

Multimedia Collaboration

Sudhir R. Ahuja

H. Fred Halsch

Ram S. Ramamurthy

Lucinda M. Sanders

Today's business activities require collaboration and teamwork among groups of people, some of whom are geographically dispersed. This paper describes a functional architecture that supports remote collaborations. AT&T has introduced leading-edge products and services in this field, and this paper presents how they are evolving to meet the need for a common strategy and architecture.

Introduction

Teamwork and collaboration are the underpinnings of most business projects. Today, business activity requires some form of collaboration among people, from designing cars to arranging weddings. As people and companies have become more geographically dispersed, the cost of collaboration has increased significantly. Not only is the cost of travel a significant factor, but also the loss of productive time associated with it. These factors have brought attention to and created the opportunity for remote collaborations.

More people than ever are interested in collaborating from a distance using electronic means. It allows them to stay in their local environment and minimizes disruptions. Telephony and fax already provide people with mechanisms to communicate in real time at a distance. However, they do not incorporate one of the most important technologies of today's businesses—desktop computing. Increasingly, businesses rely on computers for information handling and flow. The use of computers for sharing and exchange of information in remote collaborations has engendered many products that allow screen or application sharing between computers. Multimedia desktop conferencing systems allow real-time sharing of data, voice, and video. However, these address only part of the need for multimedia collaborations.

Collaborations are inherently long-term interactions, implying the need for support of various kinds over long periods of time. Generally, a collaboration involves real-

time and asynchronous interactions, such as electronic mail and fax. It includes both shared and individual work. To be effective, remote collaborations require support in the areas of scheduling, reminders, call coverage, shared information repositories, meeting notes, and the ability to suspend and resume interaction. Embedded within all these factors is the need for communication to access people, programs, and data.

AT&T has the technology and the network to provide support for global multimedia collaborations. To establish its presence, AT&T has already launched leading products and services such as the Vistium™ system and WorldWorx™ Network Services. AT&T Global Business Communications Services (GBCS) and Global Information Systems (GIS) are developing products and testing concepts that will strengthen AT&T's presence at customer premises and help support end-to-end multimedia communication and collaboration. (See Panel 1 for definitions of abbreviations, acronyms, and terms.)

This paper discusses a functional architecture that addresses various issues supporting remote collaborations. AT&T has introduced leading-edge products and services in this field, and this paper also discusses how they are evolving to meet the need for a common strategy and architecture.

Networking

Collaborations are characterized by a common goal and length of time over

which a given set of participants collaborate to achieve that goal. They involve many two-party or group interactions. In local collaborations, these interactions are conducted in face-to-face meetings, without the need for a communication infrastructure. In remote collaborations, however, the first big issue is support for "ubiquitous" multimedia communication. "Can I really find all the people in my project, wherever they may be, and connect to them?" To a large extent, this is what telephony has accomplished for voice communication. For information access, this is partially accomplished by the Internet.

Multimedia collaboration requires support for seamless interworking among various networks users need to access different types of media. For instance, collaboration should be possible between groups of people in two different locations. Such integration requires a robust architecture and protocol specifications so that applications can be built to satisfy the needs of customers. Figure 1 shows a conceptual view of domains of interest in long-term collaboration.

Users of long-term collaboration should be able to communicate from wherever they may be and with whatever capabilities they may have. Multiple users need the ability to communicate and/or collaborate among themselves at different times and from different places. Figure 1 illustrates three important functional criteria:

- **Platform independence.** Participants in a long-term collaborative meeting should be able to use different desktop platforms, such as PCs, workstations, televisions, and cellular phones.
- **Access independence.** Participants should be able to logically connect through a variety of network access technologies, such as narrowband integrated services digital network (n-ISDN), asynchronous transfer mode (ATM), "plain old telephone service" (POTS), wireless, and local area network (LAN).
- **Media independence.** Participants should be able to use the various capabilities they have access to, such as audio (voice), video, and computer information. Further complications arise if transcoding or live speech translation needs to take place to achieve friendly and useful collaboration models.

