

Crossroads

The ACM Student Magazine

ACM / Crossroads / Xrds9-3 / Where Digital Meets Physical: Distributed Augmented Reality Environments

Where Digital Meets Physical: Computer-based Distributed Collaborative Environments

by *Felix G. Hamza-Lup, Larry Davis, Charles E. Hughes, and Jannick P. Rolland*

Introduction

Collaborative Virtual Environments are systems that transform computer networks into navigable and populated 3D spaces. Technological advances in optics and computer graphics, as well as the rapid development of network technology and distributed systems [8], open new doors for building effective distributed augmented reality systems. Such systems allow the development of interactive interfaces that will make space and time transparent for people located remotely, giving birth to new ways of interaction. This article presents a collaborative environment that is utilized as a test bed for interdisciplinary research in the Optical Diagnostics and Applications Laboratory (ODALab), co-located in the School of Optics and the School of Electrical Engineering and Computer Science at the University of Central Florida (UCF). The system described uses augmented reality techniques to improve human-to-human interaction. Central to this approach is an infrastructure for distribution of three-dimensional objects shared by interested communities of users. This sharing within an augmented reality increases their communication capacity adding new dimensions to their collaborations.

What is Augmented Reality?

Augmented reality (AR) systems are used to enhance the perception of the real world. Visually, this means that the real scene a person sees is augmented with computer-generated objects. These virtual objects are placed in the scene in such a way that the computer-generated information appears in the correct

location with respect to the real objects in the scene. AR can be classified along a virtuality continuum [6] (Figure 1).

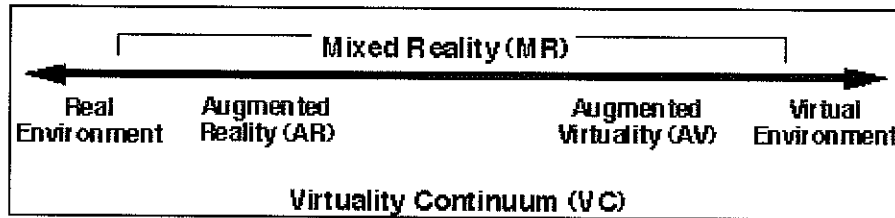


Figure 1: Augmented Reality's place in the Mixed Reality domain.

Distributed Augmented Reality Project

Envision a world where people from remote locations actively participate in a live three-dimensional demonstration instead of just watching a broadcast. Imagine being able to learn concepts by manipulating three-dimensional models that represent those concepts. Researchers push augmented environments and human-computer interaction technologies [2] beyond the current envelope of understanding and application by allowing the users of these environments to share knowledge through three-dimensional computer-generated objects embedded in the environment.

Distributed Augmented Reality Environment (DARE) is a computer supported collaborative environment based on AR. The system offers distributed access to electronic data and enhanced visualization. DARE system allows real-time, remote demonstrations. Several applications built on top of this system support these assertions about the system's capabilities. A prototype of the system being implemented was first tested for remote medical diagnostics and remote medical procedure demonstration.

The applications of the DARE technology can be categorized based on the bandwidth requirements:

- High communication bandwidth applications (above 50 Mbps): 3D Face-to-Face collaboration. In this application a stereo pair of video images as well as 3D sound is captured through the head mounted display and can be visualized remotely with 3D visualization hardware, such as, but not restricted to, a head mounted display.
- Medium communication bandwidth applications (below 50 Mbps), for example, medical diagnostics and procedure demonstration. The first

application will allow medical personnel to diagnose a patient at a remote location by visualizing the three-dimensional models generated from the magnetic resonance imaging (MRI) data. In this collaborative environment, medical personnel from geographically dispersed sites can visualize the same three-dimensional models and interact with the main site. In the procedure demonstration scenario, a team of experts will perform medical procedures, while at different geographically dispersed locations, the users of the system will be able to visualize in three-dimensions how the procedure evolves and will be able to interact with the main team.

