

Vaudeville: A High Performance, Voice Activated Teleconferencing Application

Jyoti K. Parwatikar, A. Maynard Engebretson, T. Paul McCartney,
John D. DeHart and Kenneth J. Goldman
Department of Computer Science
Washington University
St. Louis, Missouri 63130
{jp, ame, paul, jdd, kjg}@cs.wustl.edu

ABSTRACT

We present a voice-activated, hands-free, ATM-based video conferencing application. The application, called Vaudeville, features high quality NTSC video, voice-activated audio transmission, audio bridging of two audio streams, and voice-activated video switching. It supports multiple simultaneous multi-party conferences using a scalable multicast mechanism. Users can move freely between conferences without knowledge of the underlying network structure. We describe how Vaudeville was built using a component-based distributed programming environment. We also describe the algorithms used to control the audio and video of the application. Audio and video are encoded in hardware using a platform-independent, ATM hardware multimedia interface.

KEYWORDS: distributed computing, multimedia, video conferencing, ATM

1 INTRODUCTION

The goal of this research was to create a teleconferencing application using ATM that facilitates a natural conferencing environment that approaches the feeling a real-life conference in a conventional meeting room. Requirements included high quality video, a simple intuitive user-interface, hands-free operation, scalability, and support for navigation among multiple simultaneous conferences.

This paper describes Vaudeville, a teleconferencing application that we have constructed based on this goal. In Vaudeville, audio transmission is voice-activated, and the video automatically switches to reflect which person is speaking. The voice-activated audio means the user does not need to make any indication to the system that they wish to speak other than making noise. The video switching mimics the pattern of focus of an average conference participant. Also since the video is studio quality NTSC (30 frames/sec), the video is similar to what a participant would see if they were sitting in the same room with the speaker. Audio and video is described in more detail in section 1.2.

To allow conferences to be scalable, a two speaker paradigm is used. At any given moment, a conference has two designated speakers, a primary speaker and a secondary speaker. The designated speakers may change dynamically throughout the conference with voice-activated switching. Video and audio streams are associated with each speaker. The video of the primary speaker is shown by default and the audio streams of the two speakers are mixed. Two speakers are used rather than one to allow the case in which one person is speaking and one person wishes to interrupt. We do not have more than two speakers at once, because presently the hardware only supports the bridging of two audio streams. In practice, we have not seen a need for more than two speakers. Contention between audio streams is resolved in hardware, this resolution is described in section 1.2. Video for a participant is transmitted only when they are a speaker (either primary or secondary), and independent of the size of the conference, only two videos are transmitted per conference. Participants also have the opportunity to see the video transmission of anyone in the conference by explicitly selecting that participant. This feature uses additional bandwidth at the viewer and the node being viewed, but does not affect the bandwidth of the rest of the conference.

The Vaudeville software was developed using The Programmer's Playground, a distributed programming environment. A brief overview of Playground is described in Section 1.1 of this paper. Details may be found elsewhere [4,5].

Unlike the MBone tools[8,9,10,11], we decided to concentrate on high quality video. There are two reasons. First, we believe that the high quality video makes one less conscious that the conference is virtual. Second, for conferences in which participants want to use a video source as a center of discussion it is important for video to be high quality especially for medical applications. Like the video MBone tool *vic* [8], the video is voice-activated, but, we do not have the ability to show several video streams at once.

Other video conferencing applications using ATM transmit audio and video over an IP network on top of an ATM network [12,13]. In our case, a hardware codec, the MultiMediaExplorer [7], processes the video and audio streams directly from the ATM, and no data copying is required. Any extra information added to the data stream beyond the raw data is introduced by the codec for use in the audio and video algorithms.

This paper is organized as follows. Section 1.1 describes the programming environment. Section 1.2 describes the hardware testbed. Also included in section 1.2 is a description of the audio algorithm and the ATM configuration of a conference. In section 2, the basic functionality of the application is described giving the features of the application. The software components and configuration are described in section 3. In section 4, experience with the application is described as well as changes made as a result of the experience. Finally, section 5 provides conclusions and directions for future work.

1.1 THE PROGRAMMING ENVIRONMENT

Vaudeville is implemented in the context of *The Programmers' Playground*, a software library and run-time system that supports a new programming model for distributed applications [3,4,5]. The model, called I/O abstraction, provides a separation of computation from communication structure, making it well-suited for end-user construction of customized distributed applications from computational building blocks. Playground users do not need to write any source code to establish communication between the modules of a distributed application, nor do they need to understand the details of how communication occurs.

A distributed application consists of a set of communicating modules, written as C++ programs, using special data types from the Playground library. The library supports several standard data types such as integer, string, character, and mappings. Along with these data types, Playground supports a structure called a *tuple* that allows the combination of data types to form a user-defined type. Instances of these data types can be *published*, making information available to external modules. The communication structure among the modules of a Playground application is defined by a set of logical connections among published variables of the modules. The communication structure is configured dynamically at runtime through a graphical user interface for the connection management system. In our case, Vaudeville itself acts as a front-end for the Playground connection management system.

The connection management system consists of two parts, a protocol run by each Playground module and a connection manager that is a Playground module itself. The connection manager does type and protection checking when a connection is made, and informs the modules involved in the connection of the new link. After the modules have been informed of the new connection, all further negotiation and communication is handled individually by the endpoint modules' protocols.

Communication among modules of a distributed system occurs implicitly. When a published variable is modified, the new value of the variable is automatically sent to its connected variables. In this way, the programmer does not need to know the low level details of how the communication is actually achieved. Each module can be created independently of the modules with which it communicates.

The control portion of a module defines how its state may change. Each module may have both active and reactive control. Active control carries out the ongoing computation of the module. Reactive control carries out computation in response to changes in the values of published variables.

Prior to Vaudeville, all Playground communication was over TCP/IP using Unix sockets. To support ATM connections, the ATM connection management system was combined with the Playground connection management system. Also, audio and video data types were added to the Playground runtime library. If a connection is made between audio or video variables, the combined connection management system makes an ATM connection between the two sites.

1.2 HARDWARE TESTBED

Vaudeville was built using the Project Zeus ATM testbed[1] and the MultiMediaExplorer (MMX)[7]. The MMX is an ATM desktop multimedia hardware interface that is host independent, and requires no computing resources from the host workstation for processing of multimedia streams. It is currently controlled over a RS-232. It allows users to simultaneously transmit and receive audio and video. Video is full-motion, JPEG compressed, NTSC video (60 fields/30 frames per second at 640x480 resolution). The MMX has two audio modes, a default mode and a conferencing mode. The default is CD-quality stereo audio. In conferencing mode, the audio is voice-quality monaural audio.

Vaudeville uses the MMX conferencing mode. The main features of the conferencing mode are the audio-bridging of two streams and voice activated switching. Each MMX is assigned a source id by the application. The id is placed in the payload of every audio cell it transmits. On the receiving side, each MMX then chooses two audio streams from all streams received to bridge and play out. This choice is made based on the source ids of the streams. Smaller source ids have higher priority. Once the two streams have been chosen, the source ids of the streams are written into the MMX's global memory where they are accessible by the application software. Audio cells are transmitted only if a signal is present, and at least one of the audios presently chosen has a priority less than that of the transmitting MMX. A signal is considered present if the signal is larger than the voice activated threshold. This threshold can be set by an outside application.




Figure 1: ATM Configuration with Point-to-Multipoint Connections

Since the ATM software does not presently support many to many multipoint connections, a point-to-multipoint connection is created for each member of the conference. This connection is made to everyone in the conference including the transmitting member as shown in Figure 1. This ensures that all MMX's receive the same audio streams, and hence all MMX's choose the same two audio streams to play out. Speakers transmit their video on this same point-to-multipoint connection. To receive video from a non-speaker, a separate point-to-point connection is setup between the transmitter and the receiver.

Ideally, with many to many connections, we would use the ATM configuration model as shown in Fig-

ure 2. There are three multipoint to multipoint connections. One connection is for audio, and two are for video one for each of the speakers. The transmitters in the video connections change as the speakers change.




Figure 2: ATM Configuration with Many to Many Multipoint Connections

2 FUNCTIONALITY

Figure 3 shows the Vaudeville direct manipulation user interface. The user interface can be divided into four parts -- the conference menu, the people directory, the file menu and the conference window. At the top, is the conference menu that displays the active conference name, and has a pull-down list of all currently available conferences. The people directory is the list of names on the right of the figure. This is a list of all of the people currently in the system. The file menu is a pull-down menu that appears when File is selected at the top left of the interface. The rest of the interface constitutes the conference window.

Joining a Conference: A conference may be thought of as the virtual equivalent of a meeting room. Users may join a conference by selecting the name of the conference from the conference menu. This causes the user to be removed from any conference that they are currently attending, and to be added to the new conference. By selecting the New Conference option from the File Menu, the user may create a new conference. When this option has been selected, a dialog box appears that asks for the name and attributes of the new conference. After the conference has been created, the user automatically joins this conference. Users may leave a conference by selecting the Leave Conference option from the File menu.

While participating in a conference, the display shows the conference's current state. Each member of the conference is represented by a photo icon. The two current speakers are highlighted with a colored border; the primary speaker in red and the secondary speaker in gray.

Viewing the Conference: Video is viewed in a separate window or on a separate display. There are two options for receiving video. The first option is to receive the default video. In most cases, the default video is the primary speaker's video. However, for the primary speaker the default video is the secondary speaker's video. As speakers change, the video stream selection also changes. When the Become Conference Video button has been selected, the default video for the whole conference is the video belonging to the person who selected the button. This button represents a token belonging to the conference. To release the token, the user deselects the button. When no one holds the token, the default video follows the speakers. Anyone may request the token, and the token moves to whomever requests it. Users do not need to explicitly release the token in order for another user to request it. We rely on conference etiquette for floor control rather than imposing some locking mechanism.

The second option is an override option. The user selects the video received by selecting the icon of a person in the conference. At this point, the default video remains unchanged for the rest of the conference. The selected video is seen only by the user who selected it. To return to the default option from the override option, the user deselects the icon of the received video.

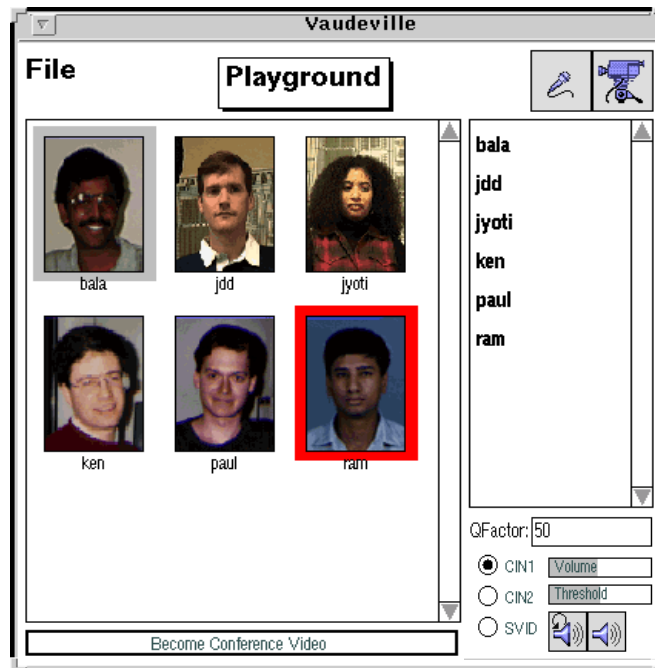


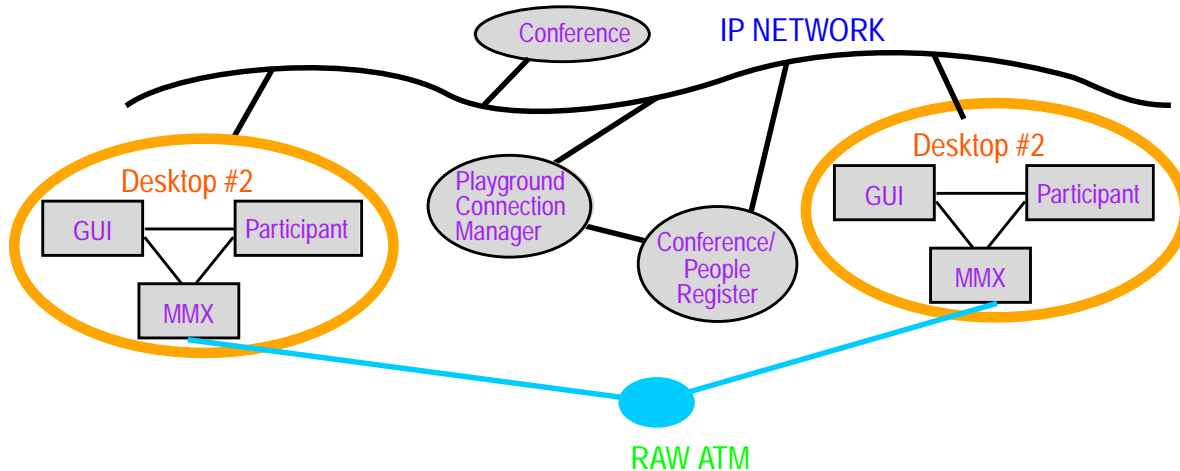
Figure 3: Vaudeville User Interface

Settings: The set of controls in the bottom right hand corner of the interface allow the user to control the quantization factor (q-factor) for the conference, the threshold, the playout volume, mute, selfmute and the video source. The q-factor is the JPEG encoding and decoding compression factor for the entire conference. This can be changed by any participant by typing in a new value for the q-factor. The threshold slider controls the threshold of the voice activated switching. The threshold is the level of sound that must be produced before an audio stream is transmitted. The mute buttons are the two loud speaker icons in the lower corner, these control the audio heard by the user. The selfmute button is the icon with the circular loop. When this button is selected the user does not hear their audio stream. When the other mute button is selected the user does not hear any of the audio streams. CIN1, CIN2, and SVID are radio buttons that allow the user to choose between any one of three different video sources.

The microphone and camera buttons in the top right-hand corner of the interface allow the user to control outgoing audio transmission. When transmission is turned off, a red slash appears through the icon. Note that enabling video or audio only makes reception possible. It does not necessarily imply that others actually have chosen to see or hear the video or audio. This is useful if a user wants to take a phone call in the middle of a conference, or just wants to observe the conference and not actively participate.

Invitations: After becoming a participant in a conference, a user may invite others to the conference by selecting a person's name from the people directory. At this time, the name changes from black to green and changes back to black when a reply has been received. On the receiving end of the invitation, an invitation window is created informing the user by whom they have been invited and to which conference. The window also contains buttons which allow the user to accept or reject the invitation. If the user accepts, the user interface responds in the same way as if the user had explicitly chosen to join the conference.

Figure 4: Vaudeville Software Components



3 SOFTWARE COMPONENTS

Vaudeville consists of five Playground modules shown in Figure 4: the *graphical user interface (GUI)*, the *MMX module*, the *Participant module*, the *conference register*, and the *conference module*. Each user has an GUI, a MMX and a Participant module. There is one centralized conference register, and there is one conference module for each active conference. The MMX module pipes commands to the MMX based on changes to its published variables made by the GUI and the Participant. It also outputs the source ids of the two audio's currently being played. The GUI outputs information supplied by the user to the MMX and Participant modules, and reflects any changes in the conference state. The Participant module's main function is to coordinate connections among the published variables of the five modules, and to control which video is played based on information supplied by the MMX module and the GUI.

For each conference in progress, there is a conference module that maintains the state of the conference and the members of the conference. Each member of the conference has both an incoming and an outgoing connection to the conference module. The outgoing connection supplies the conference with the participant's state information and any request of resources managed by the conference. The incoming connection to the participant supplies the modules with the state information of the conference. The conference register maintains lists of conferences and people maintained by the Vaudeville system. The contents of this register are displayed in the GUI.

The environment can be divided into two parts the desktop environment and the outside environment. The desktop environment consists of the MMX, the video monitor, the microphone, the camera, the speaker and the user. Two modules interact directly with the desktop environment the MMX module which interacts with the MMX, and the GUI module which interacts with the user. The Participant module acts as a bridge between the desktop and the rest of the environment. It coordinates this outside information with information from the GUI and the MMX module. The outside environment consists of devices, users and other Playground modules. Here reside the conference, the conference register, the connection manager and other users of the system each with their own MMX, GUI, and Participant modules.

3.1 Desktop Modules

MMX Module. The MMX module shown in Figure 5 has two functions, one active and one reactive. The active function polls the global memory of the MMX to read the ids of the two audio streams being played out. The reactive function pipes commands to the MMX when requested. The MMX module's interface consists of three published tuples. The first tuple acts as an interface to the GUI module, and is writable by the environment. It controls parameters which affect only the desktop area namely volume, mutes, video source, and threshold. The second tuple is also writable by the environment, and acts as an

interface to the Participant module. It controls MMX functions which may depend on information provided from beyond the desktop. It consists of parameters corresponding to q-factor and a request to open/close a receiver/transmitter. Associated with each of these tuples is a set of reaction functions. When either of these tuples changes, the reaction function causes a command to be piped over an RS-232 serial line to the MMX to make the physical change requested as shown in Figure 5.

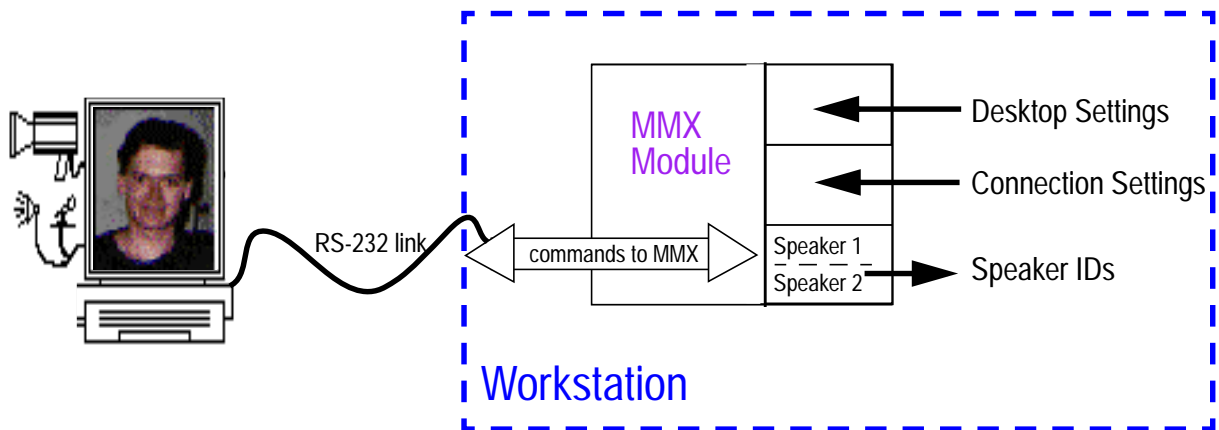


Figure 5: MMX Module

The third tuple is readable by the environment and provides the ids of the two speakers. Periodically, the module polls the global memory of the MMX, and reads the ids of the two audio streams being played out. If there is a change in the ids, the module updates the tuple to reflect the most recent set of speakers in the following manner. If there are two ids in the MMX's global memory, the module updates the tuple with the two new ids. The primary speaker is the highest priority id, and the secondary speaker is the other id. If there is only one id in the MMX's global memory, the module updates the tuple with the new id as the primary speaker, and the next highest priority id from the old set of speakers as the secondary speaker. If no ids are in the MMX's global memory, the tuple remains unchanged. The information provided by this tuple is used only by the video algorithm, and has no effect on the audio algorithm

GUI Module. The GUI module manages the user interface shown in Figure 3. Its purpose is to pass information supplied by the user to the rest of the system, to keep the user informed of the state of the system, and to allow the outside world to communicate with the user. The interface of the GUI module consists of three published tuples. Fields of the tuples correspond to elements of the user interface. Two of the tuples are readable by the environment. These tuples are updated as the elements associated with them are manipulated by the user. Elements associated with the writable tuple are updated whenever a change occurs in the tuple.

The first readable tuple controls the local multimedia parameters, mutes, volume, etc. The second readable tuple allows the user to communicate with the system and make requests for resources held outside of the local environment. This includes invitations, conference selection, and which video to play out. The writable tuple maintains state information about the system and also allows the outside environment to communicate with the user. It includes changes in the conference/people register, and changes in the conference state.

Participant Module. The Participant module is the controlling module. It coordinates information about the whole environment and takes the appropriate action. This is where most of the decision making is done. There are four published tuples. The first is a readable/writable tuple that acts as an interface to the MMX module. This tuple maintains the two current speakers, and represents requests for opening or closing receivers or transmitters. The second tuple reflects the user's current preferences. This includes which

video to play out, and which conference to join. This tuple is a writable tuple that acts as an interface to the GUI module. The third tuple is a readable tuple linked to the Playground Connection Manager. This tuple represents a request for links between modules. The fourth tuple is both readable and writable, and acts as an interface to the conference. The two main functions of this module are managing the video, and making connections between modules.

The Participant Module is responsible for deciding which video to play out and when to transmit video. Each Participant module maintains two video transmitters. When the user is a speaker, video is transmitted along the same virtual path as the audio. Otherwise video is transmitted from a second transmitter along a separate virtual path. When another person wants to view the user, a link is set up to the viewer from this second transmitter. This ensures that the incoming bandwidth attributed to the video is not a function of the number of conference participants. It also makes the conference scalable, because the number of incoming video streams remains constant as the size of the conference increases. When the user becomes a speaker or is no longer a speaker, the Participant module reacts by switching the video transmitter.

As in the case of the transmitters, the Participant module also maintains two types of receivers, one for viewing speakers, and one for viewing non-speakers. Each receiver represents a separate virtual path. If the user is viewing the default video, then each time the highest priority speaker changes the Participant module switches the video receivers to the video of the new speaker. Note this switch would not be necessary with many to many connections. The module may also have to change receivers in the case when the user is not viewing the default video. If the person the user has selected to view becomes or stops being a speaker, the module will switch between the two types of video receivers.

The Participant module is responsible for the configuration of the desktop modules under its control. Upon initialization, connections are made for the initial configuration of the desktop modules. When joining a conference, connections are made to connect the module to a conference. This includes a connection to a Conference module, and the connections to the audio calls associated with the conference. Connections are broken upon leaving a conference. The only other connections made are to support the reception of video associated with a non-speaking source.

3.2 Shared Modules

Conference Module: The Conference module maintains the state of the conference, and manages any parameters or resources that belong to the conference. It has a published list of members that is both readable and writable by the environment. Each time someone new joins the conference, an element is added to the list. Each element of the list is connected to a published variable in the corresponding Participant module. Through the elements of the list, the module assigns source ids to members, and receives requests for resources and changes to conference parameters. Currently, source ids are assigned randomly.

Conference/People Register: The Conference/People Register monitors the system, and maintains a published list of the available conferences and the people currently running Vaudeville. It has two published lists. One list is writable by the environment. This is a list of all Playground applications currently known to the Playground connection manager. Using this list, the module filters out the applications associated with conferences and people. Then it places the names and ids of these applications into a second list which is readable by the environment.

4 EXPERIENCE

Vaudeville has been running in the department since April 1996. It is currently running on both SGI and Sun platforms. All features described have been implemented. Because of the high quality video and hands-free floor control of audio, conferences have succeeded in being natural. Users have been taught to use the system in a matter of minutes, and with enough incentive (a good movie playing) they have

taught themselves. By using different video sources such as a microscope and a Silicon Graphics machine, participants have been able to meet and discuss work from their own labs and offices without any special preparation.

We discovered a few audio problems in the initial testing. At first, the range for the voice activated threshold and the default setting was too low. In large conferences, the audio was chaotic because the streams were too easily cut off by interruptions from higher priority sources. Increasing the default and maximum thresholds solved the problem. Another problem was audible pops and clicks. When a speaker paused between words, the voice activated switching caused the transmitter to be switched off which produced a click. This became a real distraction to users, so the algorithm used by the MMX to control the voice activation was changed to leave the transmitter on if there was only one speaker, and to fade out the audio when the transmission was cut. This greatly improved the audio quality of the system.

People almost immediately decided that they didn't want to have all of their conferences open to the public, we added private "invitation-only" conferences that would not appear in the conference directory. To allow people to know in which conference a person could be found, we added a feature to the people directory that prints the conference name a person is in next to their name.

Our first regular departmental faculty meeting using Vaudeville involved fourteen sites. Thirteen sites were individual offices, and one was a large room with a room microphone. The video switching worked well. One minor problem was the threshold setting. Since many of the participants had little to no experience in using the system, they did not know how and when to adjust the audio threshold. As a result, some participants would cut in if they shifted position or cleared their throat. This was most problematic for the site with the room microphone, since it picked up any sound in the room. Another problem was the variety of different transmission volumes. Some participants sounded louder than others. There were also some human factor problems. Participants complained of feeling "guilty" if they laughed or said something insignificant, since they caused the video to switch. Most people commented that they would have liked to have been able to simulate a whisper by initiating a side conversation. Another comment was that it would be nice to have a way to scan the audience randomly to be able to gauge non-verbal reactions and see facial expressions. Many felt paranoid at not being able to know when someone had selected their video. However, despite these problems, the meeting lasted over an hour, and useful work was accomplished.

5 CONCLUSIONS AND FUTURE WORK

We have presented a teleconferencing application that facilitates a natural conferencing environment through voice-activated audio, "video follows voice" switching and a simple user interface. We have discussed the components of the application, and how it was built using the Programmers' Playground.

At the moment, the only option for allocation of source ids is random. As a result, floor control assumes that all participants of the conference are peers, since there is no scheme for priority allocation. We would like to use the system for other settings such as a remote classroom setup or a conference with a chairman and a particular ordering of the speakers. For these applications, we would need more control over the allocation of the source ids. Since the conference module controls the assigning of source ids, this would require a small change in the module to accommodate a new method for allocating ids.

Another problem is the voice activated threshold, since this parameter must be set by the user and depends on the individual office setup and voice. Through experience we have discovered a good default setting that works for most setups. In addition, we plan to incorporate an automated mechanism for setting the threshold to a level just above the measured background noise.

Audio and video conferences are useful, but ideally one would like to have the use of other collaborative tools in a conference. We intend to introduce exhibits to the conferences that allow the use of a set of

commercially available tools as well as collaborative tools built using the Programmers' Playground. Each exhibit would represent a different tool. Conferences would be allowed to have an arbitrary number of exhibits associated with them, and exhibits could be added and deleted from the conference dynamically. Individuals would only receive those exhibits their workstation could support.

We would also like to incorporate some of the features suggested by participants of the faculty meeting described in Section 4. We plan to implement side conversations by creating a point-to-point connection between the two people and turning off their audio transmitters to the conference, but still allowing the primary conference audio to be mixed in with the side conversation. We also plan to indicate through the user interface when a person is being viewed.

Presently, the connection management system is centralized, so this limits the scalability of the system. This centralization occurs in both the Playground and the ATM connection management systems. In the next implementation of Playground, the connection management system will be distributed so this will solve the problem in Playground. The ATM connection management system is centralized because we are using dynamic PVC's. By switching to SVC's, we should also be able to distribute the ATM system.

ACKNOWLEDGMENTS

We thank Bala Swaminathan and Ram Sethuraman for their work in developing the Playground library. This research was supported in part by the National Science Foundation under grants CCR-91-10029, CCR-94-12711, National Challenges Award grant NCR-9318178 and Research Infrastructure grant CDA-9401159. It was also supported by the Advanced Research Projects Agency under contract DABT-63-95-C-0083, the National Institute of Health under contract 263-94-C-0617 and Southwestern Bell TRI.

REFERENCES

- [1] J. R. Cox, Jr., M. Gaddis, and J. S. Turner. "Project Zeus: Design of a Broadband Network and its Application on a University Campus," *IEEE Network*, pages 20-30, March 1993.
- [2] K. J. Goldman, et. al. "Welcome to the Programmers' Playground!" <http://www.cs.wustl.edu/cs/playground/>.
- [3] K. J. Goldman, T. P. McCartney, R. Sethuraman, and B. Swaminathan. "The Programmers' Playground: A Demonstration," In *Proceedings of the 1995 ACM International Conference on Multimedia*, November 1995, pp. 317-318 See also the conference CD-ROM proceedings for a longer version.
- [4] K. J. Goldman, T. P. McCartney, R. Sethuraman, B. Swaminathan, and T. Rodgers. "Building Interactive Distributed Applications in C++ with The Programmers' Playground," Washington University Department of Computer Science technical report WUCS-95-20, July 1995.
- [5] K. J. Goldman, B. Swaminathan, T. P. McCartney, M. D. Anderson, and R. Sethuraman. "The Programmers' Playground: I/O Abstraction for User-Configurable Distributed Applications," *IEEE Transactions on Software Engineering*. Vol. 21, No. 9, pp. 735 -746, September 1995.
- [6] T. P. McCartney, K. J. Goldman, and D. E. Saff. "EUPHORIA: End-User Construction of Direct Manipulation User Interfaces for Distributed Applications," *Software Concepts and Tools*, Vol. 16, No. 4, pp. 147-159, December 1995.
- [7] W. D. Richard, J. R. Cox, Jr., A. M. Engebretson, J. Fritts, B. Gottlieb and C. Horn. "Production Quality Video Over Broadband Networks: A Description of the System and Two Interactive Applications," *IEEE Journal on Selected Areas in Communications*, Vol. 13, No. 5, pp. 806-815, June 1995.
- [8] S. McCanne and V. Jacobson. "vic: A Flexible Framework for Packet Video," In *Proceedings of the 1995 ACM International Conference on Multimedia*, November 1995, pp.511-522

- [9] R. Frederick. "Network Video (nv)," Xerox Palo Alto Research Center. Software on-line. <ftp://ftp.parc.xerox.com/net-research>
- [10] T. Turlitti. "The INRIA videoconferencing system(IVS)," *ConneXions*, Vol. 8, No. 10, pp. 20-24, October 1994.
- [11] S. Casner and S. Deering. "First IETF Internet audio cast," *ConneXions* Vol. 6, No. 6, pp. 10-17, 1992.
- [12] D. Saha, D. Kandlur, T. Barzilai, X. Shae, and M. Willebeek-LeMair. "A Video Conferencing Testbed on ATM: Design, Implementation and Optimizations," In *Proceedings of the International Conference on Multimedia Computing and Systems*. May 1995, pp. 23-30.
- [13] Y-H. Puztaszeri, M. Aloou, E. W. Biersack, P. Dubois, J-P. Gaspoz, P. Gros , J-P. Hubaux. "Multimedia Tele-tutoring over a Trans-European ATM Network," In *Proceedings of Multimedia: Advanced Teleservices and High-Speed Communication Architectures. Second International Workshop, IWACA '94*. September 1994, pp. 315-327.