The use of collaborative virtual environments in the mine countermeasures mission

Glen H. Wheless [a], Cathy M. Lascara [a], Donna Cox [b], Robert Patterson [b], Stuart Levy [b], Andrew Johnson [c], Jason Leigh [c], Ahbinov Kapoor [c]

[a] Virtual Environments Laboratory, Old Dominion University, 768 W. 52nd St., Norfolk, VA 23529

[b] National Center for Supercomputer Applications, University of Illinois, Urbana-Champaign, Urbana, IL 61801

[c] Electronic Visualization Laboratory, University of Illinois at Chicago, 851 S. Morgan St., Rm. 1120 Chicago IL 60607-7053

ABSTRACT

We describe our work on the development and use of collaborative virtual environments (CVEs) in support of very shallow water mine remediation, mission planning and rehearsal activities. Incorporating multiple data streams, these CVEs allow the user to view, navigate, and interact with data in a 3-D environment, including graphical representation of bathymetry/topography, above-surface images, in-water objects (e.g mines, bridges), and hydrographic characteristics (e.g. currents, water levels, temperature). Asynchronous collaborative capability allows users at many distributed sites to partake in a many-to-many session that takes place in a common virtual world. Cave5D, a tool for immersive visualization of data to support oceanographic, meteorological and CFD studies, has been integrated with the VR application Virtual Director and the underlying collaborative architecture known as CAVERNSoft. Cave5D provides visualization techniques to display multi-dimensional numerical data from atmospheric, oceanographic, and other similar models, including isosurfaces, contour slices, volume visualization, wind/trajectory vectors, and various image projection formats. Virtual Director is a software framework that enables real-time VR data exploration, shared camera choreography between distributed sites, and animation creation capability. CAVERNSoft supports persistent CVEs for collaborative visualizations that involve supercomputers, massive data stores and the integration of additional 'real-time' observations collected by autonomous sensors or swimmer scouts.

Keywords: Virtual environments, visualization, data fusion, mine countermeasures, AUV's

1. INTRODUCTION

Knowledge of the undersea bathymetry and bottom composition, currents, water levels, water mass distribution and other environmental variables in the littoral zone is a tactical necessity for the planning and operational phase of any mission in the littoral zone. 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 littoral zone in support of mine warfare. Historically, this data has been limited to that collected by swimmer scouts, local ships or simple models. Today, much higher resolution data covering larger spatial areas is routinely collected by new observational platforms including high frequency (HF) radar, synthetic aperture radar (SAR), towed undulating vehicles (SCANFISH or SEASOAR), and autonomous underwater vehicles (AUVs). New simulation models now produce more types of data at smaller scales and encompassing longer time and larger space domains. Clearly, environmental assessment of the shallow water regime requires the analysis of large quantities of multidisciplinary data that span a wide range of time and space scales and involve a wide variety of collection modalities.

It is not enough, however, to be able to collect or generate large amounts of data. How the data is presented to mine countermeasures (MCM) or Naval Special Warfare (NSW) forces is often of much more importance. Up until very recently, most of the environmental information gathered from swimmers, moored instrumentation, towed data collection devices, AUVs or remote sensing platforms was presented by portraying it in visual form as a two-dimensional plot or a picture. The development of advanced visualization hardware, interactive software and networking systems have made real-time interactions with three-dimensional graphical representations of data. These advances have spurred the creation and use of realistic large scale virtual environments for interaction with and analysis of this data. Viewing and interacting with these large multivariate datasets in a geo-referenced virtual environment provides an actual sense of presence that inherently changes the way the data is analyzed thereby aiding in the mental process of assimilating complex mission critical information. [1] 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.

As the technological capabilities within the Navy have been improved, there has been a recent thrust toward more organic, innovative operational MCM methods. The organic requirements suggest that geographically distributed MCM forces should have the ability to collaboratively view, analyze and interact with the data in a way that is understandable, repeatable and intuitive. One of the most efficient ways to do this is through the use of a collaborative, shared tele-immersive application constructed from observations, model output, analysis products and various data streams. High performance networks such as DoDs Defense Research and Engineering Network (DREN) or the NSF-funded very High Bandwidth Network Service (vBNS) now provide an infrastructure for low latency multi-user access to these data and other assets. [2]

One of the key enabling technologies crucial to the use of the emerging high performance computational and communications infrastructure now and to its evolution in the future is the *Collaborative Virtual Environment* (CVE), where the appropriate middleware allows users to jointly interact with compute clusters, supercomputers, simulations, large databases and various real time data streams in an intuitive visual fashion independent of location or time. These 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 leading to a higher level mode of collaboration. [3] Multiple ocean, air and space platforms may now be linked into an information grid to enhance shared awareness. Clearly, the near real-time feedback provided by such collaborative virtual environments have the potential to change the way mission critical information is presented and viewed across geographical distances.