Hardware system components

The ARC Display

The ARC Display represents a significant advancement towards the authors' vision for a multi-modal augmented reality system that includes 3D visual, 3D audio and haptic capability. The display consists of a curved, retroreflective wall, a teleportal head-mounted projective display (THMPD), a commercially available optical tracking system, custom designed optical probes, and a Linux-based PC. The retroreflective material affixed to the wall is also pliable enough to be attached to clothing while still maintaining its retroreflective properties. The retroreflective material is manufactured by 3M and uses micro-structures of either corner-cubes or beads of about 100 microns.




Figure 2: ODALab custom-built ARC Display.

The teleportal head mounted projective display

To see the computer-generated virtual objects and to allow remote virtual teleportation of a user to another collaborative site, users wear a teleportal head mounted projective display [3]. The THMPD, which was designed in the ODALab, takes advantage of a revolutionary set of lightweight optics. This set in its current implementation allows for a 52° field of view with optics that weigh eight grams per eye. The transparent beam splitter allows the user to see the real world and the virtual objects at the same time. The virtual 3D images are currently rendered on two 640 x 480 pixel liquid crystal displays (LCDs) encased inside the THMPD. Higher resolution systems are being implemented in the ODALab. The THMPD also allows stereoscopic face capture of a user via two mirrors strategically placed in front of the user and two miniature video cameras placed close to the head.

The computer-generated images are projected through the THMPD's optics onto the curved retroreflective wall. The beams of light that compose the images are retro-reflected (reflected onto the same path) to the user's eyes allowing him or her to see the three-dimensional virtual images superimposed on the real scene.

The tracking system

When the users move their heads, their viewpoint position and orientation will change. The 3D virtual objects are correctly rendered to each user's point of view, using the transformations supplied by the Northern Digital Polaris optical tracking system. This tracking system will not only track the user's movements, but also will track the movement of other objects in the real environment.

The graphic rendering systems

Each node is equipped with a graphical rendering device. Used here is a Linux-based PC, which applies transformations to the models in the virtual world and then renders the 3D scene to both left and right eyes. The PC is composed of dual, 1.5 GHz processors, 1 GB of RAM, and a GeForce4 Ti4600 with twin view, which allows independent rendering for the left and for the right eye.

In ideal AR applications, the user should not be able to distinguish the rendered virtual objects from the real objects in the scene. In actual fact, this problem is extremely difficult to resolve, primarily because of illumination issues, so deftly handled by nature for the real objects, but so hard to achieve in real-time for the virtual ones.

Application level

The DARE application layer uses a diverse group of application programming interfaces in conjunction with the ARC display. These include Open Inventor, Open Performer, and Java3D.

3D distributed visualization module

The software visualization module was built on top of Open Performer libraries and allows visualization of three-dimensional computer-generated images. Basically, the application allows professionals from different domains to share information and interact by manipulating three-dimensional models.

One or more users wearing THMPDs can be linked to a DARE node. The users can see a 3D model (anatomical model of a mandible in this case) from their own viewpoint and they can manipulate the object in space. Meanwhile, at another remote location, another user connected to a DARE node and wearing a THMPD can see the same 3D model.

