Alive on the Grid

Dave PAPE
Buffalo, NY 14222, US
dave.pape@acm.org

and

Dan SANDIN
EVL, University of Illinois at Chicago
Chicago, IL 60607, US

ABSTRACT

"Alive on the Grid" is a collection of virtual art worlds where local and distant participants alike can interact in shared virtual spaces. Enabled by the Grid – collections of networks, computers and virtual reality displays that span the globe – users interact with one another and the models contained within each piece. This paper discusses both the artistic and technical aspects of Alive on the Grid.

Keywords: Virtual reality, shared environments, computer art

1. INTRODUCTION

Alive on the Grid is a networked virtual reality art show that was first demonstrated as part of the 2001 Ars Electronica Festival[1]. It is effectively a group show - within the overall virtual world were ten distinct pieces by contributors from several institutions. It is also a shared world, in that participants at seven international sites visited the common environment together; visitors from the different sites saw, manipulated, and interacted with the same virtual objects, and were visible to each other as 3D avatars. The environment was created to be shown in CAVEs [2], ImmersaDesks [3], and similar projection-based virtual reality displays.

There have been several virtual reality artworks created in the past as shared, multi-user worlds; however, due to the many difficulties involved, very few have attempted to function as widely distributed environments over the Internet. The typical form of a shared VR art environment has been a pair of users at a single site together; for example, Videoplace [4] can take the images of two people and put them to-

gether to interact. Placeholder [5] similarly allowed a pair of participants to enter the space at one time; however, by letting them leave marks and recordings in the virtual space, a sense of many more visitors being there over time was created. The HiPArt project's Artworld [6], a collaborative space constructed by a large group of New England artists, can operate over the Internet; the dynamics and interaction within the environment are kept relatively simple, such as keyframed animations activated by users' presence.

2. CONTENT

The space of Alive on the Grid consists of a central connecting area (called the Confluxus – figure 1) and ten worlds that participants can reach from there. The central Confluxus was valuable because of the 'show' nature of Alive on the Grid — when new groups of visitors arrived at the Festival site, they and the remote participants could assemble at the Confluxus for introductions and orientation, before heading off into the actual pieces. The individual pieces are "Home", "Looking for Water", "Super Spectacular", "PAAPAB", "Infinite Studio", "Through the Looking Glass", "syn.aesthetic", "Excavation", "Beatbox", and "Incarnation of a Divine Being". These pieces explored a wide range of possible approaches to VR art. The environments varied from complex collections of pre-built models and sounds to highly dynamic particle effects to spaces where visitors created objects or music themselves. The interactions varied from activating simple triggers to motion-captured dancing to a collaborative, improvised drama.

Many of the artists involved in Alive on the Grid have been involved in creating VR works

for several years. However, almost none had created networked pieces before. The question of how to involve multiple people at widely separate locations in a single virtual art piece is still a difficult one. In many cases, the participants were simply in the same space and experiencing it together, similar to a group visiting a traditional museum exhibit together. In some cases, though, interaction between the people was an important part of the piece. In the following sections we describe a selection of the environments and highlight some of the interesting features of each one.

Home

Home (figure 2) is an environment that explores ideas about a dwelling and its relation to the psyche. Visitors enter a house that has been abandoned by its owners; while inside this space, the visitors (that is, their avatars) become ghosts. The rooms and hallways include a soundscape composed of spoken narrative fragments and a musical ambient score.

Home began as an on-line VRML environment, and was remade for the CAVE environment. It contains links to the work of numerous artists who reflect on the meaning of domestic space and identity, including a screenwriter, photographer, set designer, film and video makers, and computer artists.

Looking for Water

Looking for Water (figure 3) consists of a series of sub-environments based around the theme of water. Participants begin in outer space, surrounded by sun and Earth images which are animated playbacks of real-time satellite weather data. The images are refreshed daily to form a long-term visual history. They then fall to Earth and land on a northern Lake Michigan archipelago of islands. The 3D models of the archipelago of islands are based on panoramic video images taken while kayaking. The 3D information has been extracted from the moving video images to allow participants to move about and explore these worlds. In the end, the world dissolves into moving water; the visitors enter an area of dynamic, particle-system based fountains and hoses that they can manipulate — they can build water sculptures, or simply spray each other.

