

Shared Miletus: Towards a Networked Virtual History Museum

Dave Pape, Josephine Anstey†, Sarita D'Souza, Tom DeFanti

Electronic Visualization Laboratory
University of Illinois at Chicago
Chicago, IL USA
E-mail: dave.pape@acm.org

Maria Roussou, Athanasios Gaitatzes

Foundation of the Hellenic World
Athens, Greece

ABSTRACT

Shared Miletus is a networked Virtual Cultural Heritage application first demonstrated at the INET 2000 conference. In this application we sought to explore new tools and techniques, ones that could be useful to create a networked virtual environment which takes the place of a traditional museum. This paper describes the software used to build the environment, and the virtual tools created specifically to support remote, international visitors to the exhibition.

1. INTRODUCTION

Cultural heritage is becoming an important application for virtual reality technology. A recent EC/NSF Advanced Research Workshop identified it as one of the key application domains for driving the development of new human-computer interfaces such as VR [1]. Computer simulations can offer public access to reconstructions of historical sites or artifacts that would normally be inaccessible, due to location or the fragile condition of the artifacts. They also provide the possibility of visiting places that no longer exist at all, or of viewing the how the sites would have appeared at different times in history.

A number of research projects have delved into the problems of building accurate digital reconstructions of ancient artifacts. For example, [2] describes the work done in creating models of the tombs of Nefertari and Tutankhamun, and presenting a VR tour of the models in a public museum. Our interest extends to what to do with the basic, computerized content, once it has been created. We see the combination of virtual reality and high speed network technologies as making possible a networked virtual museum, where people can step directly into an ancient city across the Internet.

†Current affiliation: University of Buffalo



Figure 1: Miletus in the ReaCTor

2. BACKGROUND

The Foundation of the Hellenic World (FHW) is a non-profit, privately funded museum and cultural research institution in Athens, Greece. Its mission is to preserve and present Hellenic history and culture; it seeks to use state-of-the-art technology to accomplish these goals. The FHW owns two virtual reality systems, an ImmersaDesk and a ReaCTor, that are used to present a variety of content created by Foundation staff [3]. Exhibits using the systems have included 3D reconstructions of ancient cities and buildings, as well as educational, interactive environments such as the history of Hellenic costume. Experienced museum guides lead visitors through the exhibits; the guides must have both technical skills to operate the VR displays and museum education skills to explain the history of the city. The guides are an important part of the exhibits, and any Internet-based version of these exhibits must in some way take into account their role in educating visitors.

One of the first applications shown in the

VR systems at FHW was a reconstruction of Miletus, an ancient city on the coast of Asia Minor (see figure 1). Detailed models of some of the buildings of Miletus were created, and museum visitors can explore the city as it was in antiquity.

2.1. iGrid 2000

The International Grid (iGrid) is a series of research demonstrations highlighting the value of international high-speed computer networks in science, media communications, and education [4]. iGrid 2000 took place at the INET 2000 conference in Yokohama, Japan. It provided a 100 Mbps connection from the conference site to STAR TAP in Chicago; STAR TAP serves as an international connection point for several research networks in America, Europe, and Asia. For this demonstration, we created an experimental, shared version of the Miletus environment. The objective of the application was to take the content that would normally be shown in the controlled environment of FHW's museum, and let remote, networked people visit it. In particular, we did not want to simply make it something like a VRML model that visitors would download and then play with on their own; instead, it was to be a dynamically shared world, "hosted" by the conference demonstrators (or, in the future, the museum).

3. YGDRASIL

Shared Miletus is based on the Ygdrasil authoring system. Ygdrasil is a framework that we are developing as a tool for creating networked virtual environments. 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. It is presently being used in the construction of several artistic and educational applications.

Ygdrasil is built in C++, around SGI's IRIS Performer visual simulation toolkit [5] and the CAVERNsoft G2 networking library [6]. Performer provides a hierarchical scene graph representation 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 as used in Performer, any scene graph node can have behaviors built in to it, by subclassing one of the core node classes. Each particular node is considered to be owned by the host that creates it. The owning host maintains the master version of the node, and executes any behavior associated with it. All other hosts will create proxy versions of the node, and only receive data for it through CAVERNsoft; they do not directly modify the node (they can however send messages to the master copy to request changes). The proxy version of a node is typically of the simpler parent class type, without the added behavior code. Thus, if all of the main behaviors for a virtual world are executed by a single master version of the scene, remote sites can join in this world without needing anything beyond the standard, core node types.