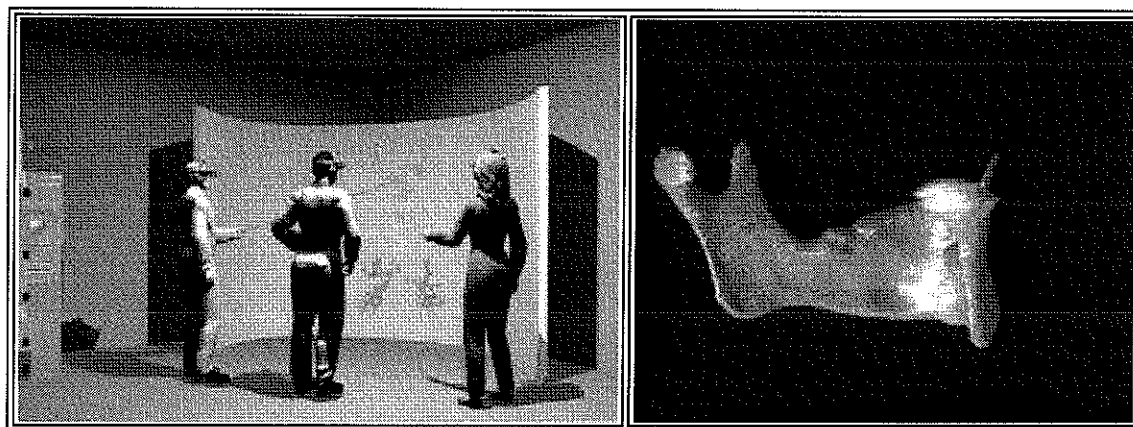


Figure 3: 3D visualization concept (left image, courtesy: S. Johnson, ODALab) and a 3D model of a mandible seen by one user (right image).

3D distributed procedure demonstration module

In the following scenario, at one location a team of experts performs a specific medical procedure. At the same time at several geographically dispersed locations, the users of the system will be able to visualize from their own viewpoint, in three-dimensions, the anatomical parts involved in the medical procedure and will be able to interact with the main team. Naturally, the position in the real scene of each anatomical 3D model has to be known. The position and orientation of the 3D models with respect to the user's viewpoint is computed with the help of an optical tracking system.

The hypothesis for this emerging technology is that the ability to visualize the dynamics of the 3D anatomical models in real-time during the medical procedure demonstration will improve learning for that particular procedure [5][7]. As a result, the emergence of an enhanced learning tool allows the sharing of knowledge in a way that has not been done before.




Figure 4: Intubation procedure concept (left image) and user's view of the 3D models of the lungs and trachea superimposed on the patient (right image).

System level

Software components

The main purpose of the DARE is to distribute data among the users of the system such that each client is able to recreate and visualize locally 3D scenes from their own unique viewpoint. The current architecture of the system is client-server. A novel architecture that will allow easy scalability of the system is currently being developed. This architecture is a hybrid between the traditional client-server and peer-to-peer decentralized schemes.

The new architecture of the systems is based on clusters. A cluster is defined as a subset of client-server nodes (CSNs) that has associated with it a tracking system (TK), a database system (DS) that will allow the 3D models persistency, and a cluster monitor (CM) (**Figure 5**).

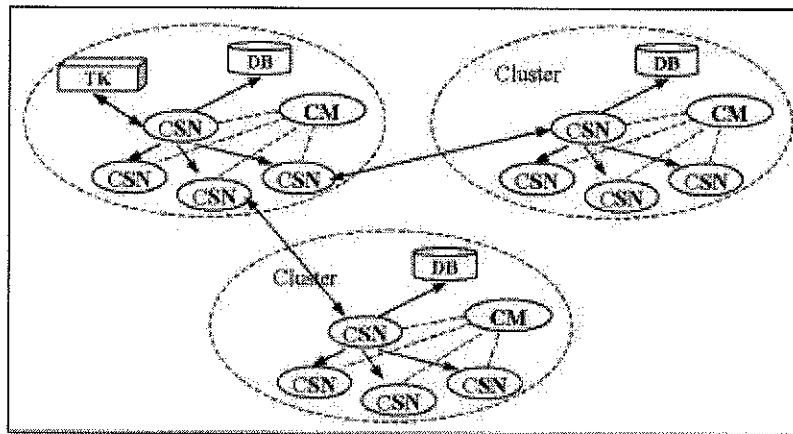


Figure 5: A snapshot in time of the DARE system.

CSNs represent the users who want to participate in the demonstration. They are named client-server nodes because they have the ability to function as a client or as a server, in case bottlenecks appear in the cluster or in the case of failures of the cluster components. In server mode, the nodes continuously interrogate the tracking systems or other parent server nodes for position information and make the information available for the clients. Clients update their local 3D environment allowing each user to see the real scene augmented with the 3D models at the current position and orientation.