Incarnation of a Divine Being

Incarnation of a Divine Being (figure 4) is an interactive, improvised play. In it, one person (a professional actor) leads the others in creating a story, based on ancient Greek drama. In this case, the entirety of the experience is focused on the distributed participants talking and interacting with each other.

Infinite Studio

Infinite Studio (figure 5) allows the users to create, in real-time, their own artwork from inside the virtual environment. Using a color palette and several drawing and modeling tools, users can create and modify virtual objects for any effect desired. Every line drawn has a lifespan causing the entire art piece to constantly evolve. Individuals can also collaborate to create group constructions. In particular, users would often comment on each other's creations and construct new ones in response, filling the space as a group.

PAAPAB

PAAPAB (Pick an Avatar, Pick a Beat) (figure 6) is a dance floor with a steady dance beat. The user joins the dance by selecting one or more life-sized puppets whose movements mirror that of the user's body through a simple motion capture system. As more users join in, the floor gets more crowded. Each puppet shape varies, and the trackers are mapped to their body parts in different ways. Hence, the user may move her arms while the puppet moves its legs, tail or wings. Over time, a diverse group of animated puppets inhabit the virtual dance floor. At each networked site, a local heartbeat keeps the dancers and the music in sync with one another.

3. IMPLEMENTATION

Alive on the Grid is based on the Ygdrasil authoring system. Ygdrasil is a framework for creating networked virtual environments, particularly by artists, educators, and other VR users who are not necessarily professional programmers. It is focused on building the behaviors of virtual objects from re-usable components, and on sharing the state of an environment through a distributed scene graph mechanism.

Ygdrasil is built in C++, around SGI's OpenGL Performer visual simulation toolkit [7] and the CAVERNsoft G2 networking library [8]. Performer provides a hierarchical scene graph rep-

Figure 1: Participants gathered in the Confluxus

Figure 2: Home - visitors become ghosts inside an abandoned house

Figure 3: Looking for Water - two local people and two remote participants view the sun

Figure 4: Incarnation of a Divine Being - users take part in an improvised drama

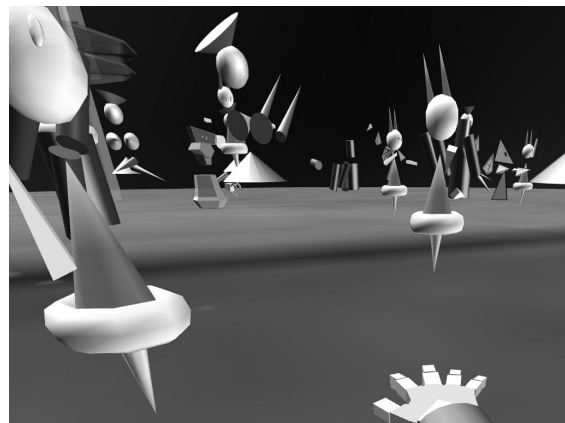Figure 5: Infinite Studio - users create models in a shared workspace

Figure 6: PAAPAB - puppets recorded by previous visitors fill the disco

resentation of the virtual world database; our framework extends this to be a shared scene graph. Necessary data, such as lists of nodes' children, transformation matrices, and model information, are automatically distributed among participants in the application via CAVERNsoft. CAVERNsoft is a networking toolkit for VR that emphasizes integrating VR with high-performance and data-intensive computing over high-speed networks. It provides a distributed database mechanism, which we use to share the scene graph data.

In Ygdrasil, in addition to the basic graphical data used in Performer, any scene graph node can have behaviors added to it. Each particular node is considered to be owned by the host that creates it. This host executes any behavior associated with the node. All other hosts will create proxy versions of the node, and only receive data for it; they do not directly modify the node, except by sending messages to the master copy to request changes. Because the basic scene graph data – that which is sufficient to render the scene – is shared automatically, new behavioral components generally do not have to include any networking themselves.

The data that are shared for any node in the scene graph are stored in the CAVERNsoft database keyed by the node name and the data members' names (see figure 7). When a client is first informed about a new node (by the node's name), it looks up the node and its type in the database. It can then retrieve all the other data as needed.

