Interactive Multimedia Services for Consumers and Businesses

Edward S. Szurkowski
Jack C. Warner

The emergence of broadband networks for residential and business customers enables a wide range of new interactive services. For consumers, these new services might include movies on demand, networked multiplayer video games, and interactive shopping via TV. Business applications might include broadcast-quality video teleconferencing, distance learning, and telecommuting. This paper presents the technical issues and infrastructure involved in delivering several of these services, along with some of the lessons learned and insights gained from AT&T's early interactive broadband services trials.

Introduction

The principal objective of AT&T's multimedia research and development activities is to create the technology required to economically deliver a variety of compelling interactive services to consumer and business customers on a large scale.

Although the specific services offered to consumer and business customers may appear quite different, they can frequently use a common set of technology elements that have been configured or applied in different ways. There will be differences in the customer premises equipment and local access facilities, but the supporting network infrastructure—that is, the transport, switching, storage, and processing elements—will be substantially the same. This sharing can lead to economies of scale that benefit both classes of end users.

In many cases, the core technologies required for large-scale delivery of interactive video services to consumers and businesses are now ready for initial deployment. The range of possible commercial services is vast, but trials and other studies have identified some key opportunities where the value delivered to the customer may be large enough to justify creating the necessary infrastructure.

This paper describes some likely interactive video-oriented multimedia services for consumers and businesses and discusses the key technical challenges involved in their delivery. It also describes some early trials of interactive multimedia services and the insights that were gained.

Interactive Consumer Video Services

Most proposed interactive consumer video services fall into the general areas of entertainment, commerce, information, communication, or education, although some services may combine elements of several of these areas. For example, a home shopping service may allow a subscriber to automatically set up a video telephone call to a customer service representative who can answer questions about the merchandise.

Some important technology considerations for interactive consumer video services are unique compared to business environments. These are described in the subsections that follow, along with the expected overall architecture of the consumer delivery infrastructure. Within that context, two consumer services that are expected to have broad consumer appeal are also described, along with their key underlying technical challenges.

Residential Broadband Networks. Figure 1 shows a general high-level view of the anticipated platform infrastructure architecture for large-scale delivery of interactive consumer video services. A global broad-

Panel 1. Abbreviations, Acronyms, and Terms

ADSL—asymmetric digital subscriber line
ATM—asynchronous transfer mode
BPSK—binary phase shift keying
BRI—basic rate interface
CATV—cable television
CO—central office
COMPASS—Communications Programs for Advanced Switched Services
CPE—customer premises equipment
DS-3—digital signal level 3
DVHT—digital video home terminal
FCC—Federal Communications Commission
FSK—frequency shift keying
FTTC—fiber-to-the-curb
GTA—general terminal adapter
HFC—hybrid fiber-coax
HSC—Home Shopping Channel
ISDN—integrated services digital network
ITV—interactive television
LATA—local access and transport area
LEC—local exchange carrier
MLS—Multiple Listing Service
MMDS—multichannel microwave distribution system
MPEG—Motion Picture Experts Group
n-ISDN—narrowband integrated services digital network
node—an individual fiber trunk and coax distribution plant
NTSC—National Television Systems Committee
OC—optical carrier
ONU—optical network unit
PCMCIA—Personal Computer Memory Card Interface Association
PDA—personal digital assistant
PDC—personal digital communicator
PON—passive optical network
POTS—"plain old telephone service"
QAM—quadrature amplitude modulation
QPSK—quadrature phase shift keying
RF—radio frequency
RISC—reduced instruction set computer
SDV—switched digital video
SMDS—Switched Multi-megabit Data Service, a service mark of Bell Communications Research (Bellcore)
SONET—synchronous optical network
STB—set-top box
STM—synchronous transfer mode
VA—Veterans Administration
VDT—video dial-tone

band asynchronous transfer mode (ATM)[1] network interconnects distributed server complexes. (See Panel 1 for definitions of abbreviations, acronyms, and terms.) These server complexes provide various services to consumer residences over broadband local access networks. Each element of this infrastructure is described in the sections that follow.

Broadband backbone network. The global broadband ATM backbone network shown in Figure 1 is not monolithic. As Figure 2 illustrates, it is composed of a hierarchy of interconnected ATM subnetworks covering various geographic areas. These ATM networks can be owned and operated by regional access network providers and/or long-haul carriers. Over time, competition will emerge among broadband local access networks, backbone networks, and server complexes. Any platform architecture must be able to cleanly accommodate competing infrastructures, as well as competing service providers.

Broadband backbone networks used for consumer video services must also be interconnected to narrowband networks such as "plain old telephone service" (POTS) dial-up and wireless. This will enable traditional narrowband services, such as electronic mail, to be used with broadband and narrowband terminal devices. For example, a person with a wireless personal digital assistant (PDA) or personal digital communicator (PDC) could send a text message to another person, who could read it using an interactive television application. This synergy, or interoperability, of narrowband and broadband services will become increasingly important.

In the platform backbone network, all transport and switching is likely to be based on ATM with underly-