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 2). 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 pre-made 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. Rather than a traditional procedural or object-oriented type of language (e.g. the use of Scheme in Avango [7]), this scripting layer 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 includes commands with each node to define or modify its behavior. This layer makes it possible for non-programming designers to create a world by simply plugging together the components, working with experienced programmers

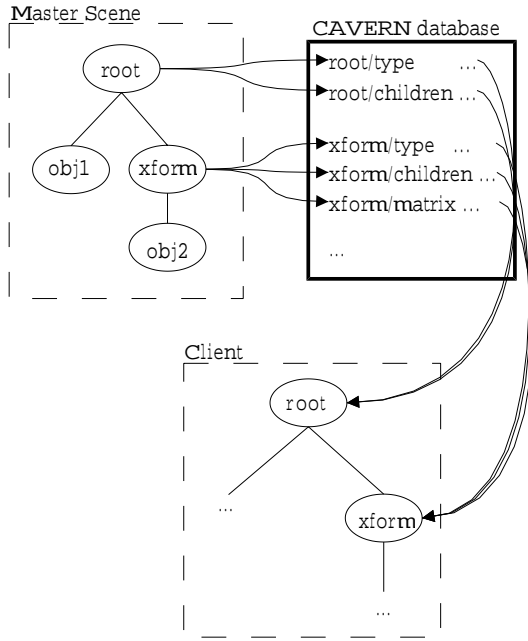


Figure 2: Ygdrasil scene graph database

when necessary to create new behavior nodes.

4. SHARED MILETUS

In creating the demonstration of the Shared Miletus environment, we focused on two issues – guiding visitors through the city, and providing them with information about what they were seeing. These features needed to work in an internationally distributed environment, where users could come and go from the space at will.

Many museum-based VR exhibits will lead visitors through the virtual world on a pre-selected path, so that users do not have to learn any special controls or know where they should be going. The river metaphor, described in [8], extends this model by allowing the users to stray somewhat from the fixed path, but always guaranteeing that they continue to progress in the right direction. In our case, we wanted to give the visitors freedom to explore Miletus at their own pace. They were given a 3D wand used for simple joystick-driven navigation; a recorded introduction when they entered the space explained how to use the wand. To make it easier to get to places of interest, we also gave them a dynamic, virtual map. This map showed the layout of the city, the user’s position in it, and also the positions of any other participants in the shared world. This helped them to drive to particular buildings, or to meet up with other

visitors or guides from the museum. In addition, the map could be used as a navigation shortcut – clicking on a particular building would summon a magic carpet that then automatically brought the user to that building’s entrance. If a visitor were to get completely lost, a special reset button would start him back at the entrance to Miletus, and repla the instructions on how to use the wand and map.

The first stage of providing visitors with information about Miletus was to include expert human guides. Guides from the actual, “official” museum could enter the shared world, just like an ordinary visitor. Through their avatars, and streaming network audio connections, the guides could then interact with the visitors, pointing out special details and answering questions.

Given the international scope of the shared space, human guides alone are unlikely to be enough – there could be large numbers of visitors, and they could be exploring the space at any time of day. So, we placed automated information kiosks within the various buildings of Miletus. These kiosks contained pre-recorded audio commentary describing each building and its history. In order to support an international audience, this audio was available in multiple languages; for the iGrid demo we provided English and Japanese commentaries, but given enough time and translation personnel, any number of languages could be supported. The multi-lingual capability was implemented by having each visitor carry their own virtual audio tool. The tool was effectively a part of the user’s avatar, and kept track of his preferred language. When the user approached a kiosk, a trigger detected the presence of an audio tool and sent the tool messages informing it of what recordings the kiosk could provide. If the user chose to listen to one of them, the tool would send a request back to the kiosk, asking for the appropriate sound file for the desired language. Other tools at the entrance to the world could be used to switch languages – clicking on a Japanese flag icon would send a message to the user’s audio tool to use Japanese, for example. The audio tool also provided the introduction and navigation instructions in the appropriate language.