Electronic Places

AT&T Bell Laboratories researchers¹ have developed a model that supports media, platform, and

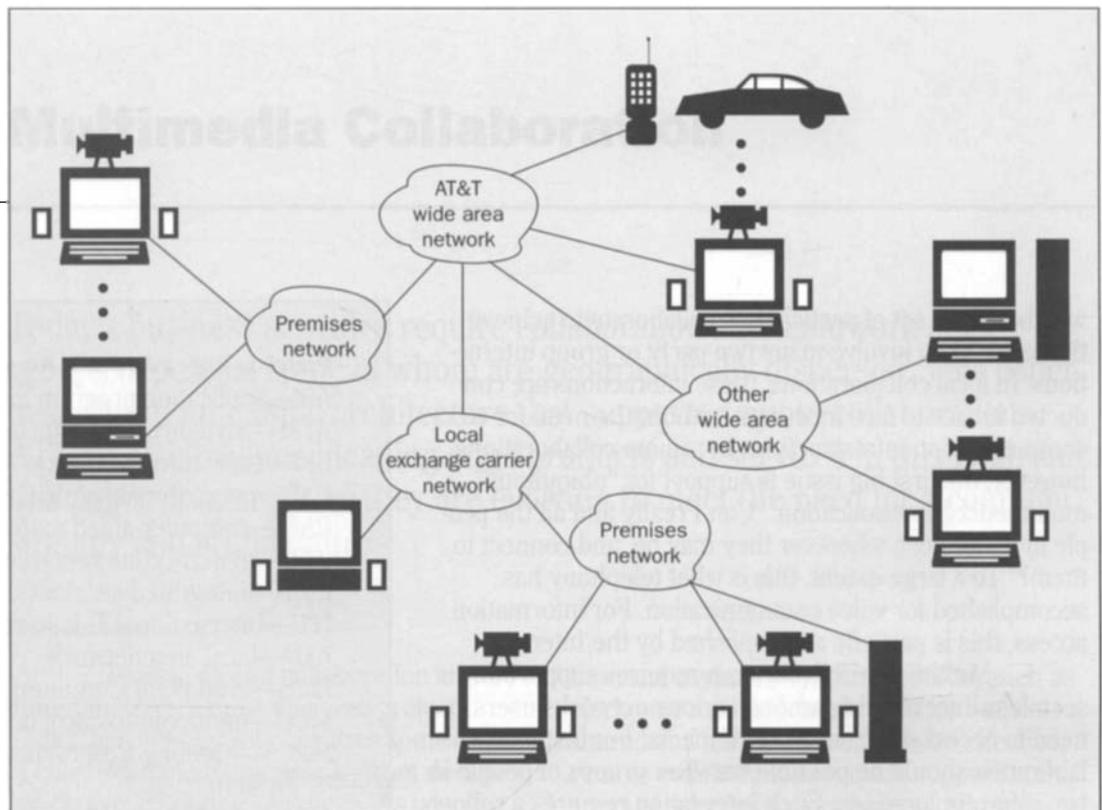
Panel 1. Abbreviations, Acronyms, and Terms

API—application program interface
ATM—asynchronous transfer mode
BRI—basic rate interface
CAD—computer-aided design
CAM—computer-aided manufacture
GCC—general conference control
ISDN—integrated services digital network
ITU—International Telecommunications Union
LAN—local area network
MCS—Multipoint Communication Service
MCU—multipoint control unit
MRCS—Meeting Reservation and Control System
PBX—private branch exchange
POTS—"plain old telephone service"
SDS—Switched Digital Service
SDDN—Software-Defined Digital Network
VMR—virtual meeting room
WAN—wide area network

transport independence. It was implemented in a multimedia collaboration system called Rapport, which uses the abstraction of an electronic place, called a *virtual meeting room* (VMR), to hold the remote collaboration. Much like physical space, an electronic place inherently accommodates many people and various types of media. A place is what gives *persistence* to meetings or information. Participants can bring information to the VMR and leave it there for any length time of time. Different people can join or leave the VMR whenever they please. Participants may use any type of media to join and take part in the meetings held in a VMR, or they may choose at times to work there alone. Like participants, programs can join the VMR, and participants can share them. A simple example of a VMR is a "session" that includes real-time voice, video, and data connection, as is the case in most desktop conferencing systems.