Most behaviors in Ygdrasil applications are built as simple components in C++. They are new node classes that extend other, existing classes. The individual node classes are compiled into dynamically loaded objects (DSOs), so that they can be rapidly added to a world or modified. The system also includes a number of premade classes (also DSOs) that implement common virtual world interactions; these include such things as users' avatars, navigation controls, and triggers that detect when a user enters an area. These built-in tools simplify the quick construction of many basic applications.

The actual composition of a virtual world in Ygdrasil is done using a higher level, scripting-like layer. Other toolkits have used traditional procedural or object-oriented scripting languages, such as Scheme in Avango[9] or VisualBasic in WorldUp[10]. The scripting layer in Ygdrasil is a simple textual representation of the scene graph layout (or a fragment of a scene graph), similar to an OpenInventor object file. It tells the system what kinds of nodes to create, and
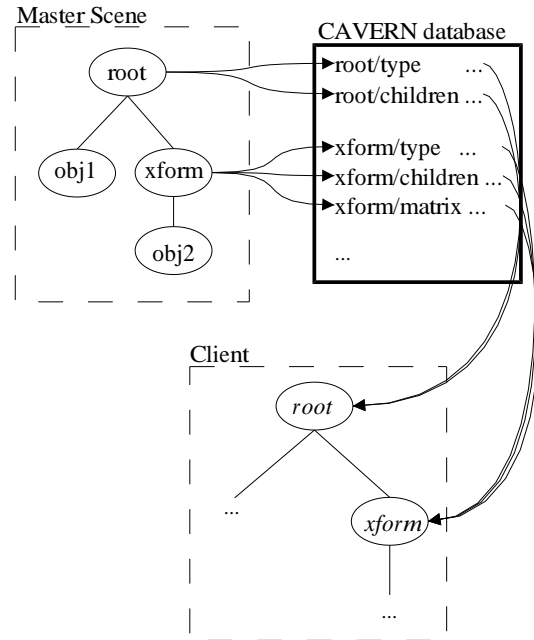

Figure 7: Ygdrasil scene graph database

includes commands with each node to control its behavior. An important aspect of the scripting system is event-driven communications — nodes such as triggers and timers generate events, and in the scene graph script these events are used to send messages to other nodes, activating or modifying their behaviors. This scripting layer makes it possible for experienced programmers and non-programming designers to work together in creating a world — experienced programmers create new behavior components when necessary, while others can create a world by simply plugging together the components.

4. NETWORKING FOR ARS ELECTRONICA

The shared world of Alive on the Grid consists of the master version of the world database being maintained on one computer, with clients in CAVEs and other VR systems connecting to it and receiving a stream of continuous updates for their local copies of the database. The clients also add data for their users' avatars to the shared database, and use a streaming audio conferencing tool to allow the remote users to talk to one another.

The final, complete world for Alive on the Grid contained 1188 modeled objects, 363 recorded sounds, and many additional textures, captured motions, and other files, for a total of 2.5 giga-

bytes of data. The scene graph included roughly 1200 dynamic nodes, and created 14000 entries in the shared CAVERNsoft database. Ideally, we would like new clients to be able to join the running world without any preparation, and simply download all the necessary data on the fly. But, given the size of the environment, this was not practical, and instead full copies of the models, sounds, etc. were sent to participating sites in advance.

For the Ars Electronica show, the master program ran on the CAVE Onyx at the Ars Electronica Center in Linz, Austria; participants in Chicago, Buffalo, Bloomington, Umeà (Sweden), Budapest, and Amsterdam, as well as an ImmersaDesk at the Center, connected at various times during the week. The Center obtained a special, added 2 megabit per second Internet connection for use during the show. However, the actual measured bandwidth from the Center to the United States was often much less than this, sometimes dropping as low as 200 kilobits/sec. Even at its maximum, this bandwidth was not sufficient to support several simultaneous clients, given the amount of dynamic data in the environment and the streaming audio communication. To get around this limitation, we ran a secondary copy of the shared database on a host at the Electronic Visualization Laboratory in Chicago with much better network connectivity. The master program in Linz linked its database to to this duplicate, and all the outside clients connected to Chicago; those in the U.S. reached it over Internet2, while the European participants came in through STAR TAP, an international networking transit point in Chicago. During the course of the show, the bandwidth used at the Chicago gateway reached roughly 5 megabits/second. Figure 8 outlines the connections that were made.

