Supporting Transcontinental Collaborative Work in Persistent Virtual Environments

Jason Leigh and Andrew E. Johnson
Electronic Visualization Laboratory
University of Illinois at Chicago

1 Introduction

One of the current long term goals at the Electronic Visualization Laboratory (EVL) is to create a persistent virtual environment enabling multiple transcontinentally-situated participants to apply collaborative Virtual Reality (VR) over high-speed and high-bandwidth networks connected to heterogeneous supercomputing resources and large data stores. Many of the applications in the GII Testbed at Supercomputing '95 focused on connecting supercomputing resources to VR devices to visualize data from specific scientific problem domains[1]. Our focus, however, was on researching the techniques required to support general collaborative work in persistent virtual environments. These techniques include providing: representations for virtual co-presence; video and audio teleconferencing; virtual VR interfaces; and database technology for sustaining persistence in virtual worlds.

These techniques have been showcased in a collaborative design laboratory called CALVIN in which participants at Supercomputing '95 in San Diego and ARPA headquarters in Arlington, Virginia were able to join collectively in an interactive design session.

2 CALVIN

CALVIN (Collaborative Architecture Via Immersive Navigation) is a system that allows multiple participants to design collaboratively in VR. CALVIN employs the concept of heterogeneous perspectives as a paradigm for collaborative interpretations of the design space. Currently the dominant perspective for viewing architectural models in VR is through an egocentric (inside-out) view[2]. This perspective allows a person to see the world as if standing inside the environment and looking out at the surroundings. However other researchers are now beginning to accompany this view with a simultaneous exocentric view in the belief that while an egocentric view affords fine manipulations of the environment, an exocentric view is preferred for gross manipulations[3].

CALVIN includes and extends these perspectives with the following:

- Perspectives that occur when multiple collaborators from perhaps differing disciplines (e.g. an architect and an engineer) offer their opinions on the design. These collaborators, who may be remotely situated, can simultaneously communicate and work in the shared design space.

- Perspectives that are produced as a result of applying differing camera parameters to view and manipulate a design (as in [3].) CALVIN defines an egocentric view as the *mortal's* view and defines an exocentric view as the *deity's* view. This provides a metaphor for defining the relationship between participants in the two views as well as determining their roles in the collaboration. For example, a mortal might be a student in the environment whereas a deity may be a teacher who possesses capabilities that are not provided to the student, such as the ability to alter their scale to alter their viewpoint on the scene.

- Perspectives that are produced as a result of applying information filters that are designed specifically for the tasks performed by individual users [4] (e.g. an engineer viewing the architecture might be more interested in viewing the stress and strain characteristics of the building.)

- Perspectives that occur from design ideas maturing over time in a system capable of sustaining a sense of temporal continuity. That is, the environment persists and evolves over time as new ideas are explored, or old ideas are reused toward the solution of new problems. Ideally collaborators may synchronously and asynchronously enter the environment to perform design work whenever inspiration strikes them.

At present, CALVIN implements only a subset of these concepts. Specifically, CALVIN implements multiple camera perspectives, allows multiple participants to design collaboratively, and communicate via voice and video, in a shared architectural space that provides a simple persistence mechanism to track evolving designs.

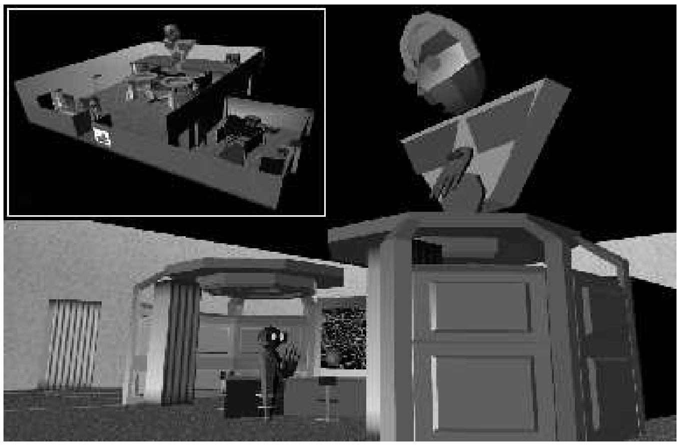

Figure 1: A deity and a mortal share a common design space- the Integration Testbed Prototype.

2.1 Hardware and Software

CALVIN was designed to run in the CAVE^{tm} virtual environment- a 10 foot cube constructed of three translucent walls. Its room-sized structure allows for multiple users to move around freely, both physically and virtually surrounded by vivid displays. Interaction is performed mainly with the use of a wand. Since the CAVE library can support heterogeneous VR display devices (ImmersaDesk^{tm}, BOOM, fish-tank VR systems) it allows collaborators to use their preferred VR hardware to work within CALVIN.