"Sessions" are held by most desktop conferencing systems to manage the association of connections in various media. Generally, the sessions deal with trans-

Figure 1. Conceptual view of domains of interest in long-term collaboration. Users of long-term collaboration should be able to communicate from wherever they may be and with whatever capabilities they may have. Multiple users need the ability to communicate and/or collaborate among themselves at different times and from different places.



port-level abstraction to facilitate easy handling of multiple connections. The VMR extends this concept to provide persistence and user-level signaling support, enabling it to exist even without connections. Such a persistent VMR allows a user (for example, a person on a mobile phone) to drop out and rejoin a VMR any number of times without forcing anyone else to drop out or reestablish connections.

Electronic places, much like physical places, are useful only because of their ambiance and the facilities that they provide, such as chairs/tokens, tables/whiteboards, and lights/video. Mechanisms have been developed that allow new services to be easily introduced within a VMR. Bringing a service into a VMR allows it be used and shared by all the participants.

In 1992 this model was modified further, becoming a network-based client-server model, and the corresponding communications broker architecture was developed by Bell Laboratories.²

The concept of an “electronic place” is supported by a VMR server that creates and manages all VMRs. A VMR may be as transient as a phone call, or it may persist for years in support of collaborations. As Figure 2 shows, each user has access to VMRs through a local client. Each service, be it transport of a particular medium or a shared program, is added as a new participant to the VMR and is embodied by a server (for example, a voice switch or a video bridge) and the service-specific clients that allow access to that service in the context of a VMR.

This layered approach supports collaboration. All physical transport is modeled as “virtual transport” and is accessed using a common application program interface (API). This allows applications to be written once for the virtual transport interface, rather than multiple times for different physical transports.

A model for long-term collaboration imposes several important architectural requirements to satisfy platform, access, and media independence:

- *Distributed intelligence* in the network that can separate a call from the connection. For example, a participant in a long-term collaboration service should be able to leave his or her office and rejoin a meeting from a car phone or a different location with different media connection capabilities.
- *Sophisticated end-to-end signaling* that can establish and coordinate many types of multimedia services. For example, with the touch of one button (a preprogrammed voice phone number, or a picture of the interested party from an electronic phone directory), the user should be able to establish a full multimedia call (voice, video, and computer connections).
- *Intelligent endpoints*, or computers, whose networks need to be able to negotiate dynamically with an endpoint intelligent agent for user’s capabilities and desired services that require support of continuous negotiation. For example, a user may wish to establish a “voice-only” call through a screen interface and later add video capability to the conversation.