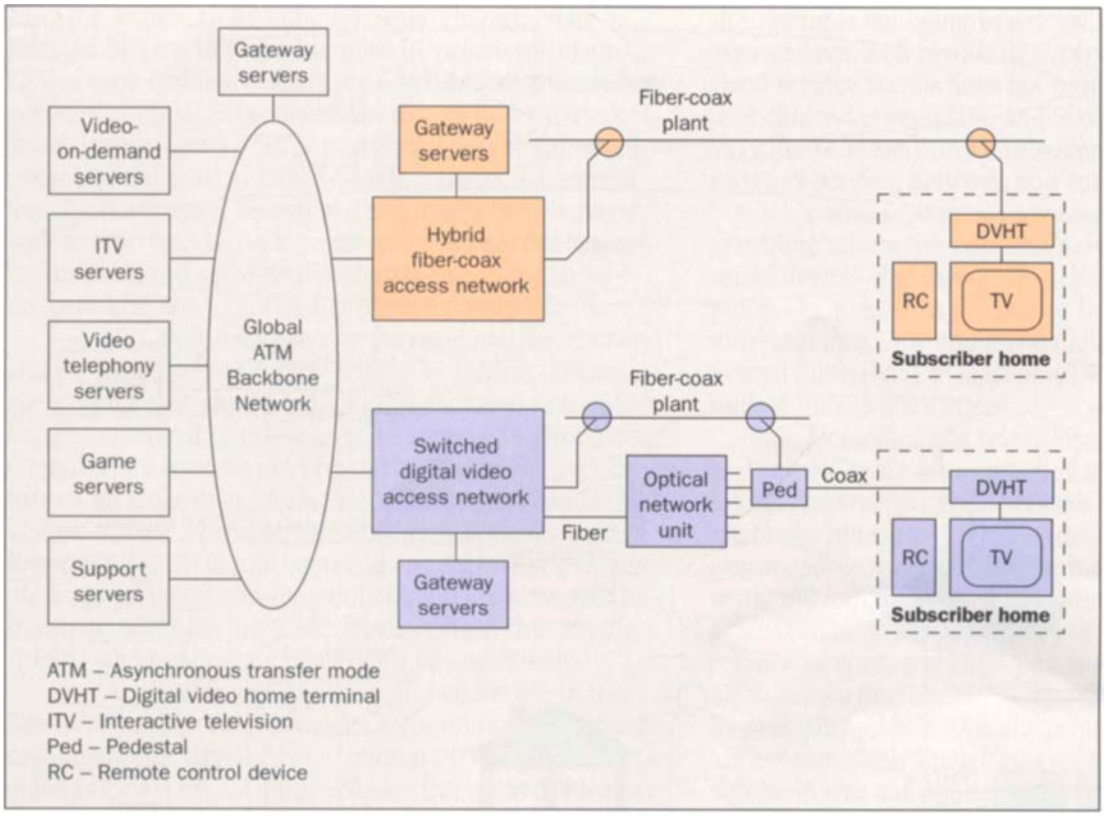

Figure 1. Integrated consumer multimedia services platform for a large-scale delivery of interactive consumer video services.

ing synchronous optical network (SONET)[2] transmission facilities. Using a common transmission and switching infrastructure, ATM and SONET enable the platform to support a wide variety of services with different communication needs.

Broadband local access networks. The platform shown in Figure 1 can support a number of different broadband local access network architectures. Figure 1 shows two of these: hybrid fiber-coax (HFC), sometimes known as fiber-to-the-feeder; and switched digital video (SDV), also known as fiber-to-the-curb (FTTC). Other local access architectures, such as passive optical networks (PONs) or asymmetric digital subscriber line (ADSL),[3] can also be accommodated, although most large-scale deployments in the United States currently focus on HFC or SDV (sometimes known as passband and baseband access networks, respectively).[4]

HFC and SDV access network architectures provide bidirectional digital channels to customer premises equipment (CPE) in subscribers' homes. Although the details of the transport and switching for each type of network differ, the higher-level elements in the platform can function with minimal specific knowledge of the underlying access scheme.

An HFC network, shown in the upper half of Figure 1, is a wide-bandwidth, bidirectional, analog radio-frequency (RF) transmission system. Signals (such as broadcast TV and digital data) are modulated onto RF carriers at different frequencies throughout the spectrum supported by the system. While the sample structure described here is based on systems carrying National Television Systems Committee (NTSC) video, the principles are the same for other formats, which may have different frequencies and bandwidths.

A typical HFC network will transport RF signals from 5 to 750 MHz, or higher. The spectrum between about 50 and 750 MHz, the upper limit of the system, is usually reserved for carrying *downstream* signals from the head-end or central office (CO) to residences. The spectrum from 5 to 30–40 MHz will typically be reserved for *upstream* sig-

