Network Requirements for Large Distributed Virtual Environments

R. B. Araujo[1] *and Dr. C. Kirner*[2]
Computer Science Department
Federal University of São Carlos
Via Washington Luiz, Km.235
13565-905 São Carlos - SP BRAZIL
Tel.: +55-162-74823 Fax: +55-162-712081
drea/dcki@power.ufscar.br

Abstract

This paper presents the Network requirements of a special class of telepresence applications: Large Distributed Virtual Environments. The characteristics of large virtual environments are introduced and their network requirements analysed from the network bandwidth and latency point of view. When multimedia information is integrated to these systems, demands over the network are even more stringent.

Keywords

Large distributed virtual environments, multimedia information, network latency, network bandwidth, multicasting.

1 INTRODUCTION

Virtual Environment - VE systems have emerged as a breakthrough in human-computer interaction, promising a revolution in areas ranging from entertainment to medicine. A VE is defined here as a system which simulates an environment so that a user can feel immersed in

---

[1] PhD Student in the Integrated Systems Laboratory of the University of São Paulo - LSI EPUSP.

[2] Visiting Professor in the Computer Science Department, University of Colorado at Colorado Springs (UCCS), CO, USA, with grant from CNPq, Brazil, Proc. 201443/93-7.

J. Rix et al. (eds.), *Virtual Prototyping*

it and interacts with it in a direct and natural way. VE applications have different characteristics, placing different demands over their supporting systems. Some of these applications involve a few participants in a shared virtual world whose 3D geometric representation is small enough to run in every node of the system, as for instance, a Computer Supported Cooperative Work - CSCW application. Other applications may use the virtual world for direct interference into the real world, such as telerobotics, which makes strict real time demands upon the supporting system.

This paper is concerned with the network requirements of yet another class of VE applications: large multi-party distributed virtual environment systems, which is characterized by the participation of thousands of users in an environment such as a virtual city, a battle field for military training, star war-like games and so on. All these applications involve a large number of participants in areas ranging from a few to thousands of miles, what makes difficult to keep the whole virtual world representation in every node of the system. Moreover, in order to add more realism to the environment, the VE is populated with simulated objects, what will bring in further demands over a communication network.

This paper is organized as follows: Section 2 presents the characteristics of large scale VEs. Section 3 describes the communication issues involved in a VE. Section 4 analyses the network requirements for large VEs. The integration of multimedia information to VE systems is discussed in section 5. The design of a distributed virtual environment is presented in section 6, followed by conclusions, acknowledgements and references.

2 CHARACTERISTICS OF LARGE DISTRIBUTED VE SYSTEMS

According to Sheperd (Sheperd, 1993), VE systems are classified in two broad classes: *telepresence,* where a common synthetic environment is shared among several people, as an extension to the CSCW concept, and *teleoperation,* where robots act upon a subject, be it a living body, a product being manufactured etc. In *telepresence* systems, maintenance of illusion is essential otherwise the sense of presence would fall in disbelief. In order to maintain illusion, the system has to provide for response time within an acceptable threshold. Besides, a participant actions have to be reflected immediate and simultaneously to all participants, and environments have to be realistic.

This implies in the following demands from the system, as summarized by VanDam (vanDam, 1993): rapid update rates and minimal lag; handling of multiple input devices; simulation of a potentially large number of interacting objects with complex behavior.

These requirements demand not only powerful tools for building the virtual world and its graphical representation. They also demand high performance graphic processors, support for the complexities and required realism of the environments to be rendered, as well as accurate and confortable 3D input and output devices, adequate operating systems and networks that can guarantee the delivery of the information in time.

The above requirements are even more stringent to a special group of applications within the telepresence class: *multi-user distributed large virtual environments.* These applications present the following characteristics:

- large population of participants and simulated objects.
- large areas, ranging from a few to thousands of miles.

- participants are geographically distributed.
- multidimensional environments.
- interaction among participants and objects in real time.
- heterogeneous computer and network resources.

Few large scale VE systems were implemented. Some of the best known systems are SIMNET (Calvin, 1993) and NPSNET (Zyda, 1992) which continue to develop towards a scale of thousands of participants. The building of large VEs is quite difficult, integrating several technological areas. This paper deals with the network requirements for large scale VEs, discussing their communication aspects and integration with multimedia information.