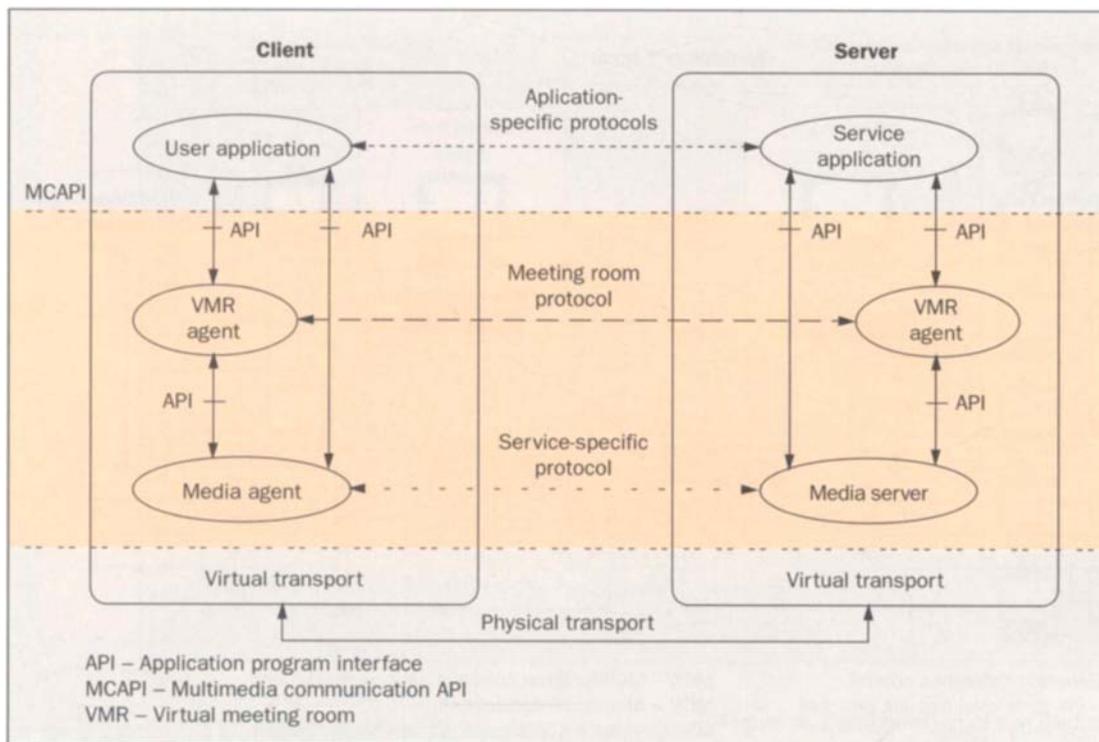


Figure 2. High-level software architecture for electronic places. Each user has access to VMRs through a local client. Each service is added as a new participant to the VMR and is embodied by a server (for example, a voice switch or a video bridge) and the service-specific clients that allow access to that service in the context of a VMR. This layered approach supports collaboration.

- *Intelligence in the network* to provide a variety of interworking and interoperable services and gateway access to other services. For example, the network should allow access to a multimedia database during a collaboration session, or it should be able to share an application between dissimilar endpoints.

These requirements are met by the communications broker architecture model, whose core functions are being trialed by GBCS in multimedia communication systems. Efforts are also under way—in various standards bodies such as the Multimedia Communications Forum—to propose and agree on this type of model. A prototype of this model was designed by Bell Laboratories and Business Communication Services (BCS) as the Virtual Meeting Service concept, which led to the introduction of WorldWorx Personal Conferencing Service. The next section discusses some of these efforts within AT&T.

WorldWorx Personal Conferencing Service

The WorldWorx Personal Conferencing Service is a network-based audio, video, and data communications service available for both two-party and multiparty calls in the United States and selected countries in Europe,

Asia, and the Caribbean. The WorldWorx Personal Conferencing Service also links people using personal workstations with people using conference-room video systems that operate at 112 or 128 kb/s. The service provides a broad set of features, including:

- Multipoint collaboration,
- Audio add-on,
- Chairperson control or voice-energy-based video switching,
- Data conferencing and PC collaboration, and
- Security.

The WorldWorx Personal Conferencing Service is available to users with endpoints (workstations or PCs with full-duplex audio/video add-on kits) that meet the International Telecommunications Union-Telecommunication Standardization Sector (ITU-T) H.320 video/audio standards and the T.120 data standards. BCS is working with a number of multimedia endpoint product vendors and standards committees to consolidate the endpoint industry and thus maximize the number of endpoints that can operate together. These endpoint vendors include AT&T GIS, Intel, Apple, Sun Microsystems, IBM,* and others. Much progress has been made in audio and video interoperability. The T.120 data standard has been accepted for multiparty data communications, and negotiations are under way with vendors to define a common data application based on ITU-T T.126.

Currently, endpoints that comply with the H.320 video/audio standard can conference together using audio and video services. Endpoints from the same vendor typically can also collaborate using a shared white-