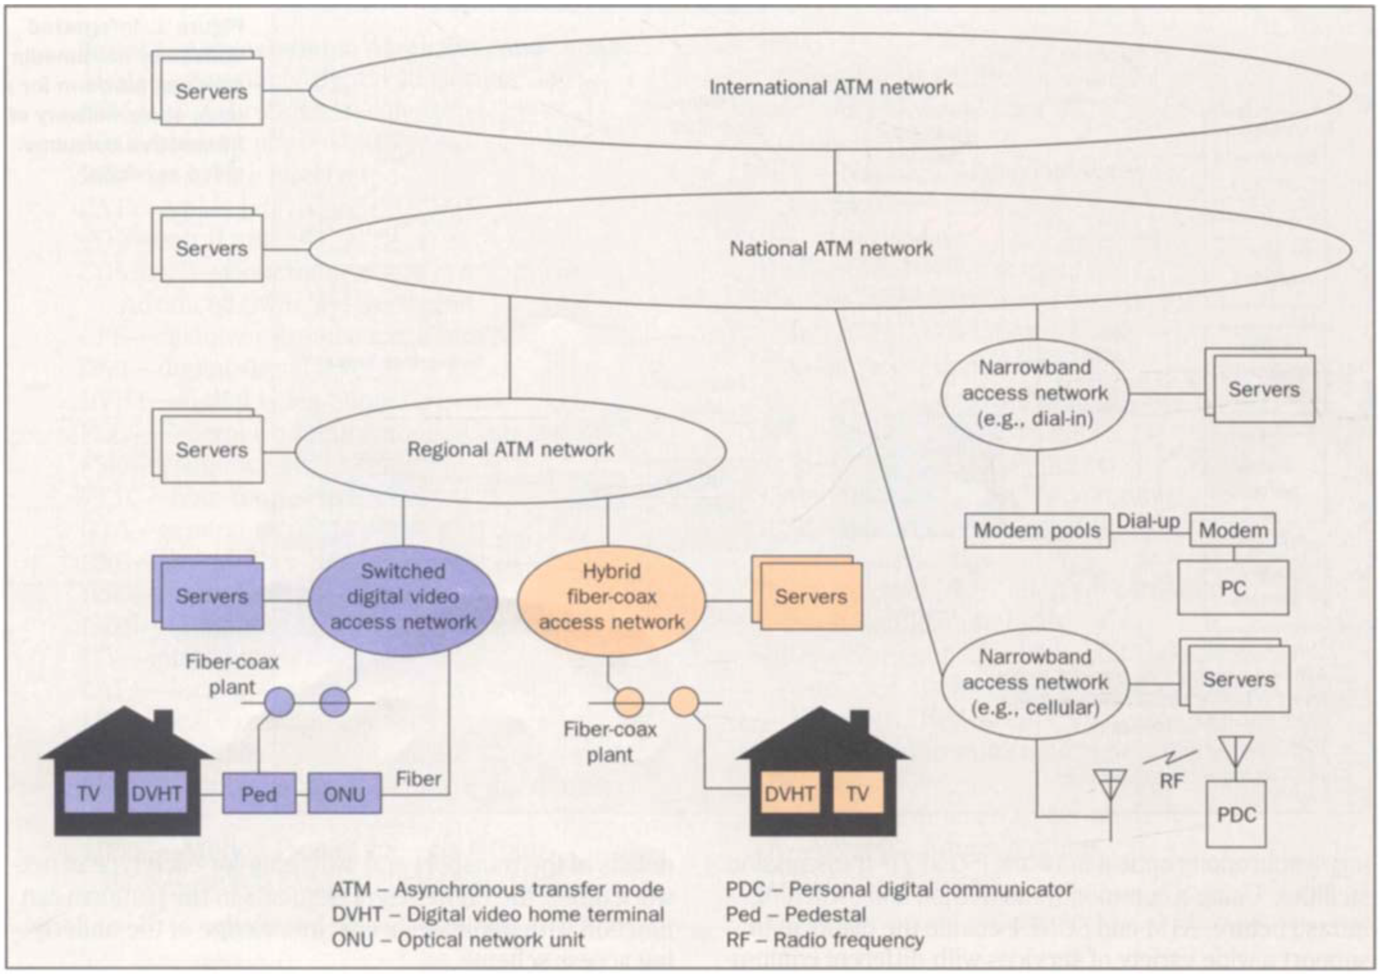

Figure 2. The backbone network structure, composed of a hierarchy of interconnected ATM subnetworks covering various geographical areas.

nals from the residences back to the head-end or CO.

The downstream bandwidth is most often divided into 6-MHz channels that can carry a standard analog TV signal. A 750-MHz HFC network will support approximately 110 of these 6-MHz downstream channels. No uniformly accepted channel plan exists for the upstream spectrum.

The asymmetry of upstream and downstream bandwidth in an HFC network fits well with currently contemplated services. For example, in a movies-on-demand service, the downstream traffic is a relatively high-bandwidth digitally compressed video stream running at a few megabits per second. The reverse traffic would be low bandwidth signaling from the user's remote control for such operations as selecting the film title, pausing the presentation, rewinding, and fast-forward. Future services, such as videotelephony or some client-server data applications, will have a more symmetrical requirement for upstream and downstream bandwidth.

Physically, a hybrid fiber-coax access network uses a combination of analog optical fiber trunks transitioning to coaxial cable for the connections to each residence—hence the name of this access network. About 200 to 1,500 homes are served by an individual fiber trunk and coax distribution plant, often called a *node.*

Each node of an HFC network is a broadcast network. Every residence served by a given trunk has access to all the same signals, although some channels may be jammed by local interdiction units. RF tuners within the CPE of each residence select the 6-MHz downstream channel(s) of interest. Subsequent electronics may then select subchannels from a multiplexed digital data stream transported on the tuner-selected 6-MHz channel. To prevent unauthorized signal reception, these digital subchannels may be encrypted. The keys for each subchannel can be dynamically and securely distributed only to the subscribers who are authorized to receive the signal.

HFC networks can carry a combination of analog and digital signals simultaneously. The 6-MHz downstream channels can be individually allocated to transport either a standard analog television signal or a digital data stream. Data streams are created by digital RF modulators using techniques such as quadrature amplitude modulation (QAM). The upstream spectrum typically carries lower-bandwidth digital signaling channels that use simple frequency and phase modulation techniques such as frequency shift keying (FSK), binary phase shift keying (BPSK), and quadrature phase shift keying (QPSK).

An SDV access network, shown earlier in the lower half of Figure 1, provides a combination of broadcast analog RF signals, conventional POTS telephony, and dedicated digital data feeds for interactive video services to each residence.

Physically, an SDV network consists of fiber trunks from a central office that terminate in small electronic subsystems known as optical network units (ONUs), located in residential neighborhoods. The ONUs are also connected to an HFC-style analog RF distribution plant carrying broadcast TV signals.

Each ONU serves about 16 to 40 homes, supplying both standard twisted-pair POTS telephony drops to the homes and twisted-pair outputs carrying bidirectional digital video signals. These digital loops connect to pedestal units typically serving up to 4 homes. In each pedestal, the digital feeds from the ONU are RF-combined with analog TV program feeds from a conventional fiber-coax distribution plant. A coax drop carries the output of the pedestal to a home, where this broadband drop is connected to the CPE.

In contrast to an HFC network, a switching system outside the home—that is, in the CO—selects the data and programming to be transported on the digital feeds to the CPE. Because all the available digital programming is not being broadcast to every home in a node service area, as it is with HFC, protection from unauthorized service access does not require signal encryption. The choice between HFC and SDV access systems can vary for different service providers, based on a wide range of service, network, and operational factors.