5. CONCLUSION

The iGrid 2000 demo of Shared Miletus was successful; however, the particular high-end research nature of the venue also points out the current limitations of the application. The general public does not have easy access to high speed international networks and large-scale VR

devices. We do not feel these problems are insurmountable, though – general broadband network access is continually improving, bring greater bandwidth to private citizens, and home computer capabilities are increasing thanks to the demands of games. We are now able to run a reduced version of Miletus on an ordinary Linux PC with a commercial graphics card.

Shared Miletus has demonstrated the possible future use of networking and virtual reality to provide a tele-immersive museum. This approach can increase the accessibility of a museum to world-wide audiences. It can also allow much larger and more in-depth exhibits than in a physical museum, and offers a chance for greater direct interaction by visitors with an exhibit. We have only just scratched the surface of what is possible in such a virtual museum; much work still remains in both the creation of tools and the construction of new experiences for the public.

6. ACKNOWLEDGMENTS

We would like to thank the entire iGrid 2000 staff, especially Hiroshi Esaki and Goro Kunito of the University of Tokyo and Akihiro Tsutsui of NTT, for their help and extreme dedication to making the show a success.

The virtual reality research, collaborations, and outreach programs at the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago are made possible by major funding from the National Science Foundation (NSF), awards EIA-9802090, EIA-9871058, ANI-9980480, and ANI-9730202, as well as the NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement ACI-9619019 to the National Computational Science Alliance. EVL also receives major funding from the US Department of Energy (DOE) Science Grid program, awards 99ER25388 and 99ER25405, and the DOE ASCI VIEWS program, award B347714. In addition, EVL receives funding from Pacific Interface on behalf of NTT Optical Network Systems Laboratory in Japan.

ImmorsaDesk is a registered trademark, and STAR TAP is a service mark, of the Board of Trustees of the University of Illinois. ReaCTor is a trademark of Trimension Systems.

7. REFERENCES

- [1] Judith R. Brown, Andy van Dam, Rae Earnshaw, José Encarnação, Richard Guedj, Jennifer Preece, Ben Shneiderman, John Vince, “Human-Centered Computing, Online Communities, and Virtual Environments,” *IEEE Computer Graphics & Applications*, 19(6), pp. 70-74, 1999.
- [2] Pierre Boulanger, John Taylor, Sabry F. El-Hakim, Marc Rioux, “How to Virtualize Reality: An Application to the Recreation of World Heritage Sites,” *Proceedings of Fourth International Conference on Virtual Systems and Multimedia*, pp. 39-45, 1998.
- [3] Maria Roussou, “Incorporating Immersive Projection-based Virtual Reality in Public Spaces,” *Proceedings of the 3rd International Immersive Projection Technology Workshop*, pp. 33-39, 1999.
- [4] Maxine D. Brown, Thomas A. DeFanti, Michael A. McRobbie, Alan Verlo, Dana Plepys, Donald F. McMullen, Karen Adams, Jason Leigh, Andrew E. Johnson, Ian Foster, Carl Kesselman, Andrew Schmidt, Steven N. Goldstein, “The International Grid (iGrid): Empowering Global Research Community Networking Using High Performance International Internet Services,” *INET '99*, 1999.
- [5] John Rohlf, Jim Helman, “IRIS Performer: A High Performance Multi-processing Toolkit for Real-Time 3D Graphics,” *Proceedings of SIGGRAPH 94 Computer Graphics Conference*, pp. 381-395, 1994.
- [6] 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.
- [7] Henrik Tramberend, “Avocado: A Distributed Virtual Reality Framework,” *Proceedings of IEEE Virtual Reality '99*, 1999.
- [8] Tinsley Galyean, “Guided Navigation of Virtual Environments,” *Proceedings of 1995 Symposium on Interactive 3D Graphics*, pp. 85-92, 1995.

- [1] Judith R. Brown, Andy van Dam, Rae Earnshaw, José Encarnação, Richard Guedj, Jennifer Preece, Ben Shneiderman, John Vince, “Human-Centered