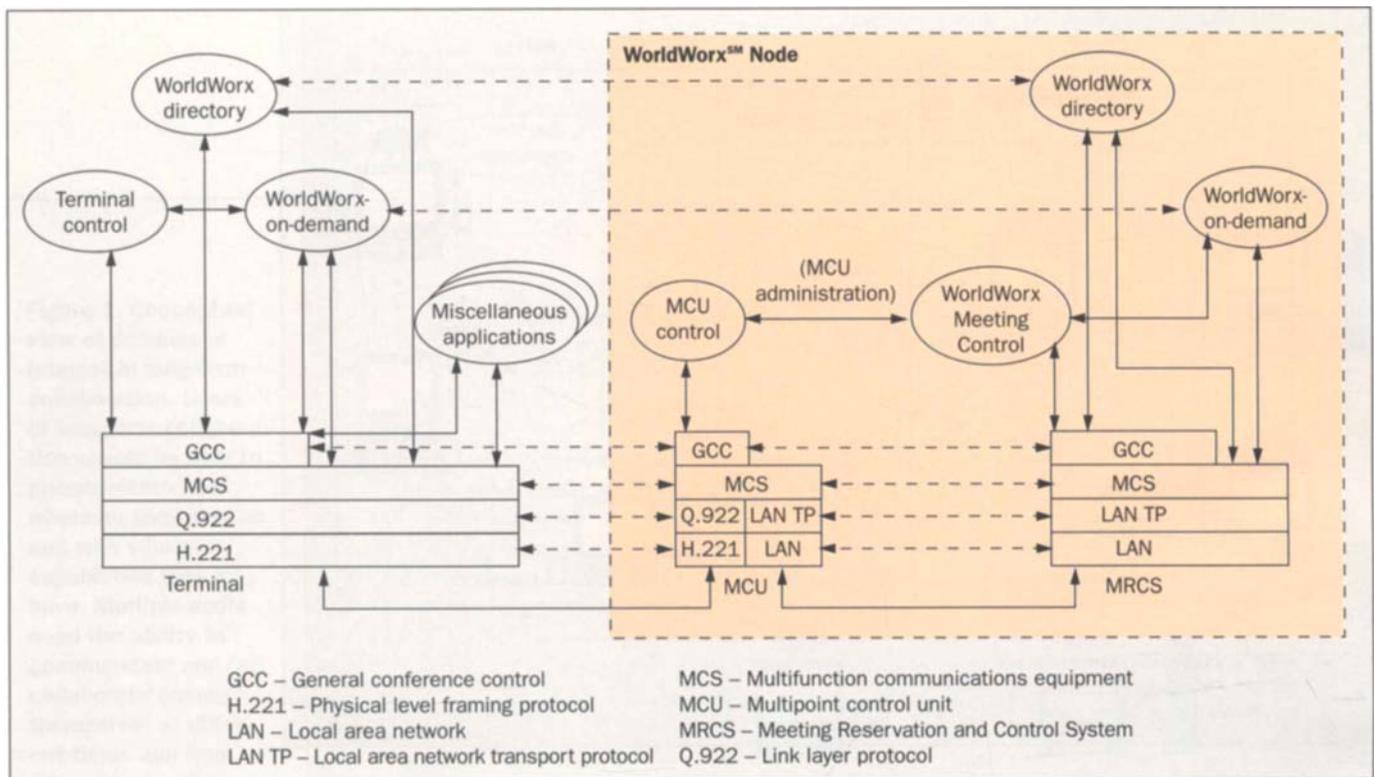


Figure 3. A protocol picture of the WorldWorxSM advanced multipoint configuration with its layered architecture.

board and application sharing.

Figure 3 shows a protocol picture of the multipoint configuration with its layered architecture. Although Advanced WorldWorx is initially implemented using the H.320 series of protocols for ISDN transport, all service definition, such as directory and meeting control, is transport independent. The layers below the Multipoint Communication Service (MCS) shown in Figure 3 represent the virtual transport of the generic architecture. The MCS layer provides multipoint protocol support for data conferencing. General conference control (GCC) supplies the middleware for session management. Most of the services are implemented as a client/server architecture. For example, Directory, Reservation, and WorldWorx-On-Demand are implemented in the network on the Directory Server or the Meeting Reservation and Control System (MRCS) and are accessed through service-specific clients.

On-demand multipoint conferencing, in which an MCU can be brought into the conference as needed, is much like the function provided by a conference button on ISDN phones in private branch exchange (PBX) or CENTREX services.