3. COMMUNICATION ISSUES IN A VE SYSTEM

3.1 The population of a VE

A virtual environment is populated with objects which are classified here as static, dynamic user-driven and dynamic simulation-driven. Large part of the world population is composed of static objects, i.e., objects that are fixed such as buildings, mountains, trees, parks, furniture etc. Simulation-driven objects can range from those programmed to act according to a pre-defined pattern (scripts) up to highly complex objects driven by expert systems, such as humanoids, robots etc. These latter objects can be developed by different companies and plugged into the VE. User actions are reflected on the VE through the user-driven objects. These may be represented as abstract elements (geometric forms), parts of a human body (hands, feet etc.), mobile vehicles, gloves etc.

3.2 The Basic Model of a VE System

A basic model of a VE system comprises the following modules: user application (APPL), virtual world simulator (VWS), 3D geometric database (3DDB), 3D rendering (visual, acoustic, tactile), and I/O devices.

The APPL is defined by the application designer who decides about the objects characteristics, the simulation-driven objects behavior, how all objects interact, if the synthetic environment is to be shared among users or if every user will have his/her own view of the VE (Isdale, 1993), and so on. A user should be able to customize the application by adding, dynamically, new objects or changing the characteristics and/or behavior of existing objects.

The simulation (VWS), is supposed to support the application and can be made in two levels: general simulation and object simulation (Araujo, 1994a). General simulation runs no stop and is responsible for the positioning of all objects in the area under its control as well as collision detection and resolution, concurrency disputes, objects translation etc. Object simulation is responsible for controlling the animation details of each dynamic object. The position of dynamic objects are calculated following physical laws. Information on dynamic objects position is sent only when unpredictable behavior occurs - as the dead reckoning concept used in SIMNET (Calvin, 1993). In a user-driven car, for instance, if the user keeps speed and orientation steady, the repositioning of the car does not need to be sent to all its replicas across the system, unless the user decides to change any of these parameters.

User actions are captured from sensors connected to them (gloves, head and body movement trackers etc.) or devices pointed towards them (cameras). The VWS translates this information into graphical and other commands which are sent to the 3D renderers for the repositioning process. The graphic renderer generates the final image to be presented to the user and the audio renderer "displays" audio which coincides with the location of the graphical objects. A user's actions, read from sensors, can trigger feedback actions reflected on other user(s) through actuators.

3.3 Distributed Nature of Large VEs

In a system where thousands of users can take part in the application, and where large areas are involved, distribution seems inherent. The geometrical information of an extensive area being synthesized, such as a city, may need to be divided in sub-areas which are distributed across the network. Environment servers - ES are responsible for the simulation and storage of sub-areas and their respective data set, as shown in Figure 1.

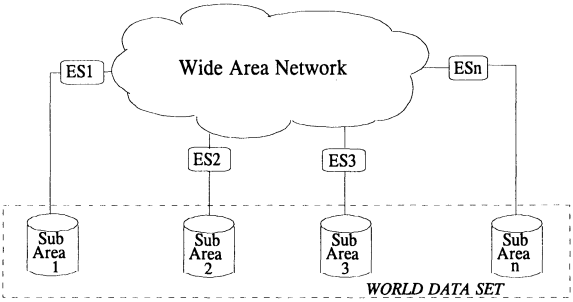

Figure 1 A Distributed VE System.

For visualization purposes, the world is stored in the environment servers in various degrees of resolution. Only objects within the user's view volume are rendered in full resolution.

The architecture to support such a system can vary from a peer-to-peer approach, client-server or a hybrid architecture. Sub-areas swapping, due to objects movements across their boundaries, may cause the loading of sub-area data sets located or maintained remotely. A hybrid communication architecture of a distributed VE system is presented in Araujo (Araujo, 1994a).

3.4 Communication in a Distributed VE

Considering a client-server architecture, where a sub-area, or part of a sub-area (area of interest) is downloaded to a client machine, replicas of the dynamic objects living in that area are also downloaded to the client machine. When there is more than one client machine for a same area of interest, i.e., there are other participants sharing the same virtual space, replicas of that area are downloaded to as many client machines as there are.