Customer premises equipment. An important issue in providing interactive consumer video services is the terminal device—for example, a television or a personal computer—to be used in the home. Televisions are far more universal than PCs: about 100 million households in the United States have at least one TV set, but only 20 to 30 million homes have PCs.

Unfortunately, televisions are only simple audio and video display devices—they have limited means for interacting with program sources other than changing channels. Interactive services delivered via television require some means of interfacing with the delivery network and for communicating bidirectionally with the program source. Efficient delivery of some services may also require local processing. Consequently, standard televisions supporting interactive applications require adjunct functionality. This is typically provided by a small appliance—commonly known as a set-top box (STB)—that is similar in size and appearance to those widely deployed in traditional cable television (CATV) systems. Interactive STBs are becoming known as digital video home terminals (DVHTs), or similar names, to distinguish them from the basic analog units.

PCs are a poor mode of delivery for many common types of entertainment programming, such as movies, which benefit from larger screens or group viewing. However, PCs can execute local applications and are well suited for games, educational applications, and some forms of communication and commerce. For many services PCs, like televisions, require a means for communicating with remote program sources. These might be as simple as a dial-up modem, or perhaps a connection to a higher-bandwidth network, such as the cable modems used with the HFC systems discussed below.

Because of its broad availability, much of AT&T's work to date on interactive consumer video services has emphasized the use of television as the home terminal device. However, other activities are also under way that will use home PCs as delivery vehicles, particularly for services such as client-server computing applications, which require more symmetric two-way communication.

The cost of new equipment required on customer premises is a prime determinant of the market acceptance of new services. More consumers will try new services if the cost of buying or leasing required CPE—for example, the DVHT—is low.

The interactive consumer video services infrastructure must be designed to support a wide range of CPE capabilities and prices simultaneously. At the low end, a basic DVHT for an HFC network would provide the traditional analog STB capabilities for tuning and addressable descrambling, along with a single channel of Motion Picture Experts Group (MPEG)-2 video decompression, stereo audio decompression, and the appropriate digital security mechanisms. These DVHTs will soon be deployed to support digitally compressed broadcast and enhanced pay-per-view programming. The anticipated price range for these low-end DVHTs is $300 to $400 in 1996.

The services delivery platform must be able to support a wide variety of interactive services, including movies on demand, home shopping, telephony, messaging, and games, even with these simple, low-cost DVHTs. Advanced—that is, more expensive—functionality is moved out of the home and back into the network, where it can be economically shared across a large base of subscribers.

The platform must also support DVHTs that have substantial local processing and video generation capability. For example, prototype DVHTs that contain three channels of MPEG-2 decompression were built for some AT&T trials, with independent real-time scaling and translation of the decompressed video, eight channels of audio decompression, a reduced instruction set computer (RISC) microprocessor, two Personal Computer Memory Card Interface Association (PCMCIA) slots, and several other features.

This design, which represents the high end of DVHT capabilities, was created to enable market testing of the broadest possible mix of services. Over time, this level of functionality could be provided at reasonable cost, but currently these prototype DVHTs cost several thousand dollars each in small quantities and are not likely to be widely deployed outside trial environments.

As telecommuting becomes more widespread, high-bandwidth data networking to the home, which supports PCs, presents another attractive service opportunity for access network providers. Data networking services can be provided either by using a special piece of CPE—sometimes called a cable modem—designed to connect a subscriber's PC directly to the broadband access network, or by additional capabilities (for example, a *10baseT data port*, which is a common form of Ethernet that uses baseband signaling on twisted-pair wiring) provided by the DVHT.

Servers. In an interactive consumer video services delivery platform, servers are collections of computing, storage, and communications equipment that cooperate with the DVHTs to implement the services offered to the subscriber. No one-to-one correspondence exists between servers and services. A service may require more than one server to implement it, or a server may help to implement more than one distinct service.

Not all servers provide services directly to subscribers. Some servers provide services to other servers, such as security, authentication, directories, and billing records collection. The platform, therefore, supports two broad classes of servers: program servers and support servers. Program servers interact directly with DVHTs to provide services directly to subscribers; support servers provide services for other servers (either program or support servers). This partitioning of server types is not absolute—some support servers can interact directly with DVHTs.

When used with simple, low-cost DVHTs, program servers might emulate compressed digital program sources. In this mode, the servers generate fully formed audio and video streams, which only require decompression in the STB. Subscriber remote control button-press signals are carried back to the server without alteration for direct interpretation by the service application.

For more intelligent DVHTs, which have local processing and can manipulate media elements, program servers maintain a client-server relationship. In this case, the server typically sends primitive media elements to the DVHT to be processed and combined under the control of the locally executing service application. Of course, the intelligent DVHTs can also accept the composite program feeds that are sent to basic DVHTs.

Reducing CPE costs by keeping shared functionality in the network can restrict the types of interactive applications that can be effectively supported. Highly interactive applications, such as quick reaction time computer games (often called "twitch games"), require very low latency from the time a consumer hits a button on the remote control until the effect is seen on the TV screen.

With a shared, centralized architecture supporting basic DVHTs, a consumer's remote control input is transmitted from the DVHT back to the server in the network. This input can change the flow of the program and, consequently, the video output for that consumer. The video (and audio) produced by the server is compressed, in real time, for transmission to the consumer's DVHT. The basic DVHT then simply decompresses these transmitted audio/video feeds to create the presentation on the TV set.

The latency in this architecture includes the transmission time of the remote control input from the DVHT to the server, plus the video feed from the server to the DVHT. If the communication bandwidths are sufficiently large to keep these latencies low, the principal component of the total latency becomes the time lag inherent in the real-time video compression in the server and the decompression in the DVHT.