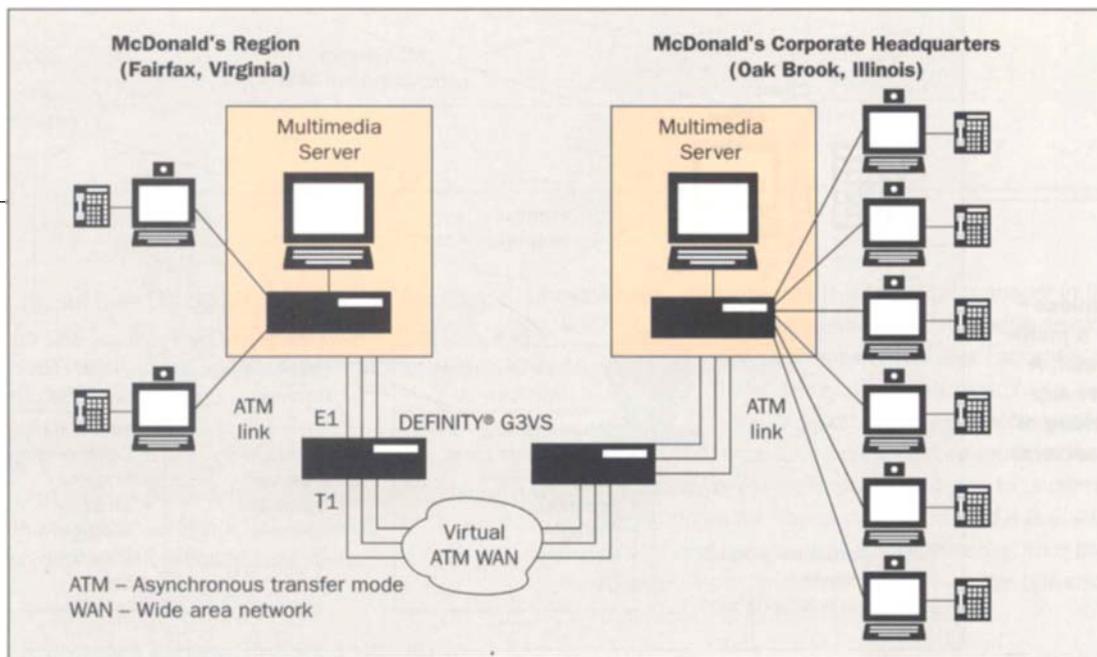
WorldWorx Personal Conferencing Service initially allows ISDN endpoints to hold desktop multimedia conferences globally. It also provides Switched Digital

Service (SDS) and Software-Defined Digital Network (SDDN), services that have the necessary physical connectivity for such conferencing.

Multimedia Calling Concept Trial

During the past two years, GBCS and Bell Laboratories have been conducting a multimedia concept trial to explore a new communications model to support a variety of multimedia calls. This AT&T initiative trialed the broad-based usefulness of this model for real-time multimedia communications by providing multimedia calling services, interworking services, and wide area network (WAN) enterprise networking to McDonald's Corporation. The trial, begun in July 1994 and ended in May 1995, augmented basic desktop collaboration tools (voice, video, and application sharing) with such calling services as call forwarding, conferencing, and coverage. It also provided an interworking service to the voice telephony network that enabled McDonald's to use multimedia endpoints to communicate with any voice set, anywhere in the world. These types of services, in addition to bandwidth management and WAN management, are key to making multimedia calling an indispensable business tool.

The McDonald's trial included two prototype ATM LAN hubs. The first hub supported six users in the headquarters building (Oakbrook, Illinois), and the second hub supported two users in a regional office in Fairfax, Virginia. The two hubs were connected using a dedicated T1 trunk between the two locations. ATM cells were sent over the T1 trunk to create a "virtual" ATM WAN.



In each location, a DEFINITY® telecommunications system G3VS provided calling features and WAN trunking access. The call server and endpoints were all running on SUN* workstations. To improve the audio quality, basic rate interface (BRI) phones were attached to the SUN endpoints. The application focused exclusively on sharing a sophisticated computer-aided design/computer-aided manufacture (CAD/CAM) package, but other applications could be shared as well. Figure 4 shows the configuration.

Using ATM connectivity to the desktop on SUN SPARCstations,* the trial provided several types of services:

- *Video.* Each multimedia endpoint received Joint Photographic Experts Group (JPEG) video.
- *Audio.* Each multimedia endpoint received high-quality audio.
- *Voice interworking.* Each multimedia endpoint could make voice calls to any phone on the public voice network.
- *Shared application.* Multimedia endpoints could share an application between endpoints (both point-to-point and multiparty). McDonald's shared a sophisticated CAD/CAM package across the wide area network.
- *Shared whiteboard.* A shared whiteboard capability was provided for creating and annotating images.
- *Multimedia conferencing.* Each multimedia endpoint could make not only point-to-point multimedia calls, but also multipoint calls. As many as six-party conference calls with four services were conducted—voice, video, shared applications, and shared whiteboard.
- *Wide area network.* All basic calling capabilities were provided across the wide area network.
- *Calling features.* The calling features of multimedia endpoints were enabled by DEFINITY PBX logic.