We describe in this paper our efforts to integrate this collaborative capability into an existing VR-based scientific visualization application in support of the MCM and NSW mission. Our objective is to enable multi-user immersive visualization of large multi-dimensional datasets. Cave5D, an application for visual analysis of large datasets in a virtual environment, was used as the underlying developmental application. 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, a software library that is designed to allow for easy development of similar applications.

2. 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™. [4] 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.

3. 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.

3.1 Cave5D

Cave5D is a configurable VR application framework that integrates the CAVELibrary VR software API with the Vis5D visualization API 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.

3.2 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

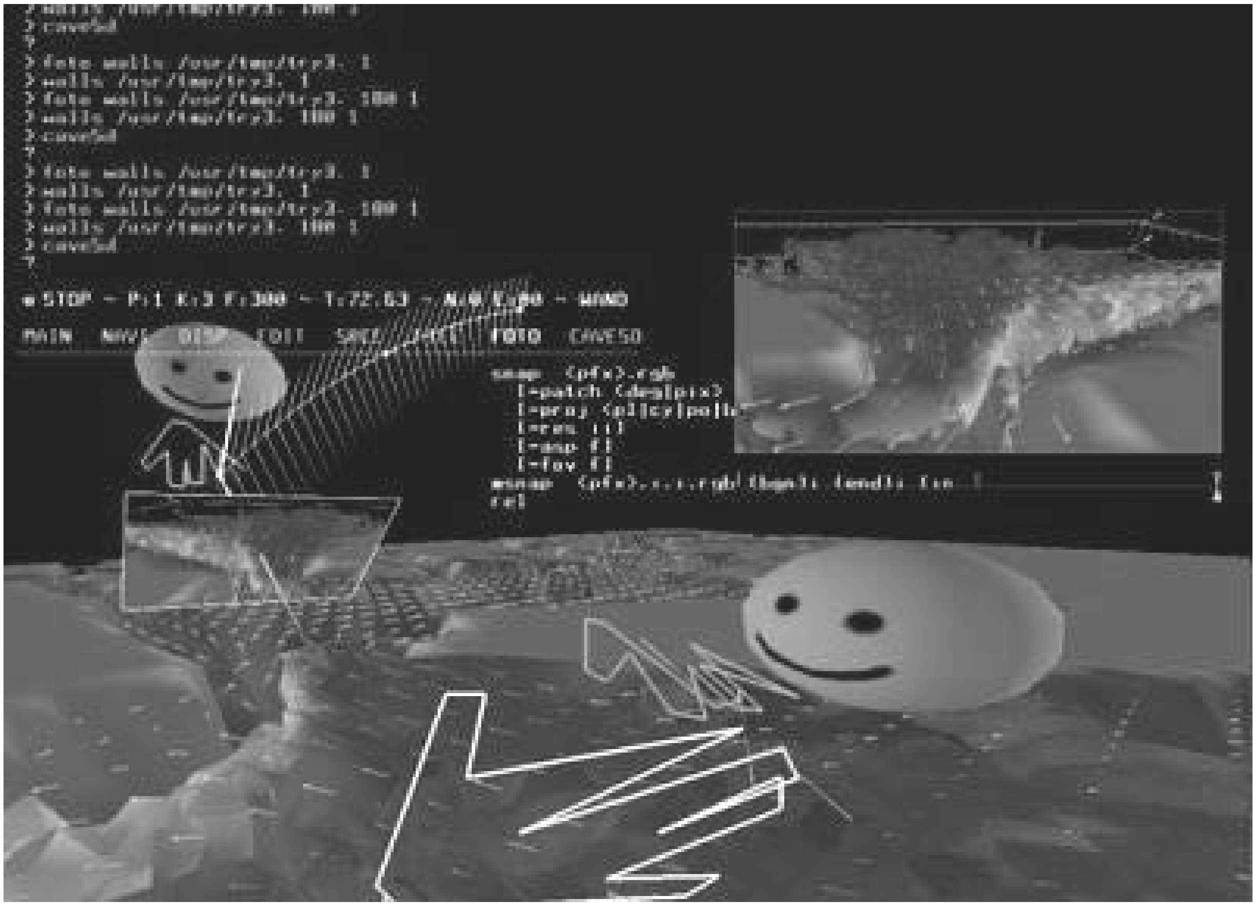

Figure 1 View into CVD showing data, avatars, virtual hand, menu items and camera spline path.

current API reads and interactively renders particle data; the interactive particle renderer can output animation frames directly to disk from Virtual Director.