This latency arises because most video compression schemes, such as MPEG, require one or more frame time(s) of delay in the encoding process to compute motion vectors, plus an additional frame time of delay in the decoder. In the NTSC television system, used in North America, a frame time is approximately 1/30 second. In interactive graphics, 100 ms, or about 3 NTSC frame times, is typically regarded as the upper bound on latency for maintaining the consumer's feeling of instant control response. With somewhat more intelligence in the DVHT, the initial response to a remote control input can be generated locally, which permits longer latencies in communication paths and in video compression and decompression processing.

In the interactive consumer video services platform, a special class of servers provides the navigation and other functions required to support subscriber selection of services and/or service providers. These servers present "menus" of providers or available services (such as movies on demand, home shopping mall, and broadcast television), accept the user's selection, and then hand off to the provider or server that supplies the selected service. This hand-off is effected by switching the DVHT ATM virtual path, as described earlier.

The visual appearance and navigation scheme of a gateway server is completely programmable by the server operator. A high-end interactive media server used as a gateway can present a rich, full-motion audio/video interface, or a low-end system may present only simple text-based menus. The gateway servers enable the access network operator to control the top-level "look and feel" of a system and, subject to regulation, to control the packaging and branding of services obtained wholesale from third-party providers.

Gateway servers allow individual users to personalize the on-screen appearance and operation of the system. Subscribers might simply designate their favorite services to be shown first, or they might engage "intelligent agents" to notify them of programs or events of potential interest. This dynamic customization could also be used to provide different views of the system for each individual in a household. For example, when the TV is turned on using a child's remote control, the system could show only those programs or services suitable for children.

Gateway servers are used to implement both Level 1 and Level 2 gateways, as required in the United States by the Federal Communications Commission (FCC) in its video dial-tone (VDT) regulations[5] for the telephone local exchange carriers (LECs). When initiating a session with the platform (that is, turning on the TV set), the subscriber is connected to a Level 1 gateway. Using this gateway, the subscriber chooses a service provider or packager from a list of those available on the system. After choosing, the subscriber would be connected to the provider's or packager's Level 2 gateway, from which he or she could choose specific services. This structure gives all service providers *identically capable* interfaces to the access network .

Bandwidth and session management are implemented as software elements in the gateway servers. Bandwidth management issues switch commands and controls the assignment of ATM virtual paths and the access network upstream and downstream channels. Session management establishes and maintains context and status information about a given subscriber's current activity on the system. These two functions are closely associated with the service selection functions, but they are logically distinct, allowing multiple virtual service providers to share management of the infrastructure.

An example of a gateway server product is the AT&T Video Manager, described elsewhere in this issue.[6]

Sample Service: Movies on Demand. The large existing market for video tape rental illustrates the substantial consumer demand for individual control over the content and scheduling of television programs. Unlike standard broadcast TV, videotape rental enables the consumer to select a program, set the start time, and directly control the presentation using pause, rewind, and fast-forward capabilities on his or her VCR.

A movies-on-demand service would provide these same capabilities without the inconvenience of traveling to the rental store to get the videotape (and possibly finding that it is not available), and then later to return the tape. The service would allow a subscriber not only to choose a movie title from a large library and view the movie immediately on his or her TV set, but also to use the same controls (pause, rewind, fast forward) available when playing a rented tape in a VCR.

Assuming that the delivery network has the required bandwidth and bidirectional communication capabilities discussed previously, the remaining key technical challenge of a large-scale movies-on-demand service is creating and managing the video server systems that can store thousands of movie titles and play any number of them simultaneously to millions of potential customers.

Clearly, large storage capacity is essential for movie-on-demand video servers. A two-hour feature film, compressed to 3 Mb/s (using MPEG-1 or MPEG-2), requires about 3 GB of storage capacity. A server offering 100 titles would need 300 GB of storage. If 1,000 people were simultaneously watching titles stored on this server, the total output bandwidth required would be about 3 Gb/s. The aggregate bandwidth required to serve large user populations is staggering.

There are significant economies of scale in video server design. Because of scale economies in the cost of bandwidth and storage of the underlying memory technologies, a large server is more cost efficient per simultaneous user than a small server.

There is a tradeoff, however, between economies of scale for servers and the cost of transport. Long-haul, high-bandwidth networking is currently very expensive, but it is expected to decline significantly as ATM facilities become widely available. The optimization of server versus transport costs for movies-on-demand service will ultimately depend on subscriber usage patterns, but these will not be reliably known until movies-on-demand service have been more broadly deployed.

The DVHT for a basic movies-on-demand service can be quite simple. All that is required is decompression of MPEG audio and video streams, along with some mechanism for selecting movie titles. The DVHT can connect to an existing analog STB furnished by the local CATV company or, eventually, DVHT functionality could be incorporated directly into the TV set.

Sample Service: Home Shopping. Based on the popularity of the Home Shopping Channel (HSC), QVC, and similar channels on traditional CATV, interactive home shopping is likely to be a widely used service. Interactive home shopping will allow the subscriber to view video catalogs of merchandise, demonstrations, and other related materials—on demand—and immediately purchase the items shown for later delivery. The shopping service might also enable the subscriber to talk on-line to a customer service representative who can answer questions about the merchandise being shown or assist in other ways.

The technology required to deliver a compelling home shopping service is similar to that required for movies on demand. Home shopping catalogs are likely to contain a substantial amount of video, thus requiring large, high-bandwidth servers and significant network bandwidth. Unlike movies, however, the video being delivered will probably not be a long, largely uninterrupted program. Subscribers are expected to interact extensively with the shopping program as it progresses, causing the video stream to change fairly rapidly.

To appear seamless and responsive to the subscriber, the service will require very short latencies throughout the system. Based on other interactive applications, such as video games, this latency should not exceed 100 ms. Some latencies are inherent in the fundamental network protocols, or they are based on distance and cannot be controlled. However, various hardware and software design strategies can reduce latencies in the server.