An action triggered by a participant, in a client machine, must have an immediate response. Moreover, this action must be broadcast to all client machines containing the same area of interest, so that all participants have a consistent view of the VE they are taking part of. The communication among objects in a VE involves basically the communication between an object and its replicas (communication type 1) and among dynamic objects (communication type 2). Figure 2 shows the flow of information in the two types of communication mentioned:

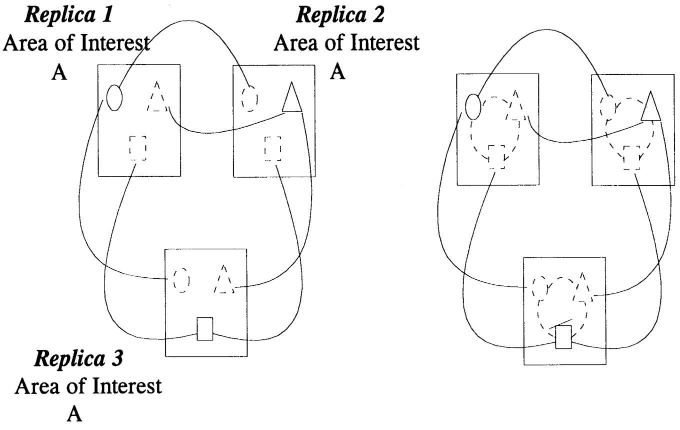

(a) Type 1 communication (b) Type 2 communication

Figure 2 Flow of information in the communication among objects in a VE

In figure 2, the squares represent three replicas of a same area of interest. The solid geometric forms inside each square represent either a user-driven object, whose user resides in the client machine housing that replicated area of interest, or a master copy of a simulation-driven object. The solid lines reflect the flow of messages from a dynamic object to its replicas, which are represented as dashed geometric forms. Consider, for instance, the solid circle as a dynamic object and the dashed circles as a dynamically modeled representation of that object in other two machines. Every movement performed by the object represented as the solid circle, which can not be predicted by the code running in the modeled representation (depicted as dashed circles), is sent to its replicas through a repositioning message.

For the communication among dynamic objects (for instance, when a participant, through a user-driven object, sings to other participants of the same area of interest), a simplification of the concept of auras (Benford, 1994) is used. Every dynamic object is assigned an aura per medium. An aura is an area surrounding an object whose range varies according to the medium reaching area it represents - the ranging area for audio may be different from the ranging area for video, for smell etc. An audio message sent, for instance, from user-driven object A, will be sent to all its replicas. When that message reaches the replicas, only the objects within the audio aura of user-driven object A replicas are affected by the message. The auras can be established dynamically by the communicating objects. Figure 2 (b) reflects the communication among dynamic objects. The dashed lines represent the communication among objects which is in fact realized through the solid lines.

Besides the communication among objects described above, communication is also realized for: synchronization between client machines and environment servers, objects movement from one sub-area to another, message exchanges in the initial loading of the areas of interest from the environment server to the client machine, concurrency disputes results broadcast by the server to its clients etc.

3.5 Communication Among Objects in Related VE Systems

The DIVE System (Carlsson, 1993) adopts a peer-to-peer approach for the interaction among objects. BrickNet (Singh, 1994) is implemented as a client-server architecture where objects, kept in different clients, are updated through different servers which together maintain information on all objects. It uses unreliable User Datagram Protocol - UDP protocol for the communication. In SIMNET (Calvin, 1993), every participant system has a copy of the whole VE. Objects communicate among themselves by broadcasting events through the Distributed Interactive Simulation protocol - DIS. The data units of this protocol - PDUs, related to the state of the objects, are transmitted periodically even when they present no change, as no centralized information about the objects states is maintained, leading to a waste of bandwidth. NPSNET also makes use of DIS but over the unreliable IP multicasting network - MBone (Macedonia, 1994a). It associates spatial, temporal and functionally related entity classes with multicasting groups (Macedonia, 1995). Araujo (Araujo, 1994b) makes a more detailed analysis of the communication approaches being developed by these systems.

4 NETWORK REQUIREMENTS FOR LARGE DISTRIBUTED VES

As far as networks are concerned, latency is perhaps the most important performance parameter for real-time communication systems. For VE applications, there is a time dependency as maintenance of illusion is crucial to the success of the system - the whole operation of handling user's input, computation, synchronization, information distribution, rendering and presentation of the rendered image, must occur in less than 0.1 seconds, which is the time a user takes to notice lag.

As all users need to have an immediate response to their actions and their actions must be seen by all users sharing the same area of interest, system success will depend not only on the processing power for the 3D graphics rendering and the simulation itself, which can be extremely heavy, but also on the network capability to deliver this information in time. This turns distributed VE systems into real-time systems and, as such, low network latency is an important requirement.

Causes for Delay

A VE has many sources of delay. The network delay itself is a major source of lag. The total network delay is contributed by:

- the propagation delay which is related to the distance and the signal propagation speed;
- the transmission delay which is related to the transmission speed (10 Mbps for the Ethernet, 100 Mbps for the FDDI etc.) and the message length;
- the protocol processing delay which is related to the processing speed of the switching nodes and end stations;
- the store-and-forward delay which is related to the number of hops, the transmission delay, and the number of packets for a segmented message;
- the queuing delay which is related to the congestion level at a switch node.