Since the most important system parameters are reliability and low latency, this hybrid architecture will allow distributed control of the system performance parameters and good scalability.

The number of CSNs in a cluster is dictated by the performance requirements of certain applications for a specific domain. The gathering and analysis of the performance parameters is the main task of the cluster monitor. The CM is responsible for the Quality of Service (QoS) parameters management, requested by the users of the system and by the application that is running at each node. One of the capabilities of the CM is to invalidate some of the CSNs in the cluster to maintain the QoS parameters. Each CSN will have a priority associated with it.

The tracking system (TK) is used in the case of applications that require tracking [1], for example the remote medical procedure demonstration. TK gives precise information regarding the objects' locations and the viewpoint of the user who performs the procedure. The server component of the CSN will pull this data from the TK or from its parent CSN at interactive speed and will make this data available to its clients. The CSN, in client mode, will use the tracking data to render the virtual 3D models at correct positions in the real scene, allowing remote three-dimensional visualization during the procedure demonstration.

The Database System (DS) assures data persistency and 3D models storage. Some of the procedures or demonstration data can be permanently saved in the database. In this way, offline (asynchronous) demonstrations can be requested by a CSN in client mode. Another use of the DS is to allow users to download the 3D models before the demonstration starts. Since the requirements for the bandwidth and latency are very stringent, we assume that at demonstration time the 3D models used are already available at each node. The only data that are synchronously sent, at interactive speed, are the scene parameters and/or the position/orientation information from the TK.

Results: Achieving the application domain requirements

Scalability

Because it is a distributed system, DARE must scale well [4] and its performance must not degrade as more users are added to the system by joining a demonstration. Adding a new user to this system will basically mean the creation of a new CSN that can be added to the current cluster or can trigger the creation of a new cluster.

Bandwidth and latency

The most stringent requirements are the ones imposed by the soft real-time attributes of the system [9]. For real-time collaborative applications, interactive speed is necessary. To achieve interactive speed, the delays must be limited to the range of ten to a few tens of milliseconds. The first category of the DARE applications, remote visualization, does not have high bandwidth requirements; however the delays must be bounded. In this case, to achieve

the interactive-speed requirements over different domains, the system distributes only the tracking data. The amount of data is very small and allows medium bandwidth nodes to join the system. However, the packet delay over several hops might still be high, and this will have a negative impact on the interactivity of the application.

In support of this real-time system, the authors have performed preliminary tests using C/C++ TCP and UDP sockets. A 100 Mbps network connection has been used, knowing that success here will guarantee future success as higher bandwidth connections are rapidly becoming widespread. The hardware consists of:

- Optical tracking system: Polaris from Northern Digital with a maximum data rate of 115 KB through a RS-232/422 interface having a maximum update rate of 60 Hz
- Client machine: a 1 GHz PC running RedHat Linux 7.2 OS with 512MB of RAM, and a GeForce 4 twin-view graphics card
- Server machine: a 1.5 GHz PC running RedHat Linux 7.2 OS with 512MB of RAM, and a GeForce 4 twin-view graphics card, linked to the optical tracking system

For data distribution, an implementation of the client-server architecture using TCP and UDP sockets was tested. Two kinds of tests were performed:

- local tests: the server and the client processes are running on the same machine
- remote tests: the server and the client processes are running on different machines located in the same LAN