Like movies on demand, the DVHTs for home shopping can also be simple, provided the server can output a succession of short video clips without noticeable pauses between them. The traffic from the DVHT to the server caused by subscriber controls will be greater than that for movies on demand, and the latency in this signaling channel must be minimized.

Early AT&T interactive television trial. To gain some preliminary insight into the attributes of interactive consumer video services that might be most attractive to consumers, AT&T conducted a trial of basic interactive television (ITV) services from the fall of 1992 through the spring of 1993.

In the Chicago area, AT&T connected commercial ITV equipment to the existing television sets of about 50 of their employees. This ITV equipment was primitive compared to the state of the art for ITV technology; it provided only still frame images with monaural, telephone-quality audio.

The distribution network was a multichannel microwave distribution system (MMDS) wireless cable system. The MMDS transmitter, located on the John Hancock tower in Chicago, broadcast the ITV programming to dish antennas installed on each home in the trial. The servers were located in an AT&T facility in New York City and connected to the transmitting equipment in Chicago by a T1 private line.

The trial offered services such as home shopping, games, information, and electronic mail. The "look and feel" of the services were substantially different than what could be delivered with the latest technology; nonetheless, some important anecdotal insights emerged from this fairly restricted testing.

One interesting insight was the importance of communication as an adjunct to other services. People would buy items through the shopping services and then post messages about the items they had bought. A spontaneous talk-radio-like activity emerged, initially based on the postings of one program developer named Dave; this later evolved to a synthetic personality created and maintained by several people for "Dave."

Another interesting insight was the unusual demographics of the people playing games on the system. Young boys are typically the heaviest users of video games; in this trial, however, the age range of the users was much broader, probably because of the games offered. A dice-like game was found to be very popular with older adults, for example, leading to the conclusion that a substantial market may exist for games that are different from those typically offered with the popular consoles from Nintendo* and Sega.*

From a technical standpoint, the clearest result was the confirmation of the requirement for full-motion video programming rather than still images. Although participants in the trial generally enjoyed the interactivity, they remarked that the interactive programming seemed somewhat dull compared to conventional noninteractive broadcasts. The public's expectations for production quality is set by the best in broadcast television, and anything less, particularly when it is delivered through the same medium, suffers by comparison.

Business Multimedia Services

The service requirements for business multimedia applications push communication network architecture in a different technical direction than consumer mass market multimedia applications do. Businesses want multimedia communications to fit well into their unique business processes, typically to save time or cost, or to improve quality. As a result, the push is in the direction of many communication networks optimized for unique capabilities. Mass markets seek a widely available single network to optimize connectivity and lower cost through volume use of media, equipment, and multimedia network traffic.

COMPASS Trials. AT&T Network Systems joined U S West in a series of advanced business multimedia application trials called Communications Programs for Advanced Switched Services (COMPASS).[7-10] The trials explored the requirements and technical characteristics of leading-edge multimedia business users based on state-of-the-art and experimental multimedia network technology. From late 1990 to mid-1994, the COMPASS effort conducted trials of five different advanced networking technologies with many different end-user applications.

To get a better understanding of the technical requirements for future interactive multimedia network design, it is useful to consider what was learned from the COMPASS effort. The sections that follow discuss the three COMPASS trials that had the strongest focus on multimedia applications. The other two trials focused primarily on the switched data networking technologies of switched T1 with prototype 5ESS® switching equipment and DACS IV, and Switched Multi-megabit Data

Service (SMDS*) with a BNS-2000 ATM switch.

Goals. The goals of COMPASS trials were to test, with real business end users, the best available state-of-the-art or prototype narrowband, wideband, or broadband switched multimedia network technology then available. The tests ended with the then new (in 1990) broadband integrated services digital network (ISDN) 150-Mb/s ATM line interface network. The team sought to:

- Assess future market needs and identify new applications;
- Assess new and emerging network technologies; and
- Understand how to move from the 1990 telephone network to tomorrow's narrowband, wideband, and broadband networks, as existing and future market demands and opportunities continued to move forward.

Multimedia data networking via narrowband ISDN. The first COMPASS trial focused on various aspects of narrowband ISDN (n-ISDN) data networking over three trial projects. The first project demonstrated the benefits of electronic document imaging over a wide area network.

Ten students from the University of Minnesota's Carlson School of Management were linked to a host computer at the university. At home, students used a single ISDN basic rate interface (BRI) line for both high-speed switched data and digital voice services. The students' personal computers were connected to the host by a B-channel packet-switched data service, accessed through n-ISDN data terminal adapters.

This arrangement allowed the graduate students to look up on-line text and scanned images and to perform collaborative computing remotely. The network supported multiple simultaneous application sessions with the host, downloading large full-screen images within 12 seconds. With their former modem connections, these tasks would have been impossible or impractical to perform over the analog voice telephone network.

The second trial bundled two high-speed B-channels used in distance learning applications. Users at the University of Minnesota's Department of Rhetoric tested PC-based video conferencing and multimedia document-sharing applications over a 128-kb/s ISDN circuit-switched network. This allowed flexible high-speed access to classroom resources from individual homes—a critical benefit to physically challenged students.

The third trial deployed real estate imaging and Multiple Listing Service (MLS) information over an ISDN B-channel to remote realtors. This application not only replaced conventional paper MLS books, but also provided on-line photos of perspective homes and neighborhoods in full color. The network was also able to offer more up-to-date listings and on-line access to more detailed information than was possible and practical with conventional paper MLS books.