CALVIN is written in C++ using OpenInventor^{tm} as the underlying graphics library. This provides a number of benefits: 1. OpenInventor provides a convenient means to organize three-dimensional scene hierarchies while at the same time offering scene culling and level-of-detail management. 2. The entire scene hierarchy can be saved at any time and trivially converted to VRML for general distribution over the World Wide Web. 3. OpenInventor and the CAVE library are portable to other platforms. 4. There is a large collection of VRML architectural models appearing on the Web which can be used as models to be explored by CALVIN.

2.2 Co-presence

Physical co-presence between distantly located collaborators is one of the main characteristics that is difficult to simulate with traditional computer media. A number of approaches in Computer Supported Cooperative Work (CSCW) that use a combination of video cameras and microphones have been developed to attempt to bridge the geographic distance between remote collaborators. Often these techniques fall short because they require the user to switch between multiple stationary cameras in order to acquire a suitable view of the environment[5, 6]- and even then the view is sub-optimal and the interface is unnatural.

Through VR we are able to address this problem by providing co-presence that uses avatars and natural navigation in the environment. An avatar is a persona that each participant dons to establish a representation in the virtual environment. Using significantly contrasting avatars, participants are able to see and identify each other in the space and hence communicate notions of relative position and direction that are important when they gather to collaborate.

In CALVIN, avatars are composed of a separate head, body and hand. This is motivated by the fact that we currently have two trackers attached to a participant: one for the head and one for the hand, allowing the environment to transmit gestures (e.g. nodding or waving.) As mortals and deities typically view the world from two very different perspectives, the deities will naturally appear like giants to the mortals, whereas mortals will appear like dwarves to the deities (figure 1.)

Another feature, important in collaborative work, is support for verbal communication, and to a lesser degree, video communication. In CALVIN, constant real-time audio and video (NTSC resolution at 30 frames per second) is streamed through a high-speed ATM connection (figure 2) between collaborators. While verbal communication has been found to be a crucial part of collaborative task completion, video has been found to be more important for mediating negotiations[5].

2.3 Persistence and Network Support

To sustain the level of persistence that will allow collaborators to work both synchronously and asynchronously, CALVIN manages a central repository of information that maintains the states of the various on-going design projects. This repository contains a collection of objects to be placed in the scenes, a collection of avatars, and a collection of scene description files

Figure 2: Participants communicate with audio and video over an ATM network using the ImmersaDesk

which store information on the state of every object within the scene. Each user maintains a user definition file for each scene in which they are involved, containing information on which avatar is portraying the user, where the user is located in the scene, and whether the user is a mortal or a deity. This information is loaded into CALVIN's central database server prior to a design session.

CALVIN's network component allows multiple networked participants to work in the same virtual space. Multiple distributed CALVINs running on separate VR systems are connected via ATM to the centralized database that guarantees consistency across the various environments. The communications library supporting CALVIN's networking is built on a client-server model with the number of remote clients limited only by bandwidth and latency. Several similar approaches have been implemented[7]. Although a centralized database approach can be a bottleneck, CALVIN uses this simple approach because it allows us to build distributed virtual environments to concentrate on the human-factors issues of collaborative interaction in persistent virtual environments. We believe that network bandwidth and latency will always pose a bottleneck in any complex distributed environment that involves a large number of participants; hence it is important to devise interaction techniques to minimize this problem at the perceptual level.

2.4 Interfaces

CALVIN mainly uses the CAVE's wand for object manipulation. For example, to move an object, users may press a wand button and then move their arm in the direction they wish the object to go. The user can walk within the confines of the CAVE or "drive" the entire CAVE through the virtual space. In both cases, downwardly directed collision detection is applied so that users may climb objects (e.g. stairs) by physically walking on them. CALVIN then automatically lowers the entire landscape to simulate the effect of raising the user. Using this scheme in a collaborative environment allows one participant to stand on an object while another participant lifts them up by lifting the object up. A more direct means of lifting another participant up is if one of the participants is a deity. As a deity one may place one's hand on the ground and ask a mortal to step on the hand. When the deity then lifts the hand up, the mortal will also be lifted.

CALVIN uses two complementary interfaces for menu selection: the Virtual Visor and speech recognition. The Virtual Visor simulates a head-up display (HUD) [8] in the virtual environment. The visor can also be used as a gaze-directed input (GDI) device as shown in figure 3. To make a selection, the user gazes at the appropriate menu option and presses a button on the wand.