Considering a packet switched network, Fluckiger (Fluckiger, 1991), in a study based on CERN experience, measured round trips delays over TCP/IP lines, linking CERN labs to many European and American sites. It was shown that round trip propagation only took a few milliseconds even for intercontinental connections (less than 25 ms for a one way trip between CERN and Cornell over a 1.544 Mbps T1 line).

Attention must be drawn to the difference between high bandwidth and high speed. Fluckiger showed that round-trip delay over 10 megabits Ethernet is 1.5 ms against 1.8ms for 100 megabit FDDI network - this delay is caused by the FDDI protocol processing overhead.

Considering that large delays are caused by the handling of input/output information[1], simulation processing, network information processing, rendering and presentation, every time unit saved is precious. NPSNET-IV (Macedonia, 1994b) uses a four processor machine for increasing frame rates where simulation, rendering, drawing and network traffic handling are run in different processors.

As technology advances, faster network intermediate equipment are employed, making internal and end system delays smaller. On the other hand, better machine architectures are needed to exploit high speed networks, as even super computers can present bottlenecks to a gigabit network.

Optimizing Bandwidth Requirements

Supposing that every participating host has an update rate of 10 frames per second and that one message is generated every 0.1 second. Considering a message of 100 bytes (a DIS Entity State PDU has at least 140 bytes (Pratt, 1993)), the following traffic is generated by each user: 10 x 100 x 8 = 8 Kbits per second.

---

[1] 10 ms is the transmission delay for mice, trackers etc.; 20-100 ms for visual output; 1-500 ms for sound and directional sound - values suggested by (Ellis, 1994).

In an application with 100.000 users, a traffic of 800 Mbits/s would be generated. Macedonia (Macedonia, 1995) wrote that network simulations using multicasting showed a traffic reduction of 90% and that dead reckoning decreases traffic by 50%. By applying these percentages on our example, the final generated traffic falls from 800 Mbps to 40 Mbps. It has to be observed that the network traffic is highly dependent on the number of dynamic objects and their behavior complexity.

From the numbers derived above and considering the emergence of the broad band highways, at first sight, VE systems does not seem to place a high demand for bandwidth upon the network. However, multicasting, although being successfully developed, presents some difficulties such as: Existence of a fixed number of multicasting addresses to serve VE applications (large scale VEs may demand large numbers of group addresses). As the use of these addresses can be highly dynamic, address collision may happen if a global addressing scheme is not used (Schooler, 1992); Time to disconnect from a multicast stream is non-deterministic, causing host resources to continue to be consumed longer than necessary (Moran, 1992); Experiments over the IP multicasting network between the Naval PostGraduate School and SRI showed delays between 100 to 1000 ms (Macedonia, 1994b), what can be very high even for rates of 10 frames per second. Also, it is not widely used nor easy to be installed (Macedonia, 1994a).

As for dead reckoning, in a large population of dynamic objects, with complex and unpredictable behavior, it may not apply. There is also the "closet syndrome" mentioned by Pratt (Pratt, 1993) where, regardless the size of the closet, it will eventually fill up (referring to the bandwidth utilization of high speed networks such as FDDI and ATM in the support of multimedia and VE applications).

ATM as the Paradigm to Support Distributed VE Systems

Asynchronous Transfer Mode - ATM networks arise as a promising alternative to the support of large VE systems because of its low latency, bandwidth on demand and support for real-time traffic. ATM is a cell switching technology where data is carried in cells of 53 bytes. However, ATM networks are virtual circuit networks in nature what require communication channels to be set up before the ATM cells can be sent. The initial virtual circuit setup time can be several hundred milliseconds. To reduce the setup delay, pre-arranged circuit setup can be performed before the VE system starts.

Hayter (Hayter, 1991) proposed the Desk Area Network (DAN), an internal architecture which uses an ATM switch to provide interconnections between the machine components and from these components to an ATM LAN and/or ATM WAN so that seamless connections are achieved, what means lower latency. Several other initiatives exist for a more effective use of high speed networks.

5 INTEGRATING MULTIMEDIA TO VE SYSTEMS

Multimedia information will certainly be a desired complement to VE systems. In a virtual city, for instance, a participant can enter a cinema to see a movie or, through hypermedia, select objects and have multimedia information being displayed on demand about those objects, within the VE environment.