From these trials, the following observations and conclusions were made:

- A switched public network that supports common standards and interfaces is ideal for these applications. These needs were met quite well by n-ISDN, especially with national n-ISDN standards.
- Ubiquity of service is essential for the wider applicability and success of these offerings. The usefulness of advanced multimedia communications increases as connectivity increases, making this lack of ubiquity a major industry problem.
- As with any new technology or service offering, considerable technical, provisioning, marketing, and customer support of the service is essential. Human resistance to change and "technophobia" are real, serious problems for multimedia applications.
- The n-ISDN equipment was operationally successful, with considerable end-user acceptance.
- Despite the common view to the contrary, n-ISDN has many applications.
- Loop plant requires individual verification of n-ISDN compatibility.

Executive desktop videophone. Another COMPASS trial focused on an application for switched broadband services. Senior U S West managers, located in five states, tested a prototype switched videophone service using a special overlay prototype network. In phase 1, a subset of executives were linked via 45-Mb/s inter-LATA (local access and transport area) and intra-LATA service. In phase 2, the entire user group tested the same system, with 384-kb/s video/audio compression and T1 facilities for the network. Calls were routed by prototype synchronous transfer mode (STM) switches, which directed traffic on a dial-up basis in much the same way that today's telephone network operates.

The AT&T prototype videophone terminals provided users with eye-to-eye contact by optically folding their images to an internal camera via half-silvered mirrors. Video calls were established using n-ISDN phone sets. This "friendly" human interface allowed users to establish two-party and multiparty conference calls and to speed-dial in the same way users make voice calls today. The multiparty videoconferencing capability supported split screen viewing of conferees via a central bridge, collocated with each switch.

The following observations and conclusions were made from this trial:

- Executive user acceptance was very good for the high-quality broadband system because a "telepresence" was achieved. The eye-to-eye capability was considered very good by the executives.
- The executives considered videophone terminals and multimedia personal computer terminals to be two different devices. In the executives' minds, these two types of terminals supported distinct applications, although from a technical viewpoint, the two may share some overlapping features and functionality. These trial users indicated that the executive desk of the future is likely to have a videophone and a multimedia personal computer with network access.
- The videophone is viewed as an evolution of the telephone and needs the features popular in modern office telephones.
- Voice quality must be very good, at least as good as modern full-duplex speakerphones.
- End-to-end delay in two-way multimedia communications must be addressed to achieve a quality video conferencing system. There is an extra burden on the multimedia communications network to minimize delay and synchronize different media types. This is a key area in which two-way networks differ from one-way, broadcast-oriented entertainment or distance learning applications. In these networks, delay is mainly observed at the start of the broadcast and the media are easily synchronized going forward. For two-way switched networks, the delay management is more complex.
- In this COMPASS trial, network delay was introduced by video compression and at digital source and sink points on interstate circuit-switched connections. This caused unexpected quality problems that precipitated reengineering. In video conferencing, compression used for the video signal causes extra delay in the two-way video. Buffering the audio signal to achieve lip-sync compensated for the video processing delay; however, the buffered signal impaired the perceived quality.
- In the COMPASS trial, the 45-Mb/s service was viewed as a desirable quality level, but the 384-kb/s service was not. This problem can be resolved with more powerful processors, better compression algorithms, and/or increased network bandwidth. In real network operations, multimedia end-to-end delay requirements add a new dimension to network configuration and design. This forces multimedia communications network engineers to think beyond the organizational line that divides network switching and transport from applications.
- Reliability and availability are essential—success breeds reliance on the system.
- Proper lighting is important. Keeping adequate front lighting on subjects was a challenge. Users of a quality videophone want to look good to the people on the other end.
- The application must be easy to use with little or no training. A telephone-like call setup was acceptable.
- Videotelephony will require new video technical skills and test equipment from telephone company operations people who install and support the service.
- Business videotelephony needs to support document imaging to allow material such as viewgraphs to be viewed at the other end.
- Widespread connectivity of the service is important to its usefulness. Executive videophone was primarily seen as a way to reduce travel and save time. It should be available wherever the executives normally travel—particularly for internal corporate meetings. (U S West Communications operates in 14 states and 3 time zones.) The transport price is a major obstacle to widespread deployment. The high cost of the interstate digital signal level 3 (DS-3) was the reason for terminating the 45-Mb/s phase of the trial.
- Excellent quality switched videophone communication can be provided now and is highly desirable.

Broadband ISDN ATM multimedia networking. The COMPASS Broadband ISDN ATM Multimedia Networking project consisted of a high-definition distance learning trial among teaching hospitals in the Minneapolis and St. Paul area. It included the University of Minnesota, Hennipen County Medical Center, and the Veterans Administration (VA) hospital.

The high-definition distance learning application involved a classroom at the University of Minnesota that was specially equipped to provide a prototype multimedia originating lecture hall of the future. The trial also had classrooms at the Hennipin County Medical Center and the VA Hospital that could receive the material from the University and from the other remote classroom. Each classroom could also provide video and voice from these students to the lecturer and students at the University and the other remote classroom.

Typically, an expert in radiology lectured resident physicians located in the classroom at the University of Minnesota, and in the Hennipin County Medical Center and the VA hospital, which are located several miles from the University Hospital. These classes were part of the residents' training. The distance learning application was conceived to save time for the students by reducing travel.

The University classroom equipment could digitize image and text media; display several medical images, viewgraph material, and/or video in high resolution; and show the lecturer and all class students to all sites simultaneously. The classrooms at the other hospitals had the equipment to see all the originating lecture hall images and sound, as well as the other remote classrooms. Each classroom could pose a question or comment to the lecturer, accompanied by video of the questioner.

All connectivity was via broadband ISDN 150-Mb/s switched access. The AT&T prototype ATM switch was located in a U S West central office several miles from the University of Minnesota. Each location accessed this overlay ATM trial network via AT&T prototype general terminal adapters (GTAs), which supported a number of inputs (such as DS-3, BRI, and 802.6), converted them to ATM cell format, and provided address information. The remote GTAs, located at the customers' premises, connected to the CO-based AT&T prototype ATM switch platform through dedicated optical carrier (OC)-3c 155-Mb/s optical links.