Virtual menuing interfaces based on their 2D desktop counterparts require menu-item search, acquisition, and then selection. GDI combines menu-item search with menu-item acquisition, thereby reducing the number of steps required to make a menu selection. GDI may provide an additional benefit when it is used with the Head-Mounted Display, Fish-Tank VR systems, and the ImmersaDesk. In such systems there is a common problem of fatigue that occurs due to prolonged raising of the user's arm to make menu selections. GDI can be used to greatly reduce this problem by reducing the number of operations requiring arm movements.

Thus far, the Virtual Visor has been applied to three additional visualization environments (a simulation of contaminant evolution through porous media- Louis Rossi, Northwestern University; a four dimensional navigator for astrophysical turbulence data- Andrea Malagoli, University of Chicago; and a network performance monitor- Kyoung Park, University of Illinois at Chicago.[1]) Preliminary feedback from users seems to empirically confirm the effectiveness of the visor.

In addition to the visor, a speaker-independent speech recognition system is provided to supplement visor menu selection.

Figure 3: The Virtual Visor being used as a gaze-directed input device in the CAVE.

Audio feedback is used to confirm visor menu selections and speech commands. Speech recognition was found to be, in general, quicker and easier to use than the visor or previous experimental menu systems. However, the visor is still necessary for situations where excessive noise levels render speech recognition unreliable.

2.5 Application of CALVIN to Design

To date, CALVIN has been used to collaboratively design a number of environments. These include the layout of:

1. a computer room at NCSA to house the Infinity Wall (a four-screen active stereoscopic projection environment and a number of workstations)[9].
2. the GII testbed room at Supercomputing '95, consisting of a CAVE, two ImmersaDesks, seating for 200 people and two Infinity Walls (one active stereo and one pasive stereo)[10].
3. octagonal command and control centers (called Octmods) in a prototype of the Integration Testbed room at ARPA headquarters (formerly known as the ARPA Enterprise.)

In the interest of brevity, we will only describe the Octmod space as an example of the type of interaction that occurred between the participants (figure 1.)

Octmods are octagonally shaped "pods" or "gazebos" that can be quickly assembled and attached to one another by two people using a screwdriver. Octmods are designed to automatically provide connectivity to control lights and networking between modules. Computers can then be easily placed in the Octmods and configured to attach to the LAN.

The virtual Integration Testbed room consisted of a number of Octmods, a CAVE, and an ImmersaDesk. During Supercomputing '95 participants at the San Diego convention center and at ARPA were able to collaborate in manipulating the components of the room to redesign its layout. Both participants were in their respective CAVEs at the two remote locations. One participant, assuming the role of a mortal, was able to configure the panels that form an Octmod while the other participant, assuming the role of a deity, watched overhead and moved entire Octmod, CAVE and ImmersaDesk structures around the space. These activities were coordinated through the real-time ATM audio/video link.

3 Discussion and Future Work

CALVIN has provided us with valuable insight into the issues involved in supporting collaboration in VR. CALVIN is currently being extended in an ongoing collaboration between EVL and industry to provide techiques for manufacturing design and rapid prototyping in VR. The CALVIN testbed will be used to further study issues in collaborative VR. These include the evaluation of the effectiveness of using heterogeneous perspectives in collaborative work; the development of database technology to support persistent virtual worlds; and the improvement of technology to support avatar, audio, and video mediated co-presence.

In addition, the techniques derived from CALVIN will be applied to two new problem domains: a narrative-based educational environment called NICE (Narrative Immersive Constructionist/Collaborative Environments), and a collaborative scientific visualization environment.

- NICE[11] is an approach for applying VR to the creation of a family of educational environments for young users. It is based on constructionism, where real and synthetic users, motivated by an underlying narrative, build persisting virtual worlds through collaboration. GULLIVR (Graphical User Learning Landscapes In VR) is a prototype system based on CALVIN[9], and the Graphical StoryWriter[12] that allows us to explore these ideas (figure 4.)

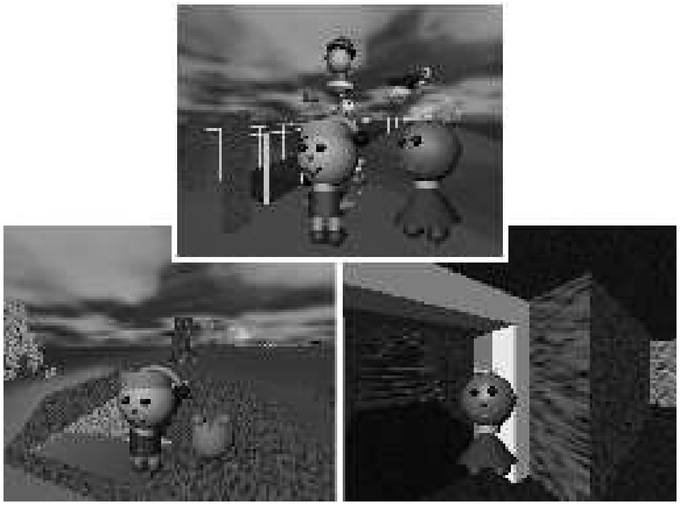

Figure 4: Three scenes from a GULLIVR story set.