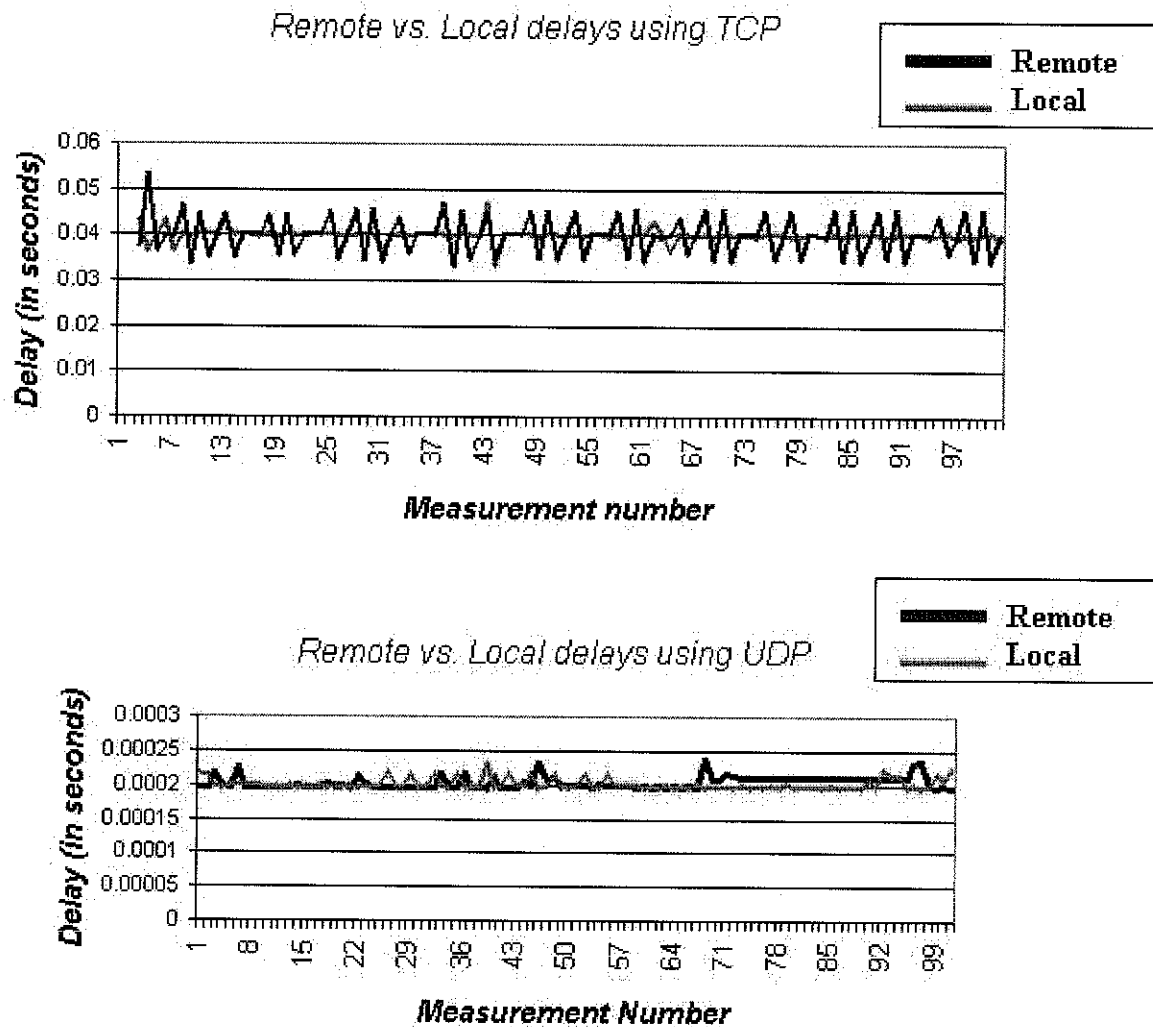


Figure 6: Delays in the current implementation using TCP and UDP.

As expected, by using UDP sockets the delays are orders of magnitude smaller than in the TCP case. TCP, on the other hand, provides a connection-based, reliable data stream. Moreover it guarantees delivery of data and also guarantees that packets will be delivered in the same order in which they were sent. In-order packet delivery is extremely important when the scene is rendered based on the tracking data because it ensures correct rendering of the virtual objects in the scene.

Conclusions

Computers and computer networks are already the primary vehicles for the distribution of information. They are also the main ingredients available to us for going beyond information dissemination, creating powerful knowledge distribution systems. Since knowledge is embedded in people and, unlike

information, occurs in a process of social interaction, this work focuses on the development of distributed collaborative environments that enhance human-to-human interactivity. Technological advances in optical projection and computer graphics allow us to augment reality with computer-generated three-dimensional objects. Moreover, the distribution of these three-dimensional objects at dispersed locations allows efficient communication of ideas through three-dimensional images.