This trial resulted in the following observations and conclusions:

- Expert physicians and students were very willing to use broadband distance learning to reduce travel. The leaders of the training program felt this technology would enable the residents to get additional training they would otherwise miss because of time constraints.
- Human factors issues are very important in the classroom setting. For example, human factors issues that caused trouble included lighting, camera tracking, teacher command console controls, audio quality from speakerphones, media format, and lead time to digitize media.
- For this teaching environment, 45 Mb/s provides good, high-resolution image quality. A well-known expert in breast cancer said the quality of the images on the trial monitors was better than that available in normal practice.
- The prototype 150-Mb/s ATM GTAs and switch were very reliable, and regularly scheduled classes were routinely conducted via the network.
- Little CPE exists to support high-end networking applications of this type, and the cost is high.
- Development of broadband multimedia applications requires significant expertise, time, and money.
- The broadband ISDN ATM technology handles a dynamic mix of traffic well and offers great flexibility. However, capacity engineering of links is complex and becomes an increasing problem as the load on the link increases.
- The availability of broadband access in today's LEC networks is limited and requires considerable time and expense to obtain.
- Some interoperability tests were run with ATM equipment from other suppliers. All connections were incompatible to some degree. Standards are still being developed.
- The expectations of multimedia users increase with usage of the technology. In the beginning of the semester everybody seemed happy with the system. As the months progressed, irritation increased with those parts of the system that did not work well, such as speakerphone sound quality and time to digitize images.

Conclusions

We are in the very early stages of a revolution in multimedia communication precipitated by advances in microelectronics, photonics, signal processing, and high-quality displays. AT&T has conducted early trials with consumers and business to understand how these technologies might be used and the impact they may have on future products and services. Several themes recurred throughout these projects.

Effective human factors design of the interface and the perceived performance of the system are very important to end-user acceptance. For example, the consumer video controller and screen display must be perceived as easy to use and understand. Synchronization between lip movements and speech sounds is just as important as video quality in a videophone application.

The future multimedia network must be planned to meet the needs of consumer mass market applications, as well as the needs of mainstream business multimedia applications. This paper highlighted trials of four different networks—a consumer video network and three business multimedia networks. Each multimedia network was optimized to some degree to serve the trial applications. In addition, these four networks are considerably different from the main-line networks of today, such as the voice telephone network, private line networks, worldwide television distribution, commercial packet networks, and the Internet. A common multimedia services infrastructure needs to evolve to achieve the required economies of scale.

In actual deployments, many people will not be interested in new services. Designing applications and equipment that are broadly compelling, along with careful management of change, are significant problems.

The basic technology exists for a wide array of multimedia networked applications. Although costs are substantial, the trends in the basic technology are favorable.

Acknowledgments

Many people, both outside and inside AT&T, contributed to the work reported in this paper. Although they are too numerous to mention individually, their efforts are recognized and appreciated.

***Trademarks**

Nintendo is a registered trademark of Nintendo of America Inc.
Sega is a registered trademark of Sega of America Inc.
SMDS is a service mark of Bell Communications Research.

References

1. M. de Prycker, *Asynchronous Transfer Mode*, Ellis Horwood, New York, 1993.
2. M. Sexton and A. Reid, *Transmission Networking: SONET and the Synchronous Digital Hierarchy*, Artech House, Boston, 1992.
3. W. Y. Chen and D. L. Waring, "Applicability of ADSL to Support Video Dial Tone in the Copper Loop," *IEEE Communications Magazine*, May 1994, pp. 102-109.
4. J. R. Jones, "Baseband and Passband Transport Systems for Interactive Video Services," *IEEE Communications Magazine*, May 1994, pp. 90-101.
5. C. H. Kennedy, *An Introduction to U.S. Telecommunications Law*, Artech House, Boston, 1994.
6. R. E. Libman, M. T. Midani, I. J. Morgan, and H. T. Nguyen, "The Interactive Video Network: An Overview of the Video Manager and the V Protocol," *AT&T Technical Journal*, Vol. 74, No. 5, September/October 1995, pp. 92-105.
7. L. J. Haas, P. K. Heldman, and J. L. Strickland, "Desktop Video: Market Finally Catches Up to Technology," *Telephone Engineer and Manager*, April 12, 1993.
8. Peter Heldman and Jack Warner, "COMPASS Points the Way to ISDN's Future," *America's Network*, Vol. 98, Nos. 19 and 20, October 1 and 15, 1994, pp. 30 and 48.
9. T. R. Nanke and B. G. Prieve, "The COMPASS Broadband Trials," *AT&T Technical Journal*, Vol. 72, No. 6, November/December 1993, pp. 67-75.
10. Barton G. Prieve, Trent R. Nanke, and William R. Walling, "Broadband ISDN/ATM Application Trials: What Did We Learn from COMPASS?," *Proceedings of the International Switching Symposium '95*, Berlin, Germany.

(Manuscript approved August 1995)

Edward S. Szurkowski *is director of the Interactive Services Delivery Laboratory at AT&T Network Systems in Murray Hill, New Jersey. He is responsible for research and development of systems that deliver interactive multimedia services to consumers over broadband networks. Mr. Szurkowski received B.E.E., M.E.E., and Ph.D. degrees in electrical engineering from the University of Delaware, Newark. He joined AT&T in 1981.*

Jack C. Warner *is director of the Denver Field Office of AT&T Bell Laboratories in Denver, Colorado. His work with U S West focuses on multimedia communications, personal communications services, and telephone network evolution. Mr. Warner holds a B.S. in industrial engineering and operations research and an M.S. in operations research from Cornell University, Ithaca, New York, as well as an M.B.A. in business economics from the University of Michigan in Ann Arbor. He joined AT&T in 1969.*