Many interactive multimedia applications are considered as soft real-time applications, which are defined as time dependent and tolerant to some amount of packet loss (Aras, 1994). Due to the inherent interactive characteristic of VE systems, they also present this time dependency and a certain tolerance to losses. However, it is important that some types of messages exchanged between environment servers and client machines be delivered reliably, for instance, messages conveying concurrency dispute results, messages exchanged among environment servers notifying object movements from one sub-area to another etc. Packet losses can be avoided but require time-out and re transmission mechanisms, what augments latency.

Network support for soft real-time systems has been extensively discussed in the literature, and is summarized as follows: interactive multimedia applications, providing services of voice and video, may need a certain guaranteed Quality of Service - QOS before being admitted to the network. The determination of the QOS depends on the characterization of the application traffic and this can be difficult to determine for both audio and video sources: they can vary over the time, and, in the case of video, be strongly dependent on the video coding algorithm employed. VE traffic is not easy to be characterized either. It can be very bursty, as shown in (Macedonia, 1994b), and have a model highly dependent on the complexity of the dynamic objects.

The integration of multimedia data in a VE system will raise further issues, such as: association of a particular medium channel to an object in the VE - for a virtual TV or radio broadcasting to user-driven objects nearby (Frécon, 1992); coordination between the presentation of the multimedia information embedded or associated to an object in a VE and the presentation of this rendered object to the users; matching of possibly different network architectures to support multimedia and VE systems - for instance, a VE network architecture may adopt a peer-to-peer approach (as in DIVE) against a client-server model for the multimedia data.

6 THE DESIGN OF A LARGE VE SYSTEM AND RELATED WORK

The first version of a distributed virtual environment system, involving a network of workstations and PCs is being designed and implemented at UCCS and UFSCar[2]. The system comprises a number of environment servers responsible for storage and control of VE sub-areas (Araujo, 1994a). Other machines connected to the network are used as client machines to the environment servers for both participation and visualization.

The emphasis here is on the distributed structure and network issues in the support of VEs with large populations and extensive areas with heterogeneous end user equipment. Systems like MASSIVE (Benford, 1994), is more concerned with the interaction aspects among users with heterogeneous equipment who communicate over an ad-hoc mixture of media. NPSNET (Zyda, 1992), one of the only really large scale VE systems implemented, do not address, as yet, support for heterogeneous equipment, as well as BrickNet (Singh, 1994), what is important if thousands of users are expected to populate a virtual world. Other systems like

---

[2] UCCS - University of Colorado at Colorado Springs; UFSCar - Universidade Federal de São Carlos.

DIVE (Carlsson, 1993) are now addressing simulation-like applications as the ones being considered here.

7 CONCLUSIONS

With the time dependency of VE applications, VE system success will depend not only on the processing power for the 3D graphics rendering and simulation, which can be extremely heavy, but also on the network capability to deliver this information in time.

This time dependency turns distributed virtual environments into real-time communication systems so that, a network to successfully support these systems must have the following characteristics: low latency with seamless connections, meaning low propagation time, low resource requirements within the network, low processing overhead in the handling of packets within the network and at the end system, high bandwidth, support for multicasting and good performance for a large number of connections. Networks as described above, allied to end systems that effectively use these high speed networks, are the candidates for supporting interactive large VE applications, specially when multimedia traffic is integrated to the VE system. ATM networks raise as promising alternative because of its low latency, bandwidth on demand and support for isochronous traffic.

A distributed VE system is being developed for PCs and workstations. The emphasis is on the distributed structure and network issues in the support to large populations with heterogeneous end user equipment and network resources.

8 ACKNOWLEDGEMENTS

We are particularly grateful to Dr. Edward Chow of the Computer Science Department at UCCS for his significant comments on this work. Thanks also to Professor Sebesta who made available to us the resources of the CS Department at UCCS.

9 REFERENCES

Aras, Ç. M. et alii (1994) Real-Time Communication in Packet Switched Networks Proceedings of the IEEE, Vol.82, No.1, January 1994, pp.122-139.

Araujo, R. B. and Kirner, C. (1994a) Distributed Large Scale Virtual Environments Discussion, Proposals and Design, paper not published.

Araujo, R. B. and Kirner, C. (1994b) An Architecture to Support Large Scale Virtual Environments, to be presented at the Brazilian Symposium on Computer Networks, 22- 26 May, Belo Horizonte, MG, 1995 (in Portuguese).

