSoft has the added ability of supporting the feature of *persistence*. That is, even when all the participants have left the environment and the virtual display devices have been switched off, the environment continues to evolve. This evolution would be caused either by behaviors previously set in motion by the participants or by dynamics from a model and/or observations from the real world that the CVE is simulating. At work might be either a central server or a network of distributed servers that continuously sustain the environment.

Cave6D

Similar in function to the CVD application described above, Cave6D is a configurable VR application based on Cave5D, the Vis5D graphics API and the CAVELibrary API. Cave6D is distinguished from CVD by its linkage with the CAVERNSoft toolkit to provide collaborative capabilities.

The CAVERNSoft IRB is the backbone of all the collaborative abilities of Cave6D, providing a unified interface to all the networking and database needs and supporting the distribution of data across the clients. The client/server at each Cave6D site has its own 'personal' IRB that is used to transparently manage sharing and caching of data retrieved from other IRBs. The data retrieval process is begun by creating a channel between the personal IRB and the remote IRB and declaring its properties. A number of uniquely identified keys are then created and linked across the channel thereby providing a handle to the data storage location in the IRBs' database. When new data arrives, callbacks are triggered to transmit the data to the remote key through the link. Any information received by a key is automatically propagated to all the other linked keys.

Multiple users are thus able to jointly visualize, discuss and interact with the data from within a shared environment. Avatars are used as representations of each user and possess long pointing rays to highlight areas of interest. All visualization objects may be remotely operated by any avatar in a global interaction mode or in a local (single site) mode. This allows each participant to customize the environment for optimum viewing based on location within the CVE. A conscious decision was made to keep time globally invariant across each Cave6D site so that temporally varying fields are viewed synchronously. This constraint forces a shared field of reference on all participants even with all other variables locally selectable.

Cave6D has been demonstrated at several venues, including the 1998 Next Generation Internet meeting held in Washington, DC, the 1998 Internet2 Spring Members meeting, and at Supercomputing 1998 as part of the iGrid project.

CONCLUSIONS

We have described our work developing and using two tele-immersive applications for knowledge discovery with large multivariate datasets. These applications allow raw data to become usable information and then cognitively useful knowledge for multiple users in many locations. The design of these collaborative applications incorporates multiple data streams from archived data stores, model results and advanced instrumentation Re-architecting both applications to more closely merge capabilities is in progress.

Clearly, the near real-time feedback provided by such collaborative applications have the potential to change the way data or model results are viewed and interpreted, the resulting information disseminated and the ensuing decisions or policies enacted.

ACKNOWLEDGEMENTS

The development of Cave5D is supported by a grant from the NSF New Technologies program (ASC-9634151). The vBNS connection to Old Dominion University is made possible by NSF grant NCR-9613910. This work is a result of research sponsored in part by NOAA Office of Sea Grant, U.S. Department of Commerce, under Grant No. NA56RG0489 to the Virginia Graduate Marine Science Consortium and the Virginia Sea Grant College Program. The U.S. Government is authorized to produce and distribute research reprints for governmental purposes notwithstanding any copyright notification that may appear hereon. The virtual reality research, collaborations, and outreach programs at EVL are made possible through major funding from the National Science Foundation, the Defense Advanced Research Projects Agency, and the US Department of Energy; specifically NSF awards CDA-9303433, CDA-9512272, NCR-9712283, CDA-9720351, and the NSF ASC Partnerships for Advanced Computational Infrastructure program.

REFERENCES

Cruz-Neira, C., J. Leigh, C. Barnes, S.M. Cohen, S. Das, R. Engelmann, R. Hudson, M. Papka, T. Roy, L. Siegel, C. Vasilakis, T.A. DeFanti and D.J. Sandin, "Scientists in Wonderland: A Report on Visualization Applications in the CAVE Virtual Reality Environment," IEEE 1993 Symposium on Research Frontiers in Virtual Reality, October 1993, pp. 59-66.

DeFanti. T. A., M.D. Brown and R. Stevens (Guest Editors), "Virtual Reality Over High-Speed Networks," IEEE *Computer Graphics & Applications*, July 1996, **16(4)**, pp. 14-17, 42-84.

Leigh, J., A. E. Johnson, T. A. DeFanti. "CAVERN: A Distributed Architectture for Supporting Scalable Persistence and Interopeability in Collaborative Virtual Environments", *Journal of Virtual Reality Research, Development and Applications*, December 1997, **22,** pp**.** 217-237.

National Science and Technology Council Committee on Computing, Information and Communications. 1998. **Technologies for the 21st Century**. Complete text available at http://www.ccic.gov.

Stevens, R., P. Woodward, T. DeFanti, C. Catlett, 1997. ``From the I-WAY to the National Technology Grid". *Communications of the ACM*, **40(11)**, pp. 51-60.

Wheless, G. H., C. M. Lascara, A. Valle-Levinson, D. P. Brutzman, W. L. Hibbard, B. Paul and W. Sherman. 1996. ``The Chesapeake Bay Virtual Ecosystem : Initial results from the prototypical system". *International Journal of Supercomputer Applications and High Performance Computing*, **10(2),** pp. 199-210.

Wheless, G. H., C. M. Lascara, D. Cox, R. Patterson, S. Levy, W. Hibbard. 1998. "Cave5D and Virtual Director: Collaborative visualization and management of large, multivariate environmental datasets". Demonstrated at Alliance'98, April 1998, Urbana-Champaign, IL.