5. CONCLUSION

Alive on the Grid demonstrated some of the possibilities for the use of high-speed networks and immersive displays in interactive art. The scale of the environment currently limits it to high-end computers and research networks that are rarely available to the public, but with time similar systems will become mainstream.

The pieces that make up Alive on the Grid cover a wide range of visual styles and styles of interaction. Presenting these pieces in a public show was vital to learning how well these approaches worked, by how other people reacted to them, and to hammering out the technical difficulties.
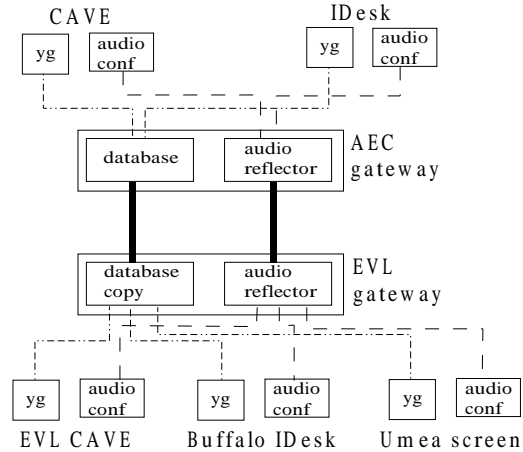

Figure 8: Network connections for Ars Electronica show

Perhaps the most important outstanding issue faced by Alive on the Grid, and its developers, is the need for the networking middleware to be more robust. We need to make it easier for users to run these sorts of applications in a distributed mode, especially when network limitations require a less-than-simple configuration, such as was used at the Ars Electronica Festival. Making it so that sharing an immersive virtual environment is as easy as sharing a web page will lead to an explosion in the creation of these environments and exploration of new concepts for interaction.

6. ACKNOWLEDGMENTS

7. REFERENCES

[1] Gerfried Stocker, Christine Schopf, *Ars Electronica 2001: Takeover - who's doing the art of tomorrow*, Vienna: Springer Verlag, 2001.

[2] Carolina Cruz-Neira, Daniel J. Sandin, Thomas A. DeFanti, "Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE," *Proceedings of SIGGRAPH 93 Computer Graphics Conference*, pp. 135-142, 1993.

[3] Marek Czernuszenko, Dave Pape, Daniel J. Sandin, Thomas A. DeFanti, Gregory L. Dawe, Maxine Brown, "The ImmersaDesk and Infinity Wall Projection-Based Virtual Reality Displays," *Computer Graphics*, Vol. 31 No. 2, pp. 46-49, May 1997.

[4] Myron W. Krueger, *Artificial Reality II*, Reading: Addison-Wesley Publishing Company, 1991.

[5] Brenda Laurel, Rachel Strickland, Rob Tow, "Placeholder: Landscape and Narrative In Virtual Environments," *Computer Graphics*, Vol. 28 No. 2, May 1994.

[6] Boston University HiPArt program, "Artworld," http://scv.bu.edu/HiPArt/ArtWorld/ artworld_home.html.

[7] John Rohlf, Jim Helman, "IRIS Performer: A High Performance Multiprocessing Toolkit for Real-Time 3D Graphics," *Proceedings of SIGGRAPH 94 Computer Graphics Conference*, pp. 381-395, 1994.

[8] Kyoung S. Park, Yong J. Cho, Naveen K. Krishnaprasad, Chris Scharver, Michael J. Lewis, Jason Leigh, Andrew E. Johnson, "CAVERNsoft G2: A Toolkit for High Performance Tele-Immersive Collaboration," *Proceedings of the ACM Symposium on Virtual Reality Software and Technology 2000,* pp. 8-15, 2000.

[9] Henrik Tramberend, "Avocado: A Distributed Virtual Reality Framework," *Proceedings of IEEE Virtual Reality '99*, 1999.

[10] Sense8 Corporation, *WorldUp Users Guide, Release 4*, 1998.