Benford, S.et alii (1994) Managing Mutual Awareness in Collaborative Virtual Environments, Proceedings of the VRST'94 Conference, 23-26 August, 1994, Singapore, pp. 223-236.

Calvin, J. et alii (1993) The SIMNET Virtual World Architecture, in IEEE Virtual Reality Annual International Symposium, September 18-22, 1993, Seattle, WA, pp. 450-455.

Carlsson, C. and Hagsand, O. (1993) DIVE - a Multi-user Virtual Reality System, in IEEE Virtual Reality Annual International Symposium, September 18-22, 1993, Seattle, WA, pp. 394-400.

Ellis, S. R. (1994) What are Virtual Environments, IEEE Computer Graphics & Applications, January 1994, pp. 17-22.

Fluckiger, F. (1991) From Megabit to Gigabit: Possible Transition Scenarios, in Computer Networks and ISDN Systems, Vol. 23, 1991, pp. 129-138.

Frécon, E., Eriksson, H. and Carlsson, C. (1992) Audio and Video Communication in Distributed Virtual Environments, Proceedings of the 5th MultiG Workshop, Royal Institute of Technology, Stocholm, December 1992.

Hayter, M. and McAuley, D. (1991) The Desk Area Network, in ACM Operating System Review, May, 1991.

Isdale, J. (1993) What is Virtual Reality? A Home-brew Introduction, electronic document in URL=ftp://sunee.uwaterloo.ca/pub/vr.

Macedonia, M. and Brutzman, D. P. (1994a) MBone Provides Audio and Video Across the Internet, IEEE Computer, April 1994, pp. 30-36.

Macedonia, M. et alii (1994b) NPSNET: A Network Software Architecture for Large Scale Virtual Environments, Presence Magazine, Vol.3 No.4, 1994.

Macedonia, M. et alii (1995) Exploiting Reality with Multicast Groups: A Network Architecture for Large Scale Virtual Environments, Submitted to the 1995 VRAIS Conference.

Moran, M. and Gusella, R. (1992) System Support for Efficient Dynamically-Configurable Multi-Party Interactive Multimedia Applications, Proceedings of the 3rd International Workshop on Network OS Support for Audio and Video, San Diego, November 1992, pp.143-156.

Pratt, D. (1993) A Software Architecture for the Construction and Management of Real-Time Virtual Worlds, PhD Dissertation, Naval PostGraduate School, Monterey, CA, June 1993.

Schooler, E. M. (1992) The Impact of Scaling on a Multimedia Connection Architecture, Proceedings of the 3rd International Workshop on Network OS Support for Audio and Video, San Diego, November 1992, pp.341-346.

Sheperd, B. J. (1993) Rationale and Strategy for VR Standards in IEEE Virtual Reality Annual International Symposium, September 18-22, Seattle, WA, 1993.

Singh, G. et alii (1994) BrickNet: A Software Toolkit for Network-Based Virtual Worlds, in Presence Magazine, Vol.3, No.1, 1994, pp.19-34.

van Dam, A. (1993) VR as a Forcing Function: Software Implications of a new Paradigm in IEEE 1993 Symposium on Research Frontiers in Virtual Reality, October 25-26, San Jose, CA, 1993.

Zyda. M. J. et alii (1992) NPSNET: Constructing a 3D Virtual World, Proceedings of the 1992 Symposium on Interactive 3D Graphics, Cambridge, MA, pp. 147-156.

10 BIOGRAPHY

REGINA BORGES DE ARAUJO received her B.S. degree in Computer Science from UFSCar-Brazil in 1983, and the M.S. in Data Communication Networks and Distributed

Systems from the University College London in 1986. She is now finishing her PhD in Electrical Engineering at LSI/EPUSP (due in October-1995). She has been a lecturer in the CS Department at Federal University of São Carlos since 1989. Her current research interests are in distributed systems and network support for multimedia systems and virtual reality. She is a member of ACM and SBC (Brazilian Computer Society).

CLAUDIO KIRNER received the B.S. degree in electrical engineering from EESC/USP-Brazil in 1973, the M.S. degree in electronic engineering from ITA - Brazil in 1978, and the Ph.D. degree in systems engineering and computer science from COPPE/UFRJ-Brazil in 1986. He has been an adjunct professor in the Computer Science Department at Federal University of São Carlos - Brazil since 1974. He is currently a visiting professor in the Computer Science Department, University of Colorado at Colorado Springs - USA. His research interests involve parallel and distributed computing, computer graphics, and virtual reality. He is a member of ACM, IEEE, and SBC (Brazilian Computer Society).