The DARE framework and early results presented here are extremely promising. The average delay due to the system components is 40 milliseconds using connection-oriented communication. This allows development of interactive applications on medium to low bandwidth networks. The authors envision a global coverage of platforms like DARE with a universal augmented world in which three-dimensional models, information, and services can be integrated as easily as pages in the World Wide Web.

References

1

Barfield, W. and Caudell, T. *Fundamentals of Wearable Computers and Augmented Reality*. Lawrence Erlbaum Associates Pub., Mahwah, NJ, 2001.

2

Billinghurst, M. et al. *Projects in VR - Real World Teleconferencing*. IEEE Computer Graphics and Applications, Vol.22/No.6, Nov/Dec, 2002.

3

Biocca, F. and Rolland, J.P. *Teleportal face-to-face system*. Patent Filed, August 2000.

4

Coulouris, G. and Dollimore, J. *Distributed Systems Concepts and Design* 3rd edition Addison Wesley, 2001.

5

Feiner, S. *Augmented reality: A new way of seeing*. Scientific American 54, April 2002.

6

Milgram, P. and Drascic, D. *Perceptual issues in Augmented Reality*. SPIE 2653, pp123-134, 1996.

7

Rolland, J.P. et al. *3D visualization and imaging in distributed*

collaborative environments. IEEE Computer Graphics and Applications, 2002.

8

Tanenbaum, A. and Steen, M. *Distributed Systems - Principles and Paradigms*. Prentice Hall, 2002.

9

Verissimo, P. and Rodrigues, L. *Distributed Systems for System Architects*. Kluwer Academic, 2001.

Biographies

Felix G. Hamza-Lup (fhamza@cs.ucf.edu) is a Ph.D. student in the School of Electrical Engineering and Computer Science at the University of Central Florida (UCF). He received his M.S. in Computer Science from University of Central Florida in 2001 and the B.A. in Computer Science from Technical University of Cluj-Napoca, Romania, in 1999. In 2001 he received the Java Developer Certification from Sun Microsystems. Felix is also a member of the International Honor Society for the Computing Sciences, SPIE, and ACM. His research interests are: distributed systems and mixed reality.

Larry Davis (davis@odalab.ucf.edu) is a Ph.D. student in the School of Electrical Engineering and Computer Science at UCF. His doctoral research focuses on conformal tracking for virtual environments. Larry's other research interests include virtual environment development and head-mounted display design. He is a student member of the ACM and IEEE.

Charles E. Hughes (ceh@cs.ucf.edu) is Professor at the School of Electrical Engineering and Computer Science at UCF. He received the B.A. in Mathematics from Northeastern University in 1966, and a M.S. and Ph.D. in Computer Science from Penn State University in 1968 and 1970, respectively. He served on the faculties of Computer Science at Penn State and the University of Tennessee prior to joining UCF in 1980. His research interests are in mixed reality, distributed interactive simulation, and models of concurrency. He is a member of the ACM, IEEE, and IEEE Computer Society.

Jannick P. Rolland (rolland@odalab.ucf.edu) is Professor at the School of Optics and the School of Electrical Engineering and Computer Science at UCF.

She received the Diploma from l'Ecole Superieure D'Optique in 1984, and her Ph.D. in Optical Science from the University of Arizona in 1990. She served on the Faculty of Computer Science at the University of North Carolina at Chapel Hill from 1990 till 1996 before joining UCF. Dr. Rolland is Associate Editor of Presence (MIT Press) since 1996, and Associat Editor of Optical Engineering since 1999. She is the UCF Distinguished Professor of year 2001 for the UCF Centers and Institutes. She is a member of the OSA, SPIE, and IEEE.

Copyright 2004, The Association for Computing Machinery, Inc.