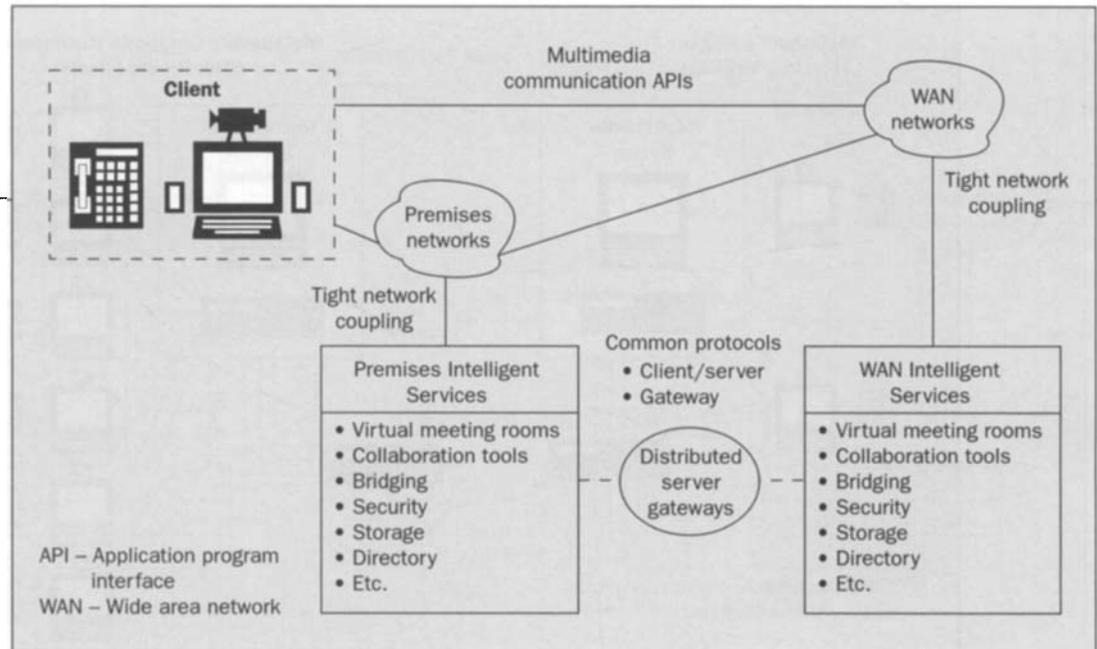
Figure 4. Configuration of the McDonald's trial, which included two prototype ATM LAN hubs.

Panel 2 describes a typical multimedia calling scenario using these trial capabilities. This trial provided AT&T with useful customer feedback about basic collaboration tools, such as shared applications, and multimedia calling services, such as coverage and conferencing. In addition, the trial gave everyone involved experience with key technology concepts related to communication models, client/server architectures, multimedia communication, and voice interworking.

Seamless Integration

In light of the variety of ongoing efforts within AT&T, it is important to seamlessly integrate products and services to offer an *end-to-end, customer-focused solution*. As mentioned earlier, remote collaboration requires integrated support for a variety of interactions. A key service identified in the WorldWorx Personal Conferencing Service and multimedia concept trials is support for the true *multimedia phone call*. This would enable users to place multimedia calls anywhere, anytime, and to anyone without worrying about connectivity, compatibility, or coverage. Such calls would have access to, and be the initiators of, collaborations. Figure 5 shows a vision of such a configuration. For a multimedia call to be considered seamless, it has to appear the same to the user, regardless of whether the call is being handled by a local server, a wide area server, or both. This can be accomplished by providing the same multimedia communication API to local and wide area call servers. Figure 5 also shows a gateway server, which is needed to support interworking of different call servers.

Figure 5. Seamless integration of a multimedia phone call. A gateway server supports interworking of different call servers.



To achieve such a vision, three important criteria need to be satisfied:

- *A common applications environment using multimedia communications APIs.* The multimedia communications API is a set of APIs that provide a high-level interface for networking clients and servers. It presents applications with a choice of standard communication objects, such as session, connection, and port, as well as persistent communication objects, such as VMRs, shared documents, and shared messages. Winsock 2.0 is included as APIs for standard transport support.
- *Two common protocols for seamless integration.* These are (1) a set of identified protocols between various client and server components. One such set of protocols, the ITU-T H.320 and T.120 suites, already deployed in WorldWorx Personal Conferencing Service, is being extended to include new media, transports, and services; (2) a set of *gateway* protocols between different network domains, necessary to mediate policies for ownership, billing, and other feature management between local servers and wide area servers.
- *Persistent network servers that provide key services.* The key services are Directory, Session Management, Security, and Bridging and Repository for shared multimedia information.