3.3 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 freely available *vat* 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. in which several users spread over thousands of miles were able to interact in a shared virtual environment.

3.3 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. [5]

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.

3.4 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.

4. A PERSISTENT CVE FOR THE LITTORAL ZONE

We continue the development of an application that portrays environmental information and other data products in a persistent CVE format to aid MCM and NSW operations. Although our original goal was to design such a system for use in mission planning activities, it has subsequently become apparent that this application would be more appropriate for the mission rehearsal phase and during actual operations. The resulting CVEs are normally constructed from observations, model output, analysis products and other data streams. Using the tele-immersive applications described above, the immediate and adjacent domain of the target area of interest are portrayed as a user-navigable, three-dimensional graphical visualization of bathymetry, above-surface images, in-water objects (e.g. mines, bridges), and hydrographic characteristics (e.g. currents, water levels, temperature). The initial information needed to construct the virtual environment are provided by archived data sets which could either be downloadable via standard Navy C4I protocols or provided on CD-ROM. The application also allows for the integration of additional real-time observations collected by autonomous sensors or swimmer scouts. The ability to rapidly create user-configurable scenario-based sessions will play an important role in the rehearsal and process of a specific mission.

CVE evolution would be caused either by behaviors previously set in motion by earlier or current 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. Offline computations running on supercomputers or compute clusters that can take hours to days to complete can be fed incrementally to a CVE during the simulation or can be stored away in secondary storage to be queried at a later time. The server would remain active at all times to allow these data transfers to occur in the absence of the user.

Persistence is also needed to allow collaborators to make recordings of previous synchronous or asynchronous sessions. Recordings allow participants to annotate their environment and to save away versions so that a history of the decisions made over time can be chronicled. A recording would consist of streaming all the avatars' gestural information as well as audio and perhaps video data to a data store, along with any meta-data needed to create a proper context under which the recording was originally made. Hence if an avatar was pointing and speaking about an object during a recording, the re-play should reproduce all the avatar's actions in relation to the object. The re-play of recordings coupled with the ability to create multiple representations that are simultaneously observed by multiple participants.

Asynchronous collaboration increases the opportunities for or is the only practical form of collaborative exploratory data analysis in cases where collaborators are separated by significant time zone differences as often happens in worldwide exercises or naval missions. This is yet another motivation for creating a persistent CVE framework. In such an environment, a persistent server may be commanded to query a large database or gather the output of a remote computation on a supercomputer. CVE participants would then be able to re-enter the environment at a later time to monitor progress, to visualize the results of the computations, or interact with earlier user actions.

Certain annotating capabilities will also be enabled within the CVE tailored for the MCM and NSW mission. These capabilities include the use of 'virtual signposts', where a remote user could place icon markers in the simulation to identify specific locations to navigate, visualize, or render as an animation. Related annotations would be communicated through pop-up 'virtual clipboards' with input from voice, keyboard, or tablet, allowing a participant to directly point at features in the visualization while simultaneously describing them with speech. Other capabilities include 'message bubbles' appearing as unobtrusive reminders in the CVE that would flash or be color coded, 'video inset windows' that display map-related still imagery of critical information, and other iconic type cueing systems. Annotations such as these have been found particularly effective. We have found that participants who viewed annotations would often address the replay as if the remote participant were really there- even though the recording was made hours ago.

Our tele-immersive CVEs now require use of a COTS-based Silicon Graphics Inc. Onyx2 graphics machine capable of real-time 3D graphics display, rapid numerical simulation capability and high-speed, high-bandwidth I/O communications. Although the application is optimized for use with large-screen projection-based VR devices such as the ImmersaDesk or CAVE™ in order to take advantage of the immersive nature of this type of hardware, methods with which to provide the information to low-end platforms are being explored.

5. 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 the Office of Naval Research through grant N00014-98-1-0231. The vBNS connection to Old Dominion University is made possible by NSF grant NCR-9613910. 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

1. G. H. Wheless, 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.

2. R. Stevens, 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.

3. 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.

4. C. Cruz-Neira., 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

5. J. Leigh, 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.