- We believe the approach of applying multiple perspectives design can be generalized to other disciplines such as collaborative engineering and scientific visualization. In scientific visualization, the data being viewed is typically multi-dimensional. By applying the notion of heterogeneous perspectives we may be able to partition the number of dimensions into smaller, more manageable pieces that multiple participants can then explore simultaneously. We intend to explore this technique further in our development of a collaborative virtual laboratory for scientific visualization.

Acknowledgements

We would like to thank Michael Kelley, Bruce Gibeson, and Steve Grinavic at University of Southern California. We would especially like to thank the other members of the NICE team: Maria Roussos, Christina Vasilakis, and Thomas Moher.

References

[1] H. Korab and M. D. Brown, eds., *Virtual Environments and Distributed Computing at SC '95: GII Testbed and HPC Challenge Applications on the I-WAY*. ACM/IEEE SUPERCOMPUTING '95, 1995.

[2] F. P. Brooks, Jr., "Walkthrough — A dynamic graphics system for simulating virtual buildings," in *Proceedings of 1986 Workshop on Interactive 3D Graphics* (F. Crow and S. M. Pizer, eds.), pp. 9–21, Oct. 1986.

[3] R. Stoakley, M. J. Conway, and R. Pausch, "Virtual reality on a WIM: Interactive worlds in miniature," in *SIGCHI '95 Proceedings*, pp. 265–272, May 1995.

[4] J. S. Olson, L. A. Mack, and P. Wellner, "Concurrent editing: The groups interface," in *Human-Computer Interaction - INTERACT'90* (D. Diaper, ed.), pp. 835–840, 1990.

[5] W. Gaver, A. Sellen, C. Heath, and P. Luff, "One is not enough: Multiple views in a media space" in *Proceedings of INTERCHI'93*, pp. 335–341, 1993.

[6] R. S. Fish, R. E. Kraut, and B. L. Chalfonte, "The videowindowsystem in informal communication," in *Proceedings of Computer Supported Cooperative Work*, pp. 1–11, 1990.

[7] M. J. Zyda, D. R. Pratt, J. G. Monahan, and K. P. Wilson, "NPSNET: Constructing a 3D virtual world," in *Proceedings of the 1992 Symposium on Interactive 3D Graphics*, pp. 147–156, 1992.

[8] C. D. Wickens and L. J. Long, "Conformal symbology, attention shifts, and the head-up display," in *Proceedings of the Human Factors and Ergonomics Society*, pp. 6–10, Oct. 1994.

[9] J. Leigh, A. E. Johnson, C. A. Vasilakis, and T. A. DeFanti, "Multi-perspective collaborative design in persistent networked virtual environments," in *Proceedings of IEEE Virtual Reality Annual International Symposium '96*, Apr. 1996.

[10] J. Leigh, A. Johnson, and T. A. DeFanti, "CALVIN: an immersimedia design environment utilizing heterogeneous perspectives," in *Proceedings of IEEE International Conference on Multimedia Computing and Systems '96*, June 1996.

[11] M. Roussos, A. E. Johnson, J. Leigh, C. A. Vasilakis, and T. G. Moher, "NICE: Narrative immersive constructionist environments (http://www.ice.eecs.uic.edu/~nice)," 1996.

[12] K. E. Steiner and T. G. Moher, "Graphic storywriter: an interactive environment for emergent story construction," in *Proceedings of CHI '92*, pp. 357–363, 1992.

Biography

Andrew E. Johnson is a post-doctoral researcher at the Electronic Visualization Laboratory at the University of Illinois at Chicago. His research interests include collaborative virtual environments and using virtual reality as a high level interface for visualizing large scientific datastores.

Jason Leigh is a PhD candidate at the Electronic Visualization Laboratory at the University of Illinois at Chicago. His PhD thesis focuses on developing techniques for enhancing collaborative work in virtual reality.

Readers may contact Leigh at the Electronic Visualization Laboratory, University of Illinois at Chicago, M/C 154, 851 S. Morgan St., Chicago, IL 60607-7053, e-mail jleigh@eecs.uic.edu