Bell Laboratories, BCS, and GBCS are working actively in these three areas.

Conclusions

Multimedia collaboration over distances requires support for persistent connections and multimedia communication. AT&T is making significant progress towards providing this capability over any distance,

Panel 2. A Typical Multimedia Calling Scenario.

John and Paula need to work on an engineering design for a new McDonald's store. Paula, who works in Fairfax, Virginia, calls John on her multimedia endpoint. John, who works in Oakbrook, Illinois, is not at his desk; he is in Chuck's office, in another part of the building. But, because John was expecting a multimedia call from Paula, he has forwarded his multimedia endpoint to Chuck's voice set. Chuck answers Paula's call, and John agrees to call her back when he has finished talking to Chuck.

When John returns to his desk, he calls Paula back on his multimedia endpoint. He starts up the call with both audio and video; when Paula answers, he can both see and hear her. Paula launches the application-sharing media, bringing up the CAD/CAM drawing on which she wants John's input. After some discussion about the drawing, John suggests various changes, and, with Paula's concurrence, he adds these to the drawing in real time. At one point, neither John nor Paula knows the exact dimensions of a drawing component. John conferences in Bob, who has a voice-only set. Both John and Paula chat with Bob as he is answering their questions. They then drop Bob from the call and continue with their collaboration session.

either locally or globally. To support multimedia collaboration effectively, AT&T must also support user and application interfaces through the use of multimedia communication APIs, end-to-end multimedia connections (common protocols over local area and wide area), and key network-

based services, such as Directory and Session Management. With its expertise and presence in these areas, and with a common vision, AT&T can provide the best network for multimedia collaboration.

***Trademarks**

IBM is a registered trademark of International Business Machines, Inc. SPARCstation is a trademark of SPARC International. Sun is a registered trademark of Sun Microsystems.

References

1. S. Ahuja, J. R. Ensor and D. Horn, "Rapport: A Multimedia Conferencing System," IEEE 2nd Annual Conference on Office Information Systems, San Jose, California, March 1988.
2. S. Ahuja, J. R. Ensor, A. K. Kuthyar, and R. Ramamurthy, "A Communications Broker Architecture for AT&T Multimedia Communications Solutions," *Proceedings of Multimedia and Collaboration System TIES*, Orlando, Florida, December 8-10, 1993.

(Manuscript approved August 1995)

Sudhir R. Ahuja is head of the Multimedia Communication



Research Department at AT&T Bell Laboratories in Holmdel, New Jersey. He is responsible for research in multimedia collaboration, communication, applications, and platforms, as well as research and technology transfer of key services

such as multimedia conferencing and multipoint audio and video bridging. Mr. Ahuja joined AT&T in 1977, after receiving a B.Tech. in electrical engineering from the Indian Institute of Technology, Bombay, and an M.S. and Ph.D. in computer science and electrical engineering from Rice University, Houston, Texas.

H. Fred Haisch is technical manager of Advanced Systems



Engineering for the WorldWorxSM Personal Conferencing Service at AT&T Bell Laboratories in Holmdel, New Jersey. Since joining AT&T in 1964, he has held positions in hardware and software development and systems engineering. Mr. Haisch

received a B.S. from the University of Tennessee, Knoxville, an M.S. from Harvard University, Cambridge, Massachusetts, and a Ph.D. from the Polytechnic Institute of Brooklyn, New York, all in electrical engineering.

Ram S. Ramamurthy is a technical manager in the Multimedia Services and Data Networking Architecture



Department at AT&T Bell Laboratories in Holmdel, New Jersey. He is responsible for business multimedia applications architecture, managing a group that develops advanced multimedia and broad-

band service concepts and prototypes to understand service implications. Mr. Ramamurthy received a B.E. (Honors) in electronics and communications engineering from the University of Madras, India, and an M.S.E.E. from the University of Missouri, Columbia. He joined AT&T in 1987.

Lucinda M. Sanders is technical supervisor of the Multimedia



Project Management and Deployment Group of AT&T Global Business Communications Services in Denver, Colorado. She is working on multimedia communication systems and is responsible for development project management. Ms. Sanders

received a B.S. in computer science and economics from Louisiana State University, Baton Rouge, and an M.S. in computer science from the University of Colorado, Boulder. She joined AT